

Dynamic Multipriority Patient Scheduling for a Diagnostic Resource

Jonathan Patrick

Telfer School of Management, University of Ottawa, Ottawa, Ontario, Canada K1N 6N5, patrick@telfer.uottawa.ca

Martin L. Puterman, Maurice Queyranne

Sauder School of Business, University of British Columbia, Vancouver, British Columbia, Canada V6T 1Z2
{martin.puterman@sauder.ubc.ca, maurice.queyranne@sauder.ubc.ca}

We present a method to dynamically schedule patients with different priorities to a diagnostic facility in a public health-care setting. Rather than maximizing revenue, the challenge facing the resource manager is to dynamically allocate available capacity to incoming demand to achieve wait-time targets in a cost-effective manner. We model the scheduling process as a Markov decision process. Because the state space is too large for a direct solution, we solve the equivalent linear program through approximate dynamic programming. For a broad range of cost parameter values, we present analytical results that give the form of the optimal linear value function approximation and the resulting policy. We investigate the practical implications and the quality of the policy through simulation.

Subject classifications: health care; approximate dynamic programming; Markov decision processes; patient scheduling; linear programming.

Area of review: Special Issue on Operations Research in Health Care.

History: Received November 2006; revisions received November 2007, March 2008; accepted April 2008.

1. Introduction

Globally, public health systems face increasing and lengthy wait times for a wide range of medical services. Although in some cases these waits may have little medical impact, in others, excessive wait times can potentially impact health outcomes (Sanmartin 2004). Thus, health-care managers and policymakers face considerable political and community pressure to better manage health-care resources in order to reduce wait times to acceptable levels without undue additional costs. One key lever for effective management is through improved patient scheduling—particularly when patients may be classified into priority categories with different medically acceptable wait times. For example, some conditions may require urgent immediate treatment, whereas in other cases it may be medically acceptable to wait up to several weeks. Because less-urgent patients are booked further into the future, this raises the question as to how much resource capacity to reserve for later-arriving but higher-priority demand? Whereas this paper focuses on scheduling diagnostic imaging resources, our methods and results apply more broadly.


Demand for a diagnostic resource (such as a computed tomography (CT) scanner) arises from multiple sources. Within the hospital, demand arrives either from the emergency department or from the wards. In both cases, requests are given varying degrees of priority, ranging from “immediate” to “within 24 hours.” The resource manager of the diagnostic facility will generally have no prior knowledge of the extent of emergency (EP) and inpatient (IP) demand

to expect. As Figure 1 illustrates, this demand can vary significantly from day to day. In addition, most hospitals also serve a significant outpatient (OP) population. In the hospital setting we studied, outpatient demand arrived in the form of faxed requisitions from specialists. These were accumulated and sent to a staff radiologist in batches for priority classification. In British Columbia, there exist three OP priority classes with allowable wait times of 7, 14, and 28 days, respectively. These targets were determined by a panel of experts in collaboration with the BC government. A booking clerk, who we refer to as a scheduler, collects the prioritized requests and assigns future appointments to each one.

The daily challenge facing the scheduler is to allocate the available capacity between the priority classes so as to minimize the number of patients whose wait time exceeds a prespecified, priority-specific target, with greater weight given to any late bookings of higher-priority demand. This requires significant foresight because each day’s decision will clearly impact what appointment slots are available for future demand. If lower-priority patients are booked too soon, then there may be insufficient capacity for later-arriving higher-priority demand. Conversely, if lower-priority patients are booked too far into the future, there is the potential for idle capacity.

This research is motivated by a study that a team (including the authors Patrick and Puterman) from the Center for Operations Excellence (COE) at the University of British Columbia carried out in collaboration with the Vancouver

Figure 1. Day-to-day variation in the number of requests for CT scans at a Vancouver hospital.



Coastal Health Authority (VCHA). VCHA management were concerned that OP wait times for CT scans were excessive. They arranged for the COE team to determine the extent of the problem and to suggest methods for improving throughput. Our analysis revealed that over a specific period, a significant proportion of the scheduled appointments for outpatients exceeded medically appropriate wait-time targets; the wait times of half of the highest-priority class, two-thirds of the second-priority class, and three-quarters of the lowest-priority class exceeded the targets. Although our initial recommendations focused on operations and system use issues such as increasing the efficiency of the porter system (Odegaard et al. 2007) and improving the scheduling of diagnostic imaging technologists, it was clear that the VCHA also faced a significant scheduling challenge. Current practice relies entirely on the expertise of the booking clerk, who has no computer system or clear procedures supporting this complex patient-scheduling challenge. Thus, we undertook to develop a more systematic approach to patient scheduling, which is described in depth here. A related nontechnical paper (Patrick and Puterman 2008) communicates our results and other observations regarding wait times to health-care managers.

1.1. Related Literature

The allocation of medical capacity in the presence of multiple patient classes has received limited attention. Comprehensive reviews of the broader appointment scheduling literature include Magerlein and Martin (1978), Cayirli and Veral (2003), Denton and Gupta (2003), and Mondschein and Weintraub (2003). In their review of surgical scheduling, Magerlein and Martin classify scheduling systems into those that schedule patients in advance of the service date, referred to as “advance scheduling,” and those that schedule available patients on the day of service, referred to as “allocation scheduling.” Our work and those we survey below fall into the first stream of “advance scheduling.” An example of allocation scheduling is the work of Green et al. (2006), who analyze the within-day scheduling of patients

to a diagnostic facility when a fixed number of outpatient scans have already been booked. Specifically, they seek to determine which patient to serve next when both inpatients and outpatients are waiting for scans.

Kolesar (1970) proposed the use of Markov decision processes for hospital admission scheduling. He formulates several models that are closely related to that considered in this paper, especially one for “scheduling reservations over a planning horizon.” However, he neither solves nor analyzes the model, but notes that “for admissions planning models that the writer envisions treating, the linear programs would be of a size that can be handled by contemporary computing capabilities.” Clearly he was not envisioning solving problems of the magnitude considered in this paper. Subsequently, Collart and Haurie (1976) develop a semi-Markov population demand model for emergency and elective patients and formulate an optimal stochastic control problem to determine an admission policy that minimizes long-run average costs. Noting that the “computation of a closed-loop solution appears to be a practically insurmountable task,” they propose an open-loop suboptimal control policy that they evaluate through simulation. Rising et al. (1973) present a case study of simulation models designed to test decision policies for a scheduling challenge with two customer classes—walk-ins and advanced appointments—for an outpatient clinic. The focus is on the impact of various decision policies on physician utilization and patient throughput.

More recently, Gerchak et al. (1996) determine the optimal number of elective patients, when facing both elective and emergency demand, to accept each day to a surgical department. They demonstrate that the optimal policy for maximizing revenue is not a strict booking limit policy, but one where the number of elective surgeries accepted increases in conjunction with the number waiting. Our paper differs in a number of respects. Most importantly, we consider an arbitrary number of priority classes rather than two. Second, although a cost is associated with each day of delay in an elective patient’s surgery, Gerchak et al.’s, model does not quantify the actual wait times for these patients, and thus does not account for multiple elective patient priority classes. Because our model includes several lower-priority classes, it requires different late booking penalty functions for each class. Our model explicitly allows for each priority class to have a viable booking window with class-specific costs for late booking.

Gupta and Wang (2008) consider the effect of patient choice on scheduling in a primary-care clinic where patients may have preference for physician and date of service. Patients are divided into those that request same-day service and those that seek an advanced appointment. Although a penalty function is included to penalize the clinic if it cannot meet the request of a patient, the model is not designed to track patient wait times.

Extensive work has been done in revenue management—particularly in the airline industry—on capacity allocation

in the presence of multiple fare classes (for examples, see Bertsimas and Popescu 2003, Brumelle and Walczak 2003, and van Ryzin and Vulcano 2008). Although helpful in our analysis, airline revenue management demonstrates some significant differences from patient scheduling. Airline revenue management has the advantage of concentrating on a small number of flights over a finite horizon. In diagnostic imaging, each potential booking day could be viewed as a flight, and although the booking horizon is finite, it is also continuously evolving, leading to an infinite-horizon problem. Moreover, passengers for a flight can choose which “priority” class to enter, whereas in diagnostic scheduling, their priority class is a function of the urgency for a scan. Finally, airline revenue management does not consider the impact of a given policy on passenger wait times.

An interesting alternative application of scheduling with multiple customer classes is presented by Bertsimas and Shioda (2003). Their work focuses on the seating of customers at a restaurant based on the size of the group and the presence of reservations. They seek to maximize revenue while controlling for customer wait time and ensuring equity.

1.2. Paper Structure

This paper proceeds as follows. We formulate the scheduling problem as a discounted infinite-horizon Markov decision process (MDP) and transform it into the equivalent linear program (LP) that, if solvable, would return the optimal value function for the MDP. However, neither the MDP nor the LP are tractable due to the size of the state space. Therefore, we use approximate dynamic programming (ADP) techniques to produce an approximate linear program (ALP) that has a manageable number of variables (although an unmanageable number of constraints). We solve the ALP through column generation on the dual to derive an estimate of the value function in the MDP. Using this approximate value function, we derive a booking policy that we test through simulation. We also present the surprising result that, under certain very reasonable conditions on the cost structure, we can determine the optimal linear approximation and the consequent policy *without having to solve the ALP*. We then discuss a fundamental unresolved issue within ADP theory—that of producing useful bounds on the “cost” associated with using an approximate value function. We conclude with potential extensions of the model and policy insights.

It could be argued that an average reward MDP would be more appropriate because the objectives are nonmonetary and the future should not be valued less than the present. We instead use a discounted model with a discount factor very close to one because it best reflects the medium-term planning horizon that is most often applicable in the hospital setting. By discounting only slightly, we insure that, over the short term, costs are relatively similar, but that far distant costs are less valued. The changing nature of both supply and demand within health care,

we would argue, makes this a reasonable model. Moreover, the discount model is tractable (in the approximate setting), whereas the average reward model is generally multiclass and requires new ADP methods.

2. A Markov Decision Process Model for the Scheduling Problem

This section formulates a discounted infinite-horizon MDP model by providing the decision epochs, state space, action sets, transition probabilities, and costs.

2.1. Decision Epochs and the Booking Horizon

We consider a system that has the capacity to perform C_1 fixed-length procedures each day. At a specific point of time in a day, referred to as the decision epoch, the scheduler observes the number of booked procedures on each future day over an N -day booking horizon and the number of cases in each priority class to be scheduled. The booking horizon consists of the maximum number of days in advance that hospital management will allow patients to be scheduled. In practice, this is usually not specified; however, we find that the length of the booking horizon is of little consequence because the policy that emerges from the model is independent of N provided that N exceeds the wait-time target of the lowest-priority class.

As mentioned in the introduction, demand arises from two sources, inpatients and outpatients. In practice, most inpatient demand is known at the beginning of each day once morning rounds have been completed on the wards. Outpatient demand arrives throughout the day, and thus is not completely known and prioritized until the end of the day. Because the scheduler will give preference to inpatients over outpatients regardless, we assume that all decisions are made once inpatient demand has been determined. Consequently, outpatient demand is never booked into day 1 (for any scenario involving inpatients and outpatients). Thus, we assume that decision epochs correspond to the beginning of each day.

Our model is complicated by the fact that the horizon is not static, but rolling. Thus, day n at the current decision epoch becomes day $n - 1$ at the subsequent decision epoch. Because no patient is scheduled more than N days in advance, at the beginning of each decision epoch, the N th day has no appointments booked.

2.2. The State Space

A typical state takes the form

$$\mathbf{s} = (\vec{x}, \vec{y}) = (x_1, x_2, \dots, x_N; y_1, y_2, \dots, y_I),$$

where x_n is the number of patients already booked on day n , y_i is the number of priority i patients waiting to be booked, and I is the number of priority classes. The state space, S , is therefore

$$S = \{(\vec{x}, \vec{y}) \mid x_n \leq C_1, 1 \leq n \leq N; \\ 0 \leq y_i \leq Q_i, 1 \leq i \leq I; (\vec{x}, \vec{y}) \in \mathbb{Z}_N \times \mathbb{Z}_I\},$$

where C_1 is the daily base capacity expressed in terms of the number of fixed-length procedures that can be performed each day and Q_i is the maximum number of priority i arrivals in a given day. (Truncating arriving demand is necessary to keep the state space finite, but the maximum number of arrivals can be set sufficiently high as to be of little practical significance.) We assume that each patient requires one appointment slot and that all appointment slots are of equal length. In our setting, the procedures required either 15, 30, 45, or 60 minutes. Because all slots were multiples of 15 minutes, it is not unreasonable to convert demand into 15-minute slots, although to be more realistic one should then consider batch arrivals. Simulation results suggest that the impact of multiple appointment lengths is minimal.

2.3. The Action Set

The scheduler's task is to decide at each decision epoch which available appointment slots to assign to each unit of waiting demand. However, if this were the only action available, then s/he would have very little recourse should base capacity prove insufficient for the realized demand. Thus, we assume the resource manager has the ability to "divert" patients to an alternative capacity source at an additional cost. This is often referred to as "surge" capacity (see Patrick and Puterman 2008). Surge capacity may be in the form of overtime or outsourcing. Alternatively, the scheduler may postpone scheduling to the next day or even reject some demand. Although the ethical implications of this last alternative would clearly depend on the availability of alternative services, it is not without precedent. In New Zealand, for example, a system has been implemented where a level of priority is prespecified for which the hospital can reasonably guarantee a wait time below a certain target level. All other demand is returned to the referring physician as insufficiently urgent to be booked at this time (MacCormick and Parry 2003).

In Vancouver, most hospitals function with limited overtime availability. If necessary, hospitals within the same health authority and even across health authorities may act as an additional source of surge capacity. To be realistic, therefore, we impose a limit on the number of patients who can be diverted per day. Thus, a vector of possible actions can be written as $(\vec{a}, \vec{z}) = \{a_{in}, z_i\}$, where a_{in} is the number of priority i patients to book on day n and z_i is the number of diverted priority i patients. To be valid, any action must satisfy the following constraints, insuring that the base capacity is not exceeded:

$$x_n + \sum_{i=1}^I a_{in} \leq C_1 \quad \forall n \in \{1, \dots, N\}, \quad (1)$$

that no more than C_2 patients are diverted,

$$\sum_{i=1}^I z_i \leq C_2, \quad (2)$$

that the number of bookings and diversions does not exceed the number waiting,

$$\sum_{n=1}^N a_{in} + z_i \leq y_i \quad \forall i \in \{1, \dots, I\}, \quad (3)$$

and that all actions are positive and integer,

$$(\vec{a}, \vec{z}) \in \mathbb{Z}_{IN} \times \mathbb{Z}_I. \quad (4)$$

We denote the action set, A_s , for any given state, $\mathbf{s} = (\vec{x}, \vec{y})$, as the set of actions, (\vec{a}, \vec{z}) , satisfying Equations (1) to (4).

2.4. Transition Probabilities

Once a decision is made, the only stochastic element in the transition to the next state consists of the number of new arrivals in each priority class. Demand that was not booked nor diverted also reappears in the next day's demand. If the number of new arrivals is represented by \vec{y}' , then the state transition,

$$\begin{aligned} & (x_1, x_2, \dots, x_N; y_1, y_2, \dots, y_I) \\ & \rightarrow \left(x_2 + \sum_{i=1}^I a_{i2}, \dots, x_N + \sum_{i=1}^I a_{iN}, 0; \right. \\ & \quad \left. y'_1 + y_1 - \sum_{n=1}^N a_{1n} - z_1, \dots, y'_I + y_I - \sum_{n=1}^N a_{In} - z_I \right), \end{aligned}$$

occurs with probability $p(\vec{y}') = \prod_{i=1}^I p(y'_i)$, where $p(y'_i)$ is the probability that y'_i priority i patients arrive on a given day. We assume that demand for each priority class is independent and that each day's demand is independent as well. Because demand arises from multiple independent sources (the hospital wards and the specialists in the region serviced by the hospital), independence between classes seems a reasonable assumption.

In practice, demand may be seasonal, but for the sake of tractability, we have chosen not to incorporate this into the model. This is not out of line with the literature as seasonality is not considered in any of the models referred to in the literature review. If seasonal patterns are significant, the model can be resolved with different demand patterns to determine the appropriate policy for each season of the year. Surprisingly, the optimal policy is extremely robust to changes in the specific data and thus re-solving may be unnecessary.

2.5. Costs

The cost associated with a given state-action pair derives from three sources: a cost associated with booking a patient beyond the priority-specific wait-time target, a cost associated with using surge capacity, and a cost associated with demand that was neither booked nor diverted. We write the costs as

$$\begin{aligned} & c(\vec{a}, \vec{z}) \\ & = \sum_{i,n} b(i, n) a_{i,n} + \sum_{i=1}^I d(i) z_i + \sum_{i=1}^I f(i) \left(y_i - \sum_{n=1}^N a_{in} - z_i \right), \end{aligned}$$

where $b(i, n)$ is the cost of booking a priority i patient on day n , $d(i)$ is the penalty for diverting a priority i patient, and $f(i)$ is the cost associated with delaying a priority i patient's booking. We represent the wait-time target for class i by $T(i)$. The choice of $b(i, n)$, although arbitrary, should include certain characteristics. It is clearly reasonable to assume that it will be decreasing in i and that $b(i, n)$ should be zero if $n < T(i)$. Furthermore, it would seem advisable to insure that the cost of delaying a patient's booking k days and then booking him/her within the target should be equal to the cost of booking the patient k days late initially. Thus, a natural form for the booking cost is

$$b(i, n) = \begin{cases} \sum_{k=1}^{n-T(i)} \gamma^{k-1} f(i) & \text{for all } n > T(i), \\ 0, & \text{otherwise,} \end{cases}$$

where $f(i)$ is a decreasing function of i . There is certainly an argument to be made for a booking cost function that increases at a faster rate in n . We have experimented with such a cost function and discovered no difference in terms of the policy dictated by the model. Even with the above booking cost function, the policy (for all reasonable values of $f(i)$ and $d(i)$) only books a patient late as a last resort. Causing the booking cost to increase at an even faster rate only further strengthens this policy conclusion. In fact, the analytical results given later provide minimal conditions on $b(i, n)$ that include any function that increases at a faster than linear rate.

The cost function explicitly balances the cost to the patient in wait time and the cost to the system in having to resort to surge capacity. The scheduler's role is to maintain reasonable wait times in a cost-effective manner. The specific value to assign to $f(i)$ is difficult to determine due to the nebulous nature of the cost of booking a patient later than the wait-time target. Determining these costs would be the role of the panel of medical experts who determined the wait time targets. Of particular difficulty is the relationship between the late penalty for each priority class and the diversion costs. The diversion cost is also potentially challenging to quantify and will clearly depend on the available source of surge capacity. The most obvious source is overtime, in which case there exists a specific overtime cost that is independent of the priority class. However, it may be more difficult to determine the cost for other sources of surge capacity. Fortunately, we show that for reasonable choices of $d(i)$ and $f(i)$, the derived policy is very robust to changes in these cost parameters so that the arbitrary nature of their specific values is of less concern.

2.6. The Bellman Equation

The value function v of the MDP specifies the minimum discounted cost over the infinite horizon for each state

and satisfies the following optimality equations for all $(\vec{x}, \vec{y}) \in S$:

$$v(\vec{x}, \vec{y}) = \min_{(\vec{a}, \vec{z}) \in A_{\vec{x}, \vec{y}}} \left\{ c(\vec{a}, \vec{z}) + \gamma \sum_{\vec{y}' \in D} p(\vec{y}') v \left(x_2 + \sum_{i=1}^I a_{i2}, \dots, x_N + \sum_{i=1}^I a_{iN}, 0; y'_1 + y_1 - \sum_{n=1}^N a_{1n} - z_1, \dots, y'_I + y_I - \sum_{n=1}^N a_{In} - z_I \right) \right\}, \quad (5)$$

where γ is the daily discount factor and D is the set of all possible incoming demand streams. It is here that "the curse of dimensionality" becomes apparent. In particular, the dimension of the state space is $C_1^N \prod_{i=1}^I Q_i$. Reasonable values of C_1 , N , I , and Q lead to very high dimensions, making a direct solution impossible.

3. Approximate Dynamic Programming

Over the past few decades, research in approximate dynamic programming has focused on developing methods for addressing the curse of dimensionality. These methods restrict the value function to lie within a specified class of functions and then seek to find the optimal value function within this class. Challenges include determining the best class of functions to use for a given problem, determining the optimal approximation within a chosen class of functions and bounding the gap between the cost of the policy determined by the approximate solution and the true cost had we been able to determine the optimal policy. Although recent work by Klabajan and Adelman (2007) promises to provide more rigor to the appropriate choice of approximating class, this issue currently remains as much an art as a science.

Simulation and analytical approaches have been used to determine the optimal approximation *within* a given class. Simulation-based solutions generate sample paths of the problem and seek to update the parameters that determine the chosen class of functions in an iterative fashion. Such methods suffer from the fact that not only is the true value function approximated, but a further source of approximation is introduced through sampling error. This paper focuses on an analytical solution first developed by Schweitzer and Seidmann (1985), with more recent work by Adelman (2005, 2004) and de Farias and Van Roy (2004b, a; 2003). The method of solution proceeds as follows:

1. Transform the MDP into its equivalent LP.
2. Approximate the value function by assuming a specific parameterized form.
3. Use the chosen approximation in the LP to create the ALP.

4. Solve the ALP to obtain the optimal linear value function approximation, v_{ALP} .

5. Use v_{ALP} to determine the “best” action for any visited state.

A fundamental result in MDP theory (Puterman 1994) implies that solving the optimality Equation (5) is equivalent to solving the following LP for any strictly positive α :

$$\max_{\vec{v}} \sum_{\vec{x}, \vec{y} \in S} \alpha(\vec{x}, \vec{y}) v(\vec{x}, \vec{y}) \quad (6)$$

subject to

$$\begin{aligned} c(\vec{a}, \vec{z}) + \gamma \sum_{\vec{d} \in D} \left[p(\vec{d}) v \left(x_2 + \sum_{i=1}^I a_{i2}, \dots, x_N + \sum_{i=1}^I a_{iN}, 0; \right. \right. \\ \left. \left. y'_1 + y_1 - \sum_{n=1}^N a_{1n} - z_1, \dots, y'_I + y_I - \sum_{n=1}^N a_{In} - z_I \right) \right] \geq v(\vec{x}, \vec{y}) \\ \forall (\vec{a}, \vec{z}) \in A_{\vec{x}, \vec{y}} \text{ and } (\vec{x}, \vec{y}) \in S. \quad (7) \end{aligned}$$

Without loss of generality, we assume that α is a probability distribution over the initial state of the system. The conversion to an LP does not avoid the curse of dimensionality because the LP has a variable for every state and a constraint for every state-action pair. A possible solution is to approximate the value function, v , with a linear combination of basis functions. As mentioned earlier, choosing a good set of basis functions remains a challenge within ADP. A reasonable starting point in our model is the following affine approximation:

$$v(\vec{x}, \vec{y}) = W_0 + \sum_{n=1}^N V_n x_n + \sum_{i=1}^I W_i y_i. \quad (8)$$

The advantage of this simple approximation is that the parameters are easily interpreted. V_n represents the marginal infinite-horizon discounted cost of an occupied slot on day n , and W_i represents the marginal infinite-horizon discounted cost of having one more patient of priority class i waiting to be booked. We impose the further restriction that all V_n and W_i are nonnegative, whereas W_0 is unconstrained. Reformulating the LP in terms of this approximate value function yields the following ALP:

$$\max_{\vec{v}, \vec{W}} \sum_{\vec{x}, \vec{y}} \alpha(\vec{x}, \vec{y}) \left(W_0 + \sum_{n=1}^N V_n x_n + \sum_{i=1}^I W_i y_i \right) \quad (9)$$

subject to

$$\begin{aligned} W_0 + \sum_{n=1}^N V_n x_n + \sum_{i=1}^I W_i y_i - \gamma \sum_{\vec{d} \in D} \left[p(\vec{d}) \left(W_0 + \sum_{n=1}^{N-1} V_n \right. \right. \\ \left. \left. \cdot \left(x_{n+1} + \sum_{i=1}^I a_{i, n+1} \right) + \sum_{i=1}^I W_i \left(y'_i + y_i - \sum_{n=1}^N a_{in} - z_i \right) \right) \right] \\ \leq c(\vec{a}, \vec{z}) \quad \forall (\vec{a}, \vec{z}) \in A_{\vec{x}, \vec{y}} \text{ and } (\vec{x}, \vec{y}) \in S, \end{aligned}$$

$$\vec{V}, \vec{W} \geq 0.$$

Rearranging terms and using the assumption that α is a probability distribution transforms the ALP into

$$\max_{\vec{V}, \vec{W}} \left\{ W_0 + \sum_{n=1}^N E_\alpha[X_n] V_n + \sum_{i=1}^I E_\alpha[Y_i] W_i \right\} \quad (10)$$

subject to

$$\begin{aligned} (1 - \gamma) W_0 + \sum_{n=1}^N V_n \left(x_n - \gamma x_{n+1} - \gamma \sum_{i=1}^I a_{i, n+1} \right) \\ + \sum_{i=1}^I W_i \left((1 - \gamma) y_i + \gamma \left(\sum_{n=1}^N a_{in} + z_i - E[Y_i] \right) \right) \leq c(\vec{a}, \vec{z}) \\ \forall (\vec{a}, \vec{z}) \in A_{\vec{x}, \vec{y}} \text{ and } (\vec{x}, \vec{y}) \in S, \\ \vec{V}, \vec{W} \geq 0. \end{aligned}$$

The additional variables x_{N+1} and $a_{i, N+1}$ are constrained to be zero (because no bookings occur beyond day N). X_n and Y_i are random variables, with respect to the probability distribution α , representing the number of appointment slots already booked on day n and the number of priority i patients waiting to be booked, respectively. Although the ALP has only $N + I + 1$ variables, the number of constraints remains intractable. We therefore formulate the dual of the ALP:

$$\min_{\vec{X}} \sum_{\substack{(\vec{x}, \vec{y}) \in S \\ (\vec{a}, \vec{z}) \in A_{\vec{x}, \vec{y}}}} X(\vec{x}, \vec{y}, \vec{a}, \vec{z}) c(\vec{a}, \vec{z}) \quad (11)$$

subject to

$$(1 - \gamma) \sum_{\substack{(\vec{x}, \vec{y}) \in S \\ (\vec{a}, \vec{z}) \in A_{\vec{x}, \vec{y}}}} X(\vec{x}, \vec{y}, \vec{a}, \vec{z}) = 1, \quad (12)$$

$$\sum_{\substack{(\vec{x}, \vec{y}) \in S \\ (\vec{a}, \vec{z}) \in A_{\vec{x}, \vec{y}}}} X(\vec{x}, \vec{y}, \vec{a}, \vec{z}) \left(x_n - \gamma x_{n+1} - \gamma \sum_{i=1}^I a_{i, n+1} \right) \geq E_\alpha[X_n] \\ \forall n = 1, \dots, N, \quad (13)$$

$$\begin{aligned} \sum_{\substack{(\vec{x}, \vec{y}) \in S \\ (\vec{a}, \vec{z}) \in A_{\vec{x}, \vec{y}}}} X(\vec{x}, \vec{y}, \vec{a}, \vec{z}) \left((1 - \gamma) y_i + \gamma \left(\sum_{n=1}^N a_{in} + z_i - E[Y_i] \right) \right) \\ \geq E_\alpha[Y_i] \quad \forall i = 1, \dots, I, \quad (14) \end{aligned}$$

$$\vec{X} \geq 0. \quad (15)$$

Solving the dual has the advantage of a reasonable number of constraints, but at the expense of creating an intractable number of variables—one for each state-action pair. Column generation solves this problem by starting with a small set S' of feasible state-action pairs to the dual and then (using the dual prices as estimates for W_0 , V_n , and W_i) finding one or more violated constraints in the primal. It then adds the state-action pair(s) associated with these violated constraints into the set S' before re-solving the dual. The process iterates until either no primal constraint is violated

or one is “close enough” to optimality to quit. In general, it may be difficult both to find an initial feasible set S' and to find a violated primal constraint. Fortunately, in our model, an initial feasible state-action pair for the dual consists of a state with no available slots and where all incoming demand is diverted. Finding a most violated primal constraint involves solving the following integer program:

$$\begin{aligned} & z(\vec{V}, \vec{W}) \\ & = \min_{\substack{(\vec{x}, \vec{y}) \in S \\ (\vec{a}, \vec{z}) \in A_{\vec{x}, \vec{y}}}} \left[\sum_{i,n} b(i, n) a_{in} + \sum_{i=1}^I d(i) z_i + f(i) \left(y_i - \sum_{n=1}^N a_{in} - z_i \right) \right. \\ & \quad - \sum_{n=1}^N V_n \left(x_n - \gamma(x_{n+1}) - \sum_{i=1}^I a_{i, n+1} \right) \\ & \quad \left. - \sum_{i=1}^I W_i \left((1-\gamma)y_i + \gamma \left(\sum_{n=1}^N a_{in} + z_i - E[Y_i] \right) \right) \right. \\ & \quad \left. - (1-\gamma)W_0 \right]. \end{aligned}$$

Rearranging terms yields

$$\begin{aligned} & z(\vec{V}, \vec{W}) \\ & = \min_{\substack{(\vec{x}, \vec{y}) \in S \\ (\vec{a}, \vec{z}) \in A_{\vec{x}, \vec{y}}}} \left[\sum_{n=1}^N \left(\sum_{i=1}^I (b(i, n) + \gamma V_{n-1} - f(i) - \gamma W_i) a_{in} \right. \right. \\ & \quad \left. \left. + (\gamma V_{n-1} - V_n) x_n \right) \right. \\ & \quad \left. + \sum_{i=1}^I ((d(i) - f(i) - \gamma W_i) z_i + (f(i) + \gamma W_i - W_i) y_i) \right. \\ & \quad \left. + \sum_{i=1}^I \gamma W_i E[Y_i] - (1-\gamma)W_0 \right]. \quad (16) \end{aligned}$$

The coefficients on a_{in} in Equation (16) have a nice intuitive interpretation in terms of balancing the costs versus the benefits of each action. For each action a_{in} there is a cost, $b(i, n) + \gamma V_{n-1}$, due to a (possibly) late scan and the loss of available capacity tomorrow as well as a benefit, $f(i) + \gamma W_i$, due to the fact that the booking decision is not delayed and the patient does not reappear in tomorrow’s demand. For each action, z_i , there is a cost, $d(i)$, due to diverting the patient, which is likewise weighed against the benefit of not delaying the booking decision and therefore not having the patient appear in tomorrow’s demand.

4. The Form of the Optimal Linear Value Function Approximation

Once the dual is solved, the prices associated with each constraint determine the coefficients in the best linear value function approximation (denoted by v_{ALP}). Investigating the properties of solutions to a wide range of numerical instances led to a conjecture of the form of the optimal primal solution. This leads to the theoretical results in this

section, which provide interpretable conditions under which the optimal solution to the primal ALP, v_{ALP} , can be solved directly.

The form of v_{ALP} depends to some extent on the nature of the cost functions. Earlier discussion suggested that a reasonable choice for the booking cost is

$$b(i, n) = \begin{cases} \sum_{k=1}^{n-T(i)} \gamma^{k-1} f(i) & \text{for all } n > T(i), \\ 0, & \text{otherwise.} \end{cases}$$

(In fact, we present some minimal restrictions on the form of $b(i, n)$ to achieve our results. These conditions include any scenario where late costs increase at a faster-than-linear rate in the days.) More critical is the form of the cost for diverting patients to an alternative capacity source. If that alternative capacity source is overtime, then it would seem reasonable to assume that the diversion cost is independent of i because overtime costs are a function of the length of the scan and not the priority of the patient. Alternatively, if diversion means that demand is sent elsewhere (i.e., rejected by the hospital in question), then it would seem reasonable to assume that the diversion cost is strictly decreasing in i . Such a cost function reflects the fact that demand that is sent elsewhere often faces an additional delay, and thus is more costly for higher-priority demand. We present two theorems that give the optimal form of v_{ALP} for these two scenarios.

4.1. The Optimal Linear Value Function Approximation with Overtime

The first theorem gives the optimal linear value function approximation, v_{ALP} , for the scenario where $d(i)$ is constant. (The proof appears in the appendix.)

THEOREM 1. *Assume that the cost of diverting a patient is constant for all i (i.e., $d(i) = d$), $T(i)$ is decreasing in i , and the late booking cost, $b(i, n)$, is nondecreasing in n and nonincreasing in i with $b(i, n) = 0$ for all $n \leq T(i)$. Assume further that*

$$b(i, n) + \gamma^{n-T(i)} d > b(i, T(i)) + \gamma^{T(i)-T(i)} d \quad (17)$$

for all $n > T(i)$ and for all i ;

$$\sum_{i=1}^I \frac{\gamma^{T(i)-n}}{1-\gamma} I_{T(i)>n} \lambda_i + \sum_{m=n}^N \gamma^{[m-n]^+} E_\alpha[X_n] < \frac{C_1}{1-\gamma} \quad (18)$$

for all $n \geq T(1)$; and

$$\begin{aligned} 0 & < \sum_{i=1}^I \frac{\gamma^{T(i)-T(1)}}{1-\gamma} \lambda_i + \sum_{n=1}^N \gamma^{[n-T(1)]^+} E_\alpha[X_n] - T(1)C_1 - \frac{\gamma C_1}{1-\gamma} \\ & < \frac{C_2}{1-\gamma}, \end{aligned} \quad (19)$$

where $I_{T(i)>n}$ is equal to one if $T(i) > n$ and zero otherwise, λ_i is the arrival rate for demand from priority class i , C_1 is equal to the base capacity, C_2 is the surge capacity (i.e., overtime), and γ is the discount rate. Then, the optimal linear value function approximation for the discounted MDP will have the following form:

$$V_n = \begin{cases} d & \text{for all } n \in \{1, \dots, T(1)\}, \\ \gamma V_{n-1} & \text{for all } n \in \{T(1) + 1, \dots, N - 1\}, \\ 0 & \text{for } n = N, \end{cases} \quad (20)$$

$$W_i = V_{T(i)} \quad \text{for all } i \in \{1, \dots, I\},$$

$$W_0 = d \left(\gamma \sum_{i=1}^I \frac{\gamma^{T(i)-T(1)}}{1-\gamma} \lambda_i - T(1)C_1 - \frac{\gamma C_1}{1-\gamma} \right).$$

The above form of v_{ALP} has considerable intuitive appeal. From a cost standpoint, the marginal cost of each slot on days up to and including $T(1)$ are identical; thus, one would expect to value these days equally. It is also intuitively appealing to assign a value equal to d for these days because the availability of this capacity allows the manager to avoid using surge capacity. After day $T(1)$, the value of an appointment slot on day n is equal to γ times the value of an appointment slot on day $n-1$ because the capacity on day n this decision epoch will be the capacity on day $n-1$ by the next decision epoch. For this reason, $V_n = \gamma V_{n-1}$ is reasonable.

Equation (17) requires that the cost of booking a patient on day $n > T(i)$ and then performing an overtime scan $n - T(1)$ days into the future be greater than the cost of booking on day $T(i)$ (assumed to be zero) and then performing an overtime scan $T(i) - T(1)$ days into the future. This reflects the fact that by booking a patient late, the scheduler has essentially only delayed the need for overtime by $n - T(i)$ days. Note that the less that future costs are discounted, the more likely that Equation (17) will be satisfied. For example, with $\gamma = 0.9$, Equation (17) will be violated if the cost of overtime, d , is approximately 10 times greater than the daily cost of a late booking, $f(i)$. If $\gamma = 0.99$, then d needs to be more than 100 times greater than $f(i)$. Therefore, the high choice of γ appropriate for the health-care setting implies that even with a small late booking penalty, Equation (17) will hold.

In traditional DP theory, the solution to the LP is known to be independent of α provided α is strictly positive for all states (Puterman 1994). However, in ADP, this is not the case (see Adelman 2004 and de Farias and Roy 2003), but the nature of the dependence of the optimal approximation on α is not very well understood. In this instance, interpreting α as a probability distribution over the initial state of the system gives Equations (18) and (19) concise interpretations. Any choice of α satisfying these two equations will yield the same v_{ALP} . We leave until later a discussion of the impact of violating these conditions.

Equation (18) requires that for any given day, $n \geq T(1)$, there is sufficient base capacity to schedule the average

demand for any priority class with a wait-time target exceeding n . In essence, this insures that overtime is only required for the highest-priority class. This condition is unlikely to be violated unless the system is either extremely undercapacitated (in which case the overtime requirements will become prohibitive) or the highest-priority class generates negligible demand in comparison to the other classes.

The first two terms in the body of Equation (19) equal the present value of the expected demand over the infinite horizon plus the present value of the expected number of appointment slots initially filled. The last two terms represent the present value of the total base capacity over the infinite horizon. (Recall that all slots are of equal value up to day $T(1)$ and are discounted by γ thereafter.) Thus, stating that the body of Equation (19) has to be greater than zero is equivalent to insuring that total expected demand exceeds total available capacity. In other words, capacity is a legitimate constraint. Stating that the body of Equation (19) has to be less than $C_2/(1-\gamma)$ insures that there is sufficient overtime capacity available to insure that appropriate scheduling can avoid exploding queues. This upper bound is of significant practical importance because it determines the necessary overtime capacity commitment for a given base capacity in order to meet the wait-time targets.

Although the three conditions place significant restrictions on the parameter values, they nonetheless allow for a wide range of realistic scenarios. Their intuitive interpretations also demonstrate their plausibility. Even if these constraints are violated, the ALP still yields a value function; it simply does not have the form given in Theorem 1.

4.2. The Optimal Linear Value Function Approximation with Rejected Demand

A similar analysis for the scenario where $d(i)$ is decreasing in i yields the following theorem.

THEOREM 2. Assume that the cost of rejecting demand, $d(i)$, satisfies

$$d(i) > \gamma^{T(i)-T(1)} d(I) \quad (21)$$

for all $i < I$; that $T(i)$ is decreasing in i ; and that the late booking cost function is nondecreasing in n and non-increasing in i with $b(i, n) = 0$ for all $n \leq T(i)$. If

$$\gamma^{T(i)-T(1)} b(i, n) + \gamma^{n-T(i)} d(I) > d(I) \quad (22)$$

for all $n > T(i)$ and for all i ;

$$\sum_{i=1}^I \frac{\gamma^{T(i)-n}}{1-\gamma} I_{T(i)>n} \lambda_i + \sum_{m=n}^N \gamma^{[m-n]^+} E_\alpha[X_n] < \frac{C_1}{1-\gamma} \quad (23)$$

for all $n \geq T(1)$; and

$$0 < \sum_{i=1}^I \frac{\gamma^{T(i)-T(1)}}{1-\gamma} \lambda_i + \sum_{n=1}^N \gamma^{[n-T(1)]^+} E_\alpha[X_n] - T(1)C_1 - \frac{\gamma C_1}{1-\gamma} < \frac{\gamma^{T(i)-T(1)} C_2}{1-\gamma}; \quad (24)$$

then the optimal linear value function approximation for the discounted MDP will have the following form:

$$V_n = \begin{cases} \gamma^{T(1)-T(i)} d(I) & \text{for all } n \in \{1, \dots, T(1)\}, \\ \gamma V_{n-1} & \text{for all } n \in \{T(1)+1, \dots, N-1\}, \\ 0 & \text{for } n = N, \end{cases}$$

$$W_i = V_{T(i)} \quad \text{for all } i \in \{1, \dots, I\},$$

$$W_0 = \gamma^{T(1)-T(i)} d(I) \left(\gamma \sum_{i=1}^I \frac{\gamma^{T(i)-T(1)}}{1-\gamma} \lambda_i - T(1) C_1 - \frac{\gamma C_1}{1-\gamma} \right). \quad (25)$$

The intuition is similar to that in the overtime scenario. Note, however, that the driving cost factor is the cost of redirecting the lowest-priority class, $d(I)$.

5. Deriving a Policy from the Approximate LP Solution

In this section, we discuss how to derive a policy from the solution to the ALP. We refer to this policy as the *approximate optimal policy* or AOP. Further, we show that the AOP has an elegant and intuitive structure for a wide range of cost parameters.

Traditional LP methods for directly solving an MDP yield the complete optimal policy by setting the probability of using action a in state s equal to the value of the dual variable, $X(s, a)$, divided by the sum of the dual variables over all possible actions in state s . The viability of this method depends on the fact that a direct solution to the LP will have at least one positive variable, $X(s, a)$, for all states s . In solving the ALP through column generation, only a very small percentage of all possible states are evaluated, and thus the traditional method for deriving a policy fails. Instead, in the ADP setting, we insert v_{ALP} into the right-hand side of the optimality equation (5) and then solve for the optimal action (\vec{a}, \vec{z}) in state (\vec{x}, \vec{y}) as needed in practice or in a simulation. This involves solving the following integer program:

$$\min_{(\vec{a}, \vec{z}) \in A_{\vec{x}, \vec{y}}} \left\{ \sum_{i=1}^I \sum_{n=1}^N b(i, n) a_{i, n} + \sum_{i=1}^I \left(d(i) z_i + f(i) \left(y_i - \sum_{n=1}^N a_{i, n} - z_i \right) \right) + \gamma \sum_{\vec{d} \in D} p(\vec{d}) \left[W_0 + \sum_{n=1}^N V_n \left(x_{n+1} + \sum_{i=1}^I a_{i, n+1} \right) + \sum_{i=1}^I W_i \left(y'_i + y_i - \sum_{n=1}^N a_{i, n} - z_i \right) \right] \right\}$$

$$= \min_{(\vec{a}, \vec{z}) \in A_{\vec{x}, \vec{y}}} \left\{ \sum_{n=1}^N \sum_{i=1}^I (b(i, n) + \gamma V_{n-1} - f(i) - \gamma W_i) a_{i, n} + \sum_{i=1}^I (d(i) - f(i) - \gamma W_i) z_i \right\} + \text{constant}$$

$$= \min_{(\vec{a}, \vec{z}) \in A_{\vec{x}, \vec{y}}} \left\{ \sum_{n=1}^N \sum_{i=1}^I A_{i, n} a_{i, n} + \sum_{i=1}^I Z_i z_i \right\} + \text{constant}. \quad (26)$$

For each priority class i , the choice of action in the AOP depends on the coefficients, $A_{i, n}$ and Z_i . Clearly, the AOP will only book patients of each priority class into those days for which $A_{i, n} < 0$ and will only use overtime for those priority classes for which $Z_i < 0$. Because the coefficient $A_{i, 1}$ is strictly less than zero for all i , any capacity available on day 1 will be used by whatever demand is available. As the following proposition states, under the conditions outlined in Theorem 1, it is possible to derive a number of relevant properties of these coefficients. When these conditions hold, the AOP can be determined *without* solving the above integer program.

PROPOSITION 1. *Suppose that the diversion costs $d(i) = d$ for all i , the booking cost is nondecreasing in n and nonincreasing in i with $b(i, n) = 0$ for all $n \leq T(i)$ and for all i ; and that Equations (17) through (19) hold. Then,*

- $A_{i, n}$ is increasing in i .
- $A_{i, n}$ is convex in n with minimum at $n = T(i)$.
- $A_{i, n} \geq 0$ for all $n \geq T(i) + 1$ with equality at $n = T(i) + 1$.
- Z_i is increasing in i .

Moreover, for $n \leq T(i)$, the coefficient $A_{i, n}$ will be negative if and only if

$$f(i) > (\gamma^{[n-T(1)-1]^++1} - \gamma^{T(i)-T(1)+1})d, \quad (27)$$

and the coefficient Z_i will be negative if and only if

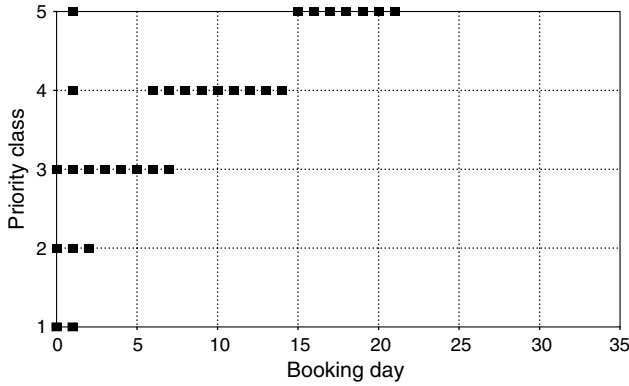
$$f(i) > (1 - \gamma^{T(i)-T(1)+1})d. \quad (28)$$

As a result of Proposition 1, there exists an interval, $(LB(i), UB(i))$, for each priority class with the property that $A_{i, n}$ is only negative for $n \in \{1 \cup (LB(i), UB(i))\}$. Moreover, $UB(i) = T(i)$ and $LB(i)$ is equal to the smallest n satisfying Equation (27). The following theorem describes how to implement the AOP.

THEOREM 3. *If the diversion costs $d(i) = d$ for all i , the booking cost is nondecreasing in n and nonincreasing in i with $b(i, n) = 0$ for all $n \leq T(i)$ and for all i ; and Equations (17) through (19) hold, then the AOP may be implemented as follows:*

- Book patients in order of priority class.
- Book as much priority 1 demand as possible into the interval $(1, T(1))$ starting with day 1 and working up to day $T(1)$.
- For each successive priority class, book incoming demand into any available slots in the interval $\{1 \cup (LB(i), UB(i))\}$ starting with day 1, then day $UB(i)$, and working down to day $LB(i)$.
- If there is any remaining demand, use overtime for a given priority class only if Equation (28) is satisfied

Figure 2. AOP with overtime ($I = 5$, $f = (20, 20, 18, 9, 6)$, $T = (1, 2, 7, 14, 21)$, $d = 100$); black squares indicate available booking days for each priority class.



and giving precedence to higher-priority demand should the overtime capacity constraint be an issue.

- For all remaining demand, delay booking.

Figure 2 provides a graphical representation of the booking intervals $(LB(i), UB(i))$ for a scenario with five priority classes and with a specific combination of parameters. Day zero represents overtime.

The policy can be fully characterized by $2I + 1$ numbers (the interval $(LB(i), UB(i))$ for each i and the lowest-priority class for which overtime is used), all of which can be determined analytically without solving the integer program given in Equation (26). Unless the facility is significantly under capacity, the lower bounds on the booking intervals will not be a factor because it is unlikely that day $T(i)$ will be filled by lower-priority classes because higher-priority demand has yet to be booked into that day. Thus, the probability of the whole interval $(LB(i), UB(i))$ being filled is remote in most realistic settings. We briefly discuss the policy implications of violations of the conditions in Theorem 1.

If Equation (18) is violated (but the other conditions are satisfied), then empirical results suggest $V_n = d$ for all $n \leq T(i')$, where $T(i')$ is the first wait-time target for which Equation (18) is satisfied. The resulting policy treats patients in classes 1 through i' similarly. If Equation (17) is violated for a given priority class j and a given day m_j , then it is violated for all priority classes lower than class j and for all $n \leq m_j$. The resulting policy will book patients of priority class j and lower late (how late depends on the value of m_j) while maintaining the wait-time targets for higher-priority demand. If the lower bound in Equation (19) is violated, the value function for the ALP is identically zero, and thus the “optimal” policy books patients on a first-come, first-served basis. In this case, capacity is not a significant limitation, and thus demand can be met as it arrives. If the upper bound in Equation (19) is violated,

then the “optimal” policy trades-off booking low-priority patients late to leave greater capacity for higher-priority demand. However, in the long run, costs explode because the system simply does not have sufficient capacity to deal with the incoming demand.

A similar analysis for the case with $d(i)$ decreasing in i yields the following partial characterization of the AOP:

- Book as much priority 1 demand as possible into the interval $(1, T(1))$ booking first on day 1, then day 2, and continuing to day $T(1)$.
- For each successive priority class except the lowest one, book incoming demand into any available slots in the interval $\{1 \cup (LB(i), UB(i))\}$ starting with day 1, then day $UB(i)$, and working down to day $LB(i)$.
- If, for any priority class $i < I$, there is insufficient capacity in $(LB(i), UB(i))$ to meet all priority i demand, then delay booking.

This is, however, not a complete policy because it fails to specify the scheduling of the lowest-priority class. This cannot be determined without solving the integer program given in Equation (26). Thus, to implement this policy requires the availability of integer programming software. Fortunately, solving Equation (26) can be done within a reasonable time frame (see discussion below), and thus the necessity of solving it repeatedly is not an insurmountable stumbling block.

5.1. Ethical Considerations

Three features of the policy may raise some concerns regarding the equitable nature of the proposed scheduling policy. First, because lower-priority demand is booked in reverse order from day $UB(i)$ to $LB(i)$, there is the potential for later-arriving demand to be given an earlier scan date. In practice, however, this is not an issue because it is only in rare instances that day $UB(i)$ is ever full for priority classes $i > 1$. Thus, invariably, the wait time for lower-priority patients is exactly $UB(i)$ days. The only instance where this might become an issue is if capacity does not exceed the average demand even after excluding the highest-priority class, but then this is precisely the condition imposed by Equation (18).

The other two potentially worrying features of the policy raise concerns only when there are no inpatients (that is, the wait-time target of the highest-priority class is not the same day the request is placed). In such a case, any demand may be used to fill available capacity tomorrow because by tomorrow that capacity will cease to be available. Thus, lower-priority demand may in fact get expedited service simply because the request is placed on the eve of a day with excess capacity. This could be argued to be unfair, but the alternative is simply to lose that capacity. One could potentially shuffle appointments to maintain equity, but the additional administrative burden hardly seems worth the effort. In the presence of inpatients, this is not an issue because only inpatients can be booked the same day the

request is placed, and therefore excess capacity is only ever used by the highest-priority class.

The final feature concerns the use of overtime. Again, if inpatients are present, outpatients are, in any realistic setting, never given overtime slots; therefore, there is no issue of equity. However, in the absence of inpatients, it is possible that a high-priority patient arriving at a time of high congestion receives an overtime slot and thus gets expedited service, when a patient of the same priority class who arrived a day earlier was perhaps forced to wait $T(1)$ days. Again, if such inequities are of significant concern, then an alternative is to shuffle appointments in order to maintain a first-come-first-served policy within each priority class. However, the added administrative burden again would seem to outweigh the benefit; especially considering that even without reshuffling no patient is waiting longer than the medically desirable wait-time target. More reasonably, one could modify the policy so that overtime slots are booked $T(1)$ days in advance, thus insuring that any patient who receives an overtime slot waits at least as long as any patient who does not. This would maintain equity and also allow for some advance planning of overtime requirements.

6. Simulation Results

Because the AOP was derived through ADP, there is no guarantee of its optimality. This section reports the results of using simulation to investigate the behavior of the policy in a variety of scenarios. Even though we optimize with respect to costs, we use the simulation to investigate operationally meaningful performance measures including the percentage of patients exceeding wait-time targets, the percentage of patients who are diverted, and system utilization. We do not consider any scenario where more than one source of surge capacity is available (i.e., overtime and rejected demand). Although this is certainly a possibility, the robustness of the overtime policy shown through the simulation results suggests that, in fact, using overtime as the sole source of surge capacity is most advantageous.

The coding of the column generation algorithm required to solve the ALP was done in AMPL with CPLEX as the solver. For scenarios booking 30 days in advance, up to five priority classes and up to 200 available scanning slots a day, the time to solution was less than five minutes. The simulation of the scheduling process was also done in AMPL/CPLEX because it involves solving the integer program given in Equation (26). The simulation length was 20,000 days with statistics collected after the first 5,000 days. Each scenario was run 10 times, with the resulting 95% confidence interval provided for each statistic. The run time for a simulation of 20,000 days was approximately 15 minutes on a good laptop PC. The discount factor is set at 0.99 for all simulation scenarios.

6.1. A Small Outpatient Clinic

Consider a small outpatient clinic with a capacity of 10 appointment slots per day with a maximum of four patients diverted per day. (To satisfy Equations (19) and (24), one only needs $C_2 \geq 1$.) The clinic divides demand into three priority classes with wait-time targets of 7, 14, and 21 days, respectively, and chooses a 30-day booking horizon. Demand for each priority class is Poisson with means 5, 3, and 2 respectively. (The Poisson distribution is truncated at three times the mean for each priority class in order to maintain a finite state space.) In such a scenario, average demand equals base capacity, and thus a lack of capacity is a legitimate concern.

As stated in Theorems 1 and 2, the simulation results are robust to changes in the exact values of the diversion costs and the late booking penalties provided Equation (17) is satisfied (if diversion costs are constant) or Equations (21) and (22) are satisfied (if diversion costs are decreasing in i). Estimates for the cost of overtime vary between \$30 per scan (the additional overtime salary of two CT technicians for a 15-minute scan) to as high as \$150 per scan (if one includes other costs). For a discount rate of 0.99, this implies that even at the high end of the scale, late booking costs need only be set at two per day or higher for the optimal linear value function approximation to have the form given in Theorem 1 and for the AOP to be as given in Theorem 3. In this scenario, overtime cost, d , is set at 100 and the late booking penalties, $f(i)$, are set at 20, 10, and 5. In the case where $d(i)$ is decreasing in i , the cost of diverting patients is set at 100, 50, and 25, and the late booking costs are kept the same.

We compare the AOP to a strict booking limit policy where each priority class is booked into a given day only if there is a predetermined number of appointment slots still available. The optimal booking limit policy (determined through enumeration) for the constant diversion costs is to book priority 2 patients into a given day only if seven slots are still available and to book priority 3 patients only if nine slots are still available (no booking limit is imposed on day one). For the scenario with decreasing diversion costs, the optimal booking limit policy is to book priority 2 patients if there are six or more slots still available and to only book priority 3 patients if there is available space on day one. The booking limit policy resorts to surge capacity only if there is no available capacity for a given priority class within the booking horizon. Results are summarized in Table 1.

Both the overtime policy and the rejection policy outperform the corresponding fixed booking limit policies with respect to the percent of cases booked late (in the case of OT); the percent of cases served through OT/diverted (both cases); and in capacity utilization (both cases). In terms of capacity utilization and the number of patients who are diverted, it is clear that the AOP in the overtime case also outperforms the AOP in the rejection case. The overtime policy *reacts* to high levels of demand by using

Table 1. A comparison of AOP and booking limit policies in the overtime and excess demand rejection cases with respect to three performance measures.

Criteria	Priority class	AOP overtime $d = 100$	Booking limit $d = 100$ BL = (1, 7, 9)	AOP removal $d = (100, 50, 25)$	Booking limit $d = (100, 50, 25)$ BL = (1, 6, -)
Percent late	P1	0.22 ± 0.04	0	4.94 ± 1.32	0
	P2	0	0.42 ± 0.17	0	0
	P3	0	47.78 ± 0.38	0	0
	Overall	0.11 ± 0.02	9.69 ± 0.13	2.47 ± 0.66	0
Percent diverted	P1	1.56 ± 0.07	0	0.33 ± 0.09	0
	P2	0	0	0	0
	P3	0	20.97 ± 0.78	11.64 ± 0.32	50.32 ± 0.66
	Overall	0.78 ± 0.07	4.20 ± 0.16	2.49 ± 0.11	10.4 ± 0.13
Utilization percentage		99.05 ± 0.08	95.73 ± 0.14	97.34 ± 0.09	89.86 ± 0.16

Note. P1 refers to priority class 1, P2 to priority class 2, and P3 to priority class 3.

overtime to serve the highest-priority class. The rejection policy, on the other hand, seeks to anticipate high levels of demand by diverting the lowest-priority class *preemptively*. Because it is easier to react than predict, the AOP in the overtime case generally performs better.

If the AOP policy is implemented *with the modification that surge capacity is used only if absolutely necessary* (that is, when there is no available capacity in the booking horizon), then the results are dramatically worse, with over 50% of priority 1 patients booked late as well as significant portions of the lower-priority classes. It is worth noting that the current practice of simply postponing the booking of any demand that cannot be met will in fact perform worse than this “reject as a last resort” policy. The sole reason that costs do not tend to infinity under the current policy is that demand is negatively correlated with expected wait time. Thus, there is an implicit rejection of demand occurring under current practice, but the decision as to when to reject demand is made not by the resource manager, but by the specialist doctors who recommend the scans in the first place.

6.2. A Large Outpatient Clinic

Increasing the size of the outpatient clinic to 60 scans per day and letting the lower-priority classes represent the larger portion of demand leads to similar results for the overtime policy with a marked improvement in the rejection policy. Table 2 presents simulation results for a scenario where demand is Poisson with rates 10, 20, and 30 for each priority class, respectively. Diversion capacity is set at 4 (although 3 is all that is required to satisfy Equations (19) and (24)). Again, varying the diversion cost (be it overtime or rejection) does not greatly affect the results, provided that the conditions given in the two theorems are satisfied.

6.3. The Hospital Setting

Theorems 1 and 2 prove that adding inpatients who require scans immediately, or by the next day at the latest, should

not affect the optimal linear value function approximation (and by extension the AOP) because the conditions are independent of changes in the number of priority classes and reasonably robust to changes in the wait-time targets. Column 1 of Table 3 presents the simulation results for a scenario where capacity is set at 120 scans per day, inpatient (IP) demand is Poisson with rate 60, and the outpatient (OP) demand is the same as in the previous scenario. Following current practice at the Vancouver hospital, the wait-time target for inpatient demand is set at zero. That is, IP demand is satisfied the same day the request is placed. Thus, day 1 now represents the current day. Because OP demand comes in the form of faxed requisitions throughout the day, it is necessary to impose the further constraint that no OP demand can be booked on day 1.

The introduction of inpatient demand leads to a dramatic increase in overtime requirements. Whereas for an outpatient clinic, an average of one overtime scan every 20 days suffices, for a hospital serving both inpatients and outpatients an average of four overtime scans per day is needed. The theoretical overtime capacity determined by Equation (19) jumps to 13. Decreasing the wait-time target of the highest-priority class from seven days to one

Table 2. Comparison of AOPs in a larger outpatient clinic.

Criteria	Priority class	Overtime policy $d = 100$	Rejection policy $d = (100, 50, 25)$
Percent late	P1	0.42 ± 0.48	0
	P2	0	0
	P3	0	0
	Overall	0.07 ± 0.02	0
Percent diverted	P1	0.48 ± 0.15	0
	P2	0	0
	P3	0	1.42 ± 0.06
	Overall	0.08 ± 0.02	0.79 ± 0.03
Utilization percentage		99.85 ± 0.04	99.23 ± 0.03

Table 3. Comparison of AOPs in the presence of inpatient demand.

Criteria	Priority class	OT policy (100% HIP) $d = 100$	OT policy (90% HIP) $d = 100$	Rejection policy $d = (100, 90, 80, 40, 20)$
Percent late	HIP	1.49 ± 0.06	1.07 ± 0.05	0.17 ± 0.01
	LIP	n/a	0	n/a
	OP1	0	0	0
	OP2	0	0	0
	OP3	0	0	0
	All	0.65 ± 0.03	0.54 ± 0.04	0.08 ± 0.01
Percent diverted	HIP	6.73 ± 0.08	5.54 ± 0.09	2.65 ± 0.04
	LIP	n/a	0	n/a
	OP1	0	0	0
	OP2	0	0	0
	OP3	0	0	29.52 ± 0.06
	All	2.93 ± 0.04	2.59 ± 0.04	8.70 ± 0.02
Utilization percentage		96.6 ± 0.03	97.47 ± 0.03	91.27 ± 0.06

greatly reduces the resource manager’s ability to “smooth” out the variability in demand, thus increasing the overtime requirement. Patrick and Puterman (2007) used simple probabilistic arguments and simulation to explore the impact of introducing some flexibility into the scheduling of the IP class. Such flexibility is available at the Vancouver hospital, where there is an IP class for patients whose scans could be delayed one day. However current practice is to ignore this flexibility and perform all IP scans on the day of the request. Consider, therefore, two IP priority classes—a high-priority class (HIP) with a wait-time target of zero and a low-priority class (LIP) with a wait-time target of 1. Column 2 of Table 3 presents the results when 10% of the total IP demand is classified as LIP. It demonstrates the significant impact of even this small amount of flexibility in IP scheduling on the cost associated with maintaining wait-time targets.

Finally, column 3 of Table 3 presents the results for the scenario where diversion costs are decreasing in i . The maximum number diverted in a day must be at least 16 to satisfy Equation (24). Not surprisingly, the ability of the rejection policy to anticipate congestion is severely hampered, leading to a higher number of diverted patients as well as a portion of late patients.

6.4. A Large Hospital

Finally, the Vancouver hospital represents a large-sized hospital that operates four CT scanners every weekday. It faces a significant inpatient demand with an average of 126 15-minute appointment requests per day. There are three outpatient priority classes with average daily demands of 9, 19, and 24 requests, respectively. The Kolmogorov-Smirnov test for goodness of fit suggests that a Poisson arrival distribution is appropriate. In the simulation, the base capacity is again set equal to the average demand with

Table 4. Comparison of AOPs at a large hospital.

Criteria	Priority class	Overtime policy (100% HIP, $d = 100$)	Overtime policy (90% HIP, $d = 100$)
Percent late	HIP	0.18 ± 0.01	0.09 ± 0.00
	LIP	n/a	0
	OP1	0	0
	OP2	0	0
	OP3	0	0
	All	0.13 ± 0.01	0.06 ± 0.00
Percent overtime	HIP	4.17 ± 0.05	2.78 ± 0.06
	LIP	n/a	0
	OP1	0	0
	OP2	0	0
	OP3	0	0
Utilization percentage		97.01 ± 0.03	98.17 ± 0.03

24 available overtime slots (each machine with an hour and a half of overtime availability). Table 4 presents the simulation results for the AOP policy in the OT case both with and without 10% of IP demand being low priority. Under both scenarios, only small fractions of patients are booked late or served through overtime with a significant improvement in performance with the added flexibility in IP scheduling.

These four scenarios demonstrate that our approximate optimal policy performs extremely well regardless of the size of the hospital and across a variety of demand streams.

7. How Good Is the Approximation?

One of the most significant challenges facing users of approximate dynamic programming is the absence of accurate bounds on the optimality gap between the cost incurred with the AOP and the cost that would have been incurred under the optimal policy. One available lower bound on the total discounted cost is the optimal objective value for the ALP (Adelman 2006). Because the form of the value function is restricted to a certain class of functions, it is clear that the optimal primal objective in the ALP will be a lower bound on the optimal objective in the original LP. By strong duality, the optimal dual objective in the approximation will also be a lower bound on the dual objective in the original LP. This dual version of the objective function represents the total discounted cost over the infinite horizon, and thus the optimal objective value in the ALP gives a lower bound on the true optimal total discounted cost. Although reasonable in theory, this bound suffers from being highly dependent on the choice of α and on the demand stream early on in the simulation run. Thus, the total discounted cost incurred during the simulation varies dramatically with α and even varies significantly between simulation runs with the same α . If α is set so that all slots during the initial booking horizon are full, then the total discounted cost in the simulation is within 18% of the lower bound ($\pm 3.08\%$). However, if α is chosen so that more capacity is available later in the booking horizon, then the optimality gap

increases to 34% ($\pm 4.83\%$). The intuition is that the LP formulation does not fully capture the time dynamics so that it best mirrors the “true costs” when the costs are incurred sooner rather than later.

As the approximation approaches the true value function, the objective of the ALP also increases towards the objective of the true LP. Thus, the optimality gap represented above is as much a result of an overly optimistic bound as it is the result of suboptimality. The true optimality gap is undoubtedly significantly lower than the bound given here.

Regardless of how well the approximation mirrors reality, one can still compare the outcome of the derived policy with current practice. Current wait times for CT scans within Canada (and elsewhere) are a major concern. As mentioned in the introduction, the wait times at the Vancouver hospital exceeded the targets for the majority of patients despite the fact that capacity is approximately equal to average demand. Thus, it is fair to claim that the AOP does in fact outperform current practice. Although it is possible that a different policy might maintain the same wait-time targets at a lower cost, this does not negate the significant improvement in outcome that the policy derived in this paper achieves.

In an effort to gain some intuition into the form of the true value function, we solved a number of very small instances of the model exactly, using LP. The scenarios involved two priority classes, booking horizons of at most three days and a daily capacity of no more than four. Even some of these small instances took two days to solve. It is unreasonable to draw too many conclusions based on such small examples, but several points seem relevant, though not conclusive.

- A linear regression with the components of the state as the independent variables and the value function as the dependent variable resulted in a R^2 value of between 0.9 and 0.95. Thus, in a small problem, a linear approximation appears to be a reasonable fit.

- When second-order terms including interactions are added to the regression, the R^2 value increased to between 0.96 and 0.99, suggesting that the true form of the value function may involve interactions between the components of the state space.

- The linear value function approximation tended to underestimate the value of a state when the system was close to being empty and to overestimate the value of a state when the system was close to capacity. This suggests that the linear value function approximation overestimates the cost of being in a congested state, thus potentially leading to a more conservative policy.

There is one final piece of evidence to suggest that the linear value function approximation is, in fact, a reasonable one. It arises from pursuing what was originally thought to be a defect in the model—namely, that a linear approximation does not take into account that on any given day, the marginal cost of booking a scanning slot should increase as the number of available slots decreases. Under a linear

approximation, reducing the number of available slots from two to one “costs” the same as reducing the number of available slots from ten to nine. A more reasonable conjecture of the value function might be nondecreasing and convex in the number already booked on each day (x_n). For simplicity the value function approximation was kept linear in the y state variable. Surprisingly, the added flexibility in this model *made no difference because the optimal value function remained the same as in the linear model given earlier*—at least for the scenarios analyzed by the authors. This would seem to support the conclusion that the linear model is an adequate representation of the true value function of the MDP and that the “cost” associated with having approximated the value function is not in fact too significant.

8. Extensions, Further Applications, and Conclusions

Although we have endeavored to insure that the model is as realistic as possible, there are inevitably enhancements that could potentially lead to interesting results. First, what would be the impact of introducing a no-show rate for outpatient appointments? Our intuition is that for any scenario involving inpatient demand, the booking policy would remain unchanged, with the advantage being a reduction in overtime. For an outpatient clinic, however, it is likely that the introduction of no-shows would cause overbooking to be beneficial. Second, one of the simplifying assumptions of the model is that all appointment slots are of equal length. As mentioned early on in the paper, reality is very different, but the impact of that difference does not appear to be significant. Third, what would be the impact of allowing service times to be random? In a small hospital setting this may cause significant deviations in the results. In a larger hospital, however, one would expect the results to be fairly similar provided that the mean service time is in fact equal to the length of the appointment slot. Finally, true overtime costs are more likely to be piecewise linear, reflecting the fact that CT technicians are paid time and a half initially, and then double time should the amount of overtime exceed a given limit. Given the reasonable nature of the analytical bounds on overtime capacity, this is less of an issue than it might have been if overtime had been left unbounded. Finally, it would be of interest to explore the impact of assuming a centralized booking system that services a number of different hospitals. The obvious complication is that the resource, as well as demand, is no longer located in one place, and thus one must take into account patient preferences and travel costs.

Although what we have presented here is adapted primarily to a diagnostic imaging resource, any situation where there are multiple priority classes with priority-specific wait-time targets would be a potential application of the model. The most obvious alternative health application is surgical scheduling. The complication is that most surgical scheduling works on a block design that assigns blocks

of time to each surgeon, who then books his/her patients into those blocks. Thus, our model might be most useful to the surgeons themselves as they seek to minimize excessive wait times for their patients based on a fixed available capacity. Diversion might consist of sending patients to other surgeons or using overtime. A second complication is that for surgical scheduling it is unreasonable to assume a deterministic service time. A second potential application is the scheduling of radiation treatment. The added complication is that radiation treatment involves a series of appointments rather than just one.

Through the use of approximate dynamic programming, we have been able to solve a multipriority patient-scheduling problem that was previously intractable. The derived policy gives clear guidelines as to when to book each priority class and when to resort to overtime (or other alternate sources of capacity). Our analysis also gives the resource manager a lower limit on the amount of overtime required given a specific base capacity, priority-specific wait-time targets, and demand stream. We have demonstrated the robustness of this model to changes in the parameter values and to variations in the demand mix. Although the use of approximations leaves open the possibility that the resulting scheduling policy is suboptimal, we have clearly demonstrated that the policies derived through this research perform significantly better than current practice, and maintain wait times within the required targets without recourse to excessive amounts of overtime. Finally, our analysis demonstrates the significant impact of even a small amount of flexibility in the scheduling of the highest-priority class. This is contrary to current practice, which tends to look for flexibility in the scheduling of the lowest-priority class instead. Such a strategy provides only temporary relief, with no long-term impact on the ability of a resource manager to maintain wait-time targets.

We trust that this research will prove useful to health-care policy makers as they seek to maintain reasonable wait times for diagnostic services.

Appendix. Proof of the Form of the Optimal Linear Value Function Approximation with Overtime

Presented in the appendix is a proof of the form of the optimal linear value function approximation for the scenario where the diversion costs are constant (overtime). The proof for the case where $d(i)$ is decreasing in i follows the same steps and is omitted. The wait-time targets, $T(i)$, for each priority class are assumed to increase with i because a high-priority patient is, by definition, a patient who must be served sooner. For completeness, we restate the theorem before giving the proof.

Restating the Theorem

THEOREM 1. *Assume that the cost of overtime, $d(i)$, is the same for all priority classes, $T(i)$ is strictly decreasing*

in i and the booking costs are nondecreasing in n and nonincreasing in i , with $b(i, n) = 0$ for all $n \leq T(i)$. If

$$b(i, n) + \gamma^{n-T(1)}d > b(i, T(i)) + \gamma^{T(i)-T(1)}d \quad (29)$$

for all $n > T(i)$ and for all i ,

$$\sum_{i=1}^I \frac{\gamma^{T(i)-n}}{1-\gamma} I_{T(i)>n} \lambda_i + \sum_{m=n}^N \gamma^{[m-n]^+} E_\alpha[X_n] < \frac{C_1}{1-\gamma} \quad (30)$$

for all $n \geq T(1)$, and

$$0 < \sum_{i=1}^I \frac{\gamma^{T(i)-T(1)}}{1-\gamma} \lambda_i + \sum_{n=1}^N \gamma^{[n-T(1)]^+} E_\alpha[X_n] - T(1)C_1 - \frac{\gamma C_1}{1-\gamma} < \frac{C_2}{1-\gamma}, \quad (31)$$

where λ_i is the arrival rate for demand from priority class i ; C_1 is the base capacity, C_2 is the diverted capacity (i.e., overtime), and γ is the discount rate, then the optimal approximate value function amongst all linear approximations for the discounted MDP will have the following form:

$$V_n = \begin{cases} d & \text{for all } n \in \{1, \dots, T(1)\}, \\ \gamma V_{n-1} & \text{for all } n \in \{T(1) + 1, \dots, N - 1\}, \\ 0 & \text{for } n = N, \end{cases} \quad (32)$$

$$W_i = V_{T(i)} \quad \text{for all } i \in \{1, \dots, I\},$$

$$W_0 = d \left(\gamma \sum_{i=1}^I \frac{\gamma^{T(i)-T(1)}}{1-\gamma} \lambda_i - T(1)C_1 - \frac{\gamma C_1}{1-\gamma} \right).$$

The Proof

The outline of the proof is as follows.

1. Prove the primal feasibility of the proposed solution.
2. Determine necessary and sufficient conditions under which a dual solution, together with the proposed primal solution, would satisfy complementary slackness.
3. Demonstrate that there exists a dual solution satisfying the necessary and sufficient conditions.

The existence of a dual solution that together with the proposed primal solution satisfies complementary slackness is sufficient to prove the optimality.

Proving Primal Feasibility

We begin by proving the feasibility of the hypothesized primal solution. Clearly, it gives nonnegative values for \vec{V} and \vec{W} . With a little algebraic manipulation, the constraint for the primal LP can be written as

$$(1-\gamma)W_0 \leq \sum_{n=1}^N \left(\sum_{i=1}^I [b(i, n) + \gamma V_{n-1} - f(i) - \gamma W_i] a_{i,n} + [\gamma V_{n-1} - V_n] x_n \right) + \sum_{i=1}^I ([d - f(i) - \gamma W_i] z_i + [f(i) + \gamma W_i - W_i] y_i) + \gamma \sum_{i=1}^I W_i \lambda_i \quad \forall (\vec{a}, \vec{z}) \in A_{(\vec{x}, \vec{y})} \text{ and } (\vec{x}, \vec{y}) \in S, \quad (33)$$

where $\lambda_i = E[Y_i]$. To prove primal feasibility, it suffices to demonstrate that the above definition for W_0 satisfies Equation (33) when the right-hand side (RHS) is at its most negative. By Equation (29) with $n = T(i) + 1$, the coefficient for y_i is positive for all i because W_i is assumed to equal $\gamma^{T(i)-T(1)}d$. Thus, the state-action pair minimizing the RHS of the above equation will always satisfy the condition $y_i = \sum_{n=1}^N a_{in} + z_i$ for all i . Substituting the hypothesized solution into the primal constraint as well as replacing y_i with $\sum_{n=1}^N a_{in} + z_i$ for all i yields

$$\begin{aligned} (1-\gamma)W_0 \leq & -\sum_{i=1}^I \gamma^{T(i)-T(1)} da_{i1} \\ & + \sum_{n=2}^N \sum_{i=1}^I (b(i,n) + (\gamma^{[n-T(1)-1]^+ + 1} - \gamma^{T(i)-T(1)})d) a_{in} \\ & - dx_1 - \sum_{n=2}^{T(1)} (1-\gamma)dx_n + \sum_{i=1}^I (d - \gamma^{T(i)-T(1)}d)z_i \\ & + \sum_{i=1}^I \gamma^{T(i)-T(1)+1} d\lambda_i \end{aligned} \quad (34)$$

for all $(\vec{a}, \vec{z}) \in A_{\vec{x}, \vec{y}}$ and $(\vec{x}, \vec{y}) \in S$. To determine which state-action pair corresponds to the most negative RHS, one can take each day and priority class separately.

For $n = 1$, the coefficient for x_1 in Equation (34) is at least as negative as the coefficient for a_{i1} with equality only if $i = 1$. Thus, any state-action pair minimizing the RHS will have $x_1 + a_{11} = C_1$.

For $n = 2, \dots, T(1)$, $b(i, n) = 0$ for all i . Thus, by inspecting the coefficient for x_n , any state-action pair minimizing the RHS will have $x_n = C_1$ if

$$(1-\gamma) \geq \gamma^{T(i)-T(1)} - \gamma,$$

which is clearly true for all i with equality for $i = 1$. Thus, any state-action pair that minimizes the RHS will have $x_n + a_{1n} = C_1$ for all $n \leq T(1)$.

For $n > T(1)$, the coefficient for a_{in} in Equation (34) equals zero at $n = T(i)$ for each $i \geq 2$. If $n < T(i)$, the coefficient is clearly positive for each priority class because $\gamma < 1$. On the other hand, if $n > T(i)$, then Equation (29) insures that the coefficient is also positive. Thus, any state-action pair minimizing the RHS of Equation (34) will have $a_{in} = 0$ for all $n > T(1)$ such that $n \neq T(i)$.

The z_i coefficient in Equation (34) implies that any state-action pair minimizing the RHS will have $z_i = 0$ for all $i > 1$, and that, again, we are indifferent in the case of $i = 1$ (although z_1 is constrained by the overtime limit).

Therefore, Equation (34) is equivalent to

$$\begin{aligned} W_0 \geq & \frac{1}{1-\gamma} \left(-dC_1 - (T(1)-1)(1-\gamma)dC_1 \right. \\ & \left. + \sum_{i=1}^I \gamma^{T(i)-T(1)+1} d\lambda_i \right) \\ \geq & d \left(\gamma \sum_{i=1}^I \frac{\gamma^{T(i)-T(1)}}{1-\gamma} \lambda_i - T(1)C_1 - \frac{\gamma C_1}{1-\gamma} \right), \end{aligned}$$

which is true with equality for the assumed value of W_0 . Therefore, if we can prove the existence of a feasible dual solution, $X(\vec{x}, \vec{y}, \vec{a}, \vec{z})$, that together with the proposed primal solution satisfies complementary slackness, then the optimality of the proposed primal solution will be demonstrated.

Conditions for a Dual Solution to Satisfy Complementary Slackness

The proof thus far has demonstrated that the proposed primal solution will have tight constraints only for those state-action pairs satisfying the following conditions:

- $x_n + a_{1n} = C_1$ for all $n \leq T(1)$,
- $z_i = 0$ for all $i > 1$,
- $a_{1n} = 0$ for all $n > T(1)$,
- $a_{in} = 0$ for all $i > 1, n \neq T(i)$, and
- $\sum_{n=1}^N a_{in} + z_i = y_i$ for all i .

Thus, to prove that the proposed primal solution is optimal, complementary slackness states that we must show the existence of a dual feasible solution that is zero for all state-action pairs that do not satisfy the above conditions and for which all the dual constraints are tight (because all primal variables are nonzero). To ease the proof, we impose the further restriction that a dual variable is positive for a given state-action pair only if x_n and a_{in} equal either zero or C_1 for all i and n .

Let $(\mathbf{s}, \mathbf{a}) = (\vec{x}, \vec{y}, \vec{a}, \vec{z})$ represent an arbitrary, feasible state-action pair and let $B = \{(\mathbf{s}, \mathbf{a}) \mid X(\mathbf{s}, \mathbf{a}) > 0\}$. Recall that the dual constraints have the form

$$(1-\gamma) \sum_{(\mathbf{s}, \mathbf{a}) \in B} X(\mathbf{s}, \mathbf{a}) = 1, \quad (35)$$

$$\sum_{(\mathbf{s}, \mathbf{a}) \in B} X(\mathbf{s}, \mathbf{a}) \left(x_n - \gamma x_{n+1} - \gamma \sum_{i=1}^I a_{i, n+1} \right) \geq E_\alpha[X_n] \quad \forall n = 1, \dots, N, \quad (36)$$

$$\sum_{(\mathbf{s}, \mathbf{a}) \in B} X(\mathbf{s}, \mathbf{a}) (y_i - \gamma \lambda_i) \geq \lambda_i \quad \forall i = 1, \dots, I, \quad (37)$$

$$X \geq 0 \quad \forall (\mathbf{s}, \mathbf{a}) \in S \times A. \quad (38)$$

Equation (37) has been simplified by recognizing that for all $X \in B$, $\sum_{n=1}^N a_{in} + z_i = y_i$. For $n = N$, Equation (36) yields

$$\sum_{(\mathbf{s}, \mathbf{a}) \in B} X(\mathbf{s}, \mathbf{a}) (x_N) = E_\alpha[X_N] \Rightarrow \sum_{\substack{(\mathbf{s}, \mathbf{a}) \in B \\ x_N > 0}} X(\mathbf{s}, \mathbf{a}) = \frac{E_\alpha[X_N]}{C_1}.$$

For $n = N - 1$, Equation (36) yields

$$\begin{aligned} \sum_{(\mathbf{s}, \mathbf{a}) \in B} X(\mathbf{s}, \mathbf{a}) (x_{N-1} - \gamma x_N) &= E_\alpha[X_{N-1}] \\ \Rightarrow \sum_{(\mathbf{s}, \mathbf{a}) \in B} X(\mathbf{s}, \mathbf{a}) (x_{N-1}) - \gamma \sum_{(\mathbf{s}, \mathbf{a}) \in B} X(\mathbf{s}, \mathbf{a}) (x_N) &= E_\alpha[X_{N-1}] \end{aligned}$$

$$\begin{aligned} \Rightarrow \sum_{\substack{(\mathbf{s}, \mathbf{a}) \in B \\ x_{N-1} > 0}} X(\mathbf{s}, \mathbf{a}) &= \frac{E_\alpha[x_{N-1}]}{C_1} + \gamma \sum_{\substack{(\mathbf{s}, \mathbf{a}) \in B \\ x_N > 0}} X(\mathbf{s}, \mathbf{a}) \\ &= \frac{1}{C_1} (E_\alpha[x_{N-1}] + \gamma E_\alpha[N]). \end{aligned}$$

Proceeding similarly, for arbitrary $n > T(I) - 1$, Equation (36) yields

$$\begin{aligned} \sum_{\substack{(\mathbf{s}, \mathbf{a}) \in B \\ x_n > 0}} X(\mathbf{s}, \mathbf{a}) &= \frac{E_\alpha[x_n]}{C_1} + \gamma \sum_{\substack{(\mathbf{s}, \mathbf{a}) \in B \\ x_{n+1} > 0}} X(\mathbf{s}, \mathbf{a}) \\ &= \frac{1}{C_1} \sum_{m=n}^N \gamma^{m-n} E_\alpha[x_m]. \end{aligned}$$

For $n = T(I) - 1$, there is the added complication that bookings may be made on day $n + 1$. In this case, Equation (36) yields

$$\begin{aligned} \sum_{(\mathbf{s}, \mathbf{a}) \in B} X(\mathbf{s}, \mathbf{a}) (x_{T(I)-1} - \gamma x_{T(I)} - \gamma a_{I, T(I)}) &= E_\alpha[X_{T(I)-1}] \\ \Rightarrow C_1 \sum_{\substack{(\mathbf{s}, \mathbf{a}) \in B \\ x_{T(I)-1} > 0}} X(\mathbf{s}, \mathbf{a}) - \gamma C_1 \sum_{\substack{(\mathbf{s}, \mathbf{a}) \in B \\ x_{T(I)} > 0}} X(\mathbf{s}, \mathbf{a}) \\ - \gamma C_1 \sum_{\substack{(\mathbf{s}, \mathbf{a}) \in B \\ a_{I, T(I)} > 0}} X(\mathbf{s}, \mathbf{a}) &= E_\alpha[X_{T(I)-1}]. \end{aligned}$$

Because $z_I = 0$ and $a_{In} = 0$ for all $n \neq T(I)$, Equation (37) yields

$$\begin{aligned} C_1 \sum_{\substack{(\mathbf{s}, \mathbf{a}) \in B \\ a_{I, T(I)} > 0}} X(\mathbf{s}, \mathbf{a}) &= \sum_{(\mathbf{s}, \mathbf{a}) \in B} X(\mathbf{s}, \mathbf{a}) y_I = \frac{1}{1-\gamma} \lambda_I \\ \Rightarrow \sum_{\substack{(\mathbf{s}, \mathbf{a}) \in B \\ a_{I, T(I)} > 0}} X(\mathbf{s}, \mathbf{a}) &= \left(\frac{1}{1-\gamma} \right) \frac{\lambda_I}{C_1}. \end{aligned}$$

Thus,

$$\begin{aligned} \sum_{\substack{(\mathbf{s}, \mathbf{a}) \in B \\ x_{T(I)-1} > 0}} X(\mathbf{s}, \mathbf{a}) \\ &= \frac{1}{C_1} \left(E_\alpha[X_{T(I)-1}] + \frac{\gamma}{1-\gamma} \lambda_I \right) + \gamma \sum_{\substack{(\mathbf{s}, \mathbf{a}) \in B \\ x_{T(I)} > 0}} X(\mathbf{s}, \mathbf{a}) \\ &= \frac{1}{C_1} \left(\sum_{m=T(I)-1}^N \gamma^{m-T(I)+1} E_\alpha[x_m] + \frac{\gamma}{1-\gamma} \lambda_I \right). \end{aligned}$$

Proceeding similarly, for $n \geq T(1)$,

$$\begin{aligned} \sum_{\substack{(\mathbf{s}, \mathbf{a}) \in B \\ x_n > 0}} X(\mathbf{s}, \mathbf{a}) \\ &= \frac{1}{C_1} \left(\sum_{m=n}^N \gamma^{m-n} E_\alpha[X_m] + \frac{1}{1-\gamma} \sum_{i=1}^I \gamma^{[T(i)-n]^+} I_{T(i) > n} \lambda_i \right). \quad (39) \end{aligned}$$

Finally, for $n < T(1)$, a_{1n} may be nonzero. Therefore, Equation (36) yields

$$\begin{aligned} \sum_{(\mathbf{s}, \mathbf{a}) \in B} X(\mathbf{s}, \mathbf{a}) (x_n - \gamma x_{n+1} - \gamma a_{1, n+1}) &= E_\alpha[X_n] \\ \Rightarrow C_1 \sum_{\substack{(\mathbf{s}, \mathbf{a}) \in B \\ x_n > 0}} X(\mathbf{s}, \mathbf{a}) \\ &= E_\alpha[X_n] + \gamma C_1 \left(\sum_{\substack{(\mathbf{s}, \mathbf{a}) \in B \\ x_{n+1} > 0}} X(\mathbf{s}, \mathbf{a}) + \sum_{\substack{(\mathbf{s}, \mathbf{a}) \in B \\ a_{1, n+1} > 0}} X(\mathbf{s}, \mathbf{a}) \right). \end{aligned}$$

However, not all priority 1 patients are necessarily booked on the same day. This would make the system unsolvable but for the fact that we know that $x_{n+1} + a_{1, n+1} = C_1$ for all $n \in \{0, \dots, T(1) - 1\}$ and for all $(\mathbf{s}, \mathbf{a}) \in B$. Hence, Equation (35) yields

$$\sum_{\substack{(\mathbf{s}, \mathbf{a}) \in B \\ a_{1, n+1} > 0}} X(\mathbf{s}, \mathbf{a}) + \sum_{\substack{(\mathbf{s}, \mathbf{a}) \in B \\ x_{n+1} > 0}} X(\mathbf{s}, \mathbf{a}) = \frac{1}{1-\gamma}.$$

Therefore,

$$\sum_{\substack{(\mathbf{s}, \mathbf{a}) \in B \\ x_n > 0}} X(\mathbf{s}, \mathbf{a}) = \frac{E_\alpha[X_n]}{C_1} + \frac{\gamma}{1-\gamma} \quad \text{for } n < T(1). \quad (40)$$

Thus, Equations (39) and (40) give the required dual weight for each day to satisfy complementary slackness and dual feasibility. All weights are positive and none are greater than the total available weight $1/(1-\gamma)$ (from Equation (35)). To determine the required dual weight for saturation pairs with positive z_1 , recall that Equation (37) yields

$$\frac{\lambda_1}{1-\gamma} = \sum_{(\mathbf{s}, \mathbf{a}) \in B} X(\mathbf{s}, \mathbf{a}) y_1.$$

Therefore,

$$\begin{aligned} \sum_{(\mathbf{s}, \mathbf{a}) \in B} X(\mathbf{s}, \mathbf{a}) z_1 \\ &= \frac{\lambda_1}{1-\gamma} - \sum_{n=1}^{T(1)} \left(\sum_{(\mathbf{s}, \mathbf{a}) \in B} X(\mathbf{s}, \mathbf{a}) a_{1, n} \right) \\ &= \frac{\lambda_1}{1-\gamma} - C_1 \sum_{n=1}^{T(1)} \left(\frac{1}{1-\gamma} - \sum_{\substack{(\mathbf{s}, \mathbf{a}) \in B \\ x_n = C_1}} X(\mathbf{s}, \mathbf{a}) \right) \\ &= \frac{\lambda_1}{1-\gamma} - \frac{C_1 T(1)}{1-\gamma} + C_1 \left(\sum_{n=1}^{T(1)-1} \left[\frac{E_\alpha[X_n]}{C_1} + \frac{\gamma}{1-\gamma} \right] \right) \\ &\quad + C_1 \left(\sum_{n=T(1)}^N \gamma^{n-T(1)} \frac{E_\alpha[X_n]}{C_1} + \sum_{i=2}^I \gamma^{T(i)-T(1)} \left(\frac{1}{1-\gamma} \right) \frac{\lambda_i}{C_1} \right) \\ &= \frac{\lambda_1}{1-\gamma} - \frac{C_1 T(1)}{1-\gamma} + \frac{\gamma C_1 (T(1)-1)}{1-\gamma} \\ &\quad + \sum_{n=1}^N \gamma^{[n-T(1)]^+} E_\alpha[X_n] + \sum_{i=2}^I \frac{\gamma^{T(i)-T(1)}}{1-\gamma} \lambda_i \\ &= \sum_{i=1}^I \frac{\gamma^{T(i)-T(1)}}{1-\gamma} \lambda_i - T(1) C - \frac{\gamma C_1}{1-\gamma} + \sum_{n=1}^N \gamma^{[n-T(1)]^+} E_\alpha[X_n]. \end{aligned}$$

Thus,

$$\sum_{(\mathbf{s}, \mathbf{a}) \in B} X(\mathbf{s}, \mathbf{a}) z_1 = \sum_{i=1}^I \frac{\gamma^{T(i)-T(1)}}{1-\gamma} \lambda_i - T(1) C_1 - \frac{\gamma C_1}{1-\gamma} + \sum_{n=1}^N \gamma^{[n-T(1)]^+} E_\alpha[X_n], \quad (41)$$

which is greater than zero and less than $C_2/(1-\gamma)$ by Equation (31) and therefore dual feasible. It is worth noting that simply for dual feasibility, we would only need the above to be greater than or equal to zero. However, if Equation (41) was identically zero, then we would have degeneracy and the proposed primal solution would no longer be unique. In fact, the optimal objective function would be zero, suggesting that an approximate value function identical to zero would be optimal.

The Existence of a Dual Solution Satisfying Complementary Slackness

The above argument suggests a weighting scheme for a dual feasible solution that, together with the proposed primal solution, satisfies complementary slackness. It remains to prove that a dual solution satisfying the above weighting scheme must exist. We can determine the state-action pairs with positive dual weight starting on day 1 and working up to day N . Because we imposed the condition that x_n and a_{in} are equal to either zero or C_1 , it follows that such a dual solution exists if

1. The total dual weight available $1/(1-\gamma)$ given by Equation (35) does not exceed the combined weight assigned to all states where x_n or $\sum_{i=1}^I a_{in}$ are greater than zero for any n .

2. The total weight assigned to dual variables where $\sum_{n=1}^N a_{in} + z_i$ is positive is equal to the weight assigned to dual variables where y_i is positive.

This turns out to be straightforward because it is easy to show that under the above weighting scheme and using Equation (30),

$$\sum_{x_n > 0} X(\mathbf{s}, \mathbf{a}) x_n + \sum_{i=1}^I \sum_{a_{in} > 0} X(\mathbf{s}, \mathbf{a}) a_{in} \leq \frac{C_1}{1-\gamma}$$

$$\Rightarrow \sum_{x_n > 0} X(\mathbf{s}, \mathbf{a}) + \sum_{i=1}^I \sum_{a_{in} > 0} X(\mathbf{s}, \mathbf{a}) \leq \frac{1}{1-\gamma}$$

for all $n \in \{1, \dots, N\}$ and

$$\sum_{n=1}^N \sum_{a_{in} > 0} X(\mathbf{s}, \mathbf{a}) a_{in} + \sum_{z_i > 0} X(\mathbf{s}, \mathbf{a}) z_i = \sum_{y_i > 0} X(\mathbf{s}, \mathbf{a}) y_i$$

for all $i \in \{1, \dots, I\}$. Thus, there exist admissible state-action pairs that satisfy the above weighting scheme, proving the existence of the required dual solution.

Acknowledgments

This research was partially supported by NSERC grant OGP0005527. The authors thank Dan Adelman for many helpful comments through the course of their research and Antoine Saure for providing exact solutions to the small instances of the problem and for assistance with the literature review.

References

Adelman, D. 2004. A price-directed approach to stochastic inventory/routing. *Oper. Res.* **52** 499–514.

Adelman, D. 2005. Dynamic bid-prices in revenue management. *Oper. Res.* **55** 647–661.

Adelman, D. 2006. Weakly coupled stochastic dynamic programs. Personal communication.

Bertsimas, D., I. Popescu. 2003. Revenue management in a dynamic network environment. *Transportation Sci.* **37** 257–277.

Bertsimas, D., R. Shioda. 2003. Restaurant revenue management. *Oper. Res.* **51** 472–486.

Brumelle, S., D. Walczak. 2003. Dynamic airline revenue management with multiple semi-Markov demand. *Oper. Res.* **51** 137–148.

Cayirli, T., E. Veral. 2003. Outpatient scheduling in health care: A review of literature. *Production Oper. Management* **12** 519–549.

Collart, D., A. Haurie. 1976. On a suboptimal control of a hospital inpatient admission system. *IEEE Trans. Automatic Control* **21**(2) 233–238.

de Farias, D., B. Van Roy. 2003. The linear programming approach to approximate dynamic programming. *Oper. Res.* **51** 850–865.

de Farias, D., B. Van Roy. 2004a. A cost-shaping linear program for average-cost approximate dynamic programming with performance guarantees. *Math. Oper. Res.* **29** 462–478.

de Farias, D., B. Van Roy. 2004b. On constraint sampling in the linear programming approach to approximate dynamic programming. *Math. Oper. Res.* **29** 462–478.

Denton, B., D. Gupta. 2003. A sequential bounding approach for optimal appointment scheduling. *IIE Trans.* **35** 1003–1016.

Gerchak, Y., D. Gupta, M. Henig. 1996. Reservation planning for elective surgery under uncertain demand for emergency surgery. *Management Sci.* **42** 321–334.

Green, L., S. Savin, B. Wang. 2006. Managing patient demand in a diagnostic medical facility. *Oper. Res.* **54** 11–25.

Gupta, D., L. Wang. 2008. Revenue management for a primary-care clinic in the presence of patient choice. *Oper. Res.* **56** 576–592.

Klabjajan, D., D. Adelman. 2007. An infinite dimensional linear programming algorithm for deterministic semi-Markov decision processes on Borel spaces. Personal communication.

Kolesar, P. 1970. A Markovian model for hospital admission scheduling. *Management Sci.* **16** 384–396.

MacCormick, A., B. Parry. 2003. Waiting time thresholds: Are they appropriate? *ANZ J. Surgery* **73** 926–928.

Magerlein, J., J. Martin. 1978. Surgical demand scheduling: A review. *Health Services Res.* **13** 418–433.

Mondschein, S., G. Weintraub. 2003. Appointment policies in service operations: Critical analysis of the economic framework. *Production Oper. Management* **12** 266–286.

Odegaard, F., L. Chen, R. Quee, M. Puterman. 2007. Improving the efficiency of hospital porter services, part 1: Study objectives and results. *J. Healthcare Quart.* **29** 4–11.

- Patrick, J., M. Puterman. 2007. Improving resource utilization for diagnostic services through flexible inpatient scheduling. *J. Oper. Res. Soc.* **58** 235–245.
- Patrick, J., M. Puterman. 2008. Reducing waiting time through operations research: Optimizing the use of surge capacity. *Healthcare Policy* **3** 75–88.
- Puterman, M. 1994. *Markov Decision Processes*. John Wiley and Sons, New York.
- Rising, E., R. Baron, B. Averill. 1973. A systems analysis of a university-health-service outpatient clinic. *Oper. Res.* **21**(5) 1030–1047.
- Sanmartin, C., F. Gendron, J.-M. Berthelot, K. Murphy. 2004. Access to health care services in Canada, 2003. Statistics Canada, Catalogue 82-575-XIE.
- Schweitzer, P., A. Seidman. 1985. Generalized polynomial approximations in Markovian decision processes. *J. Math. Anal. Appl.* **110** 568–582.
- van Ryzin, G., G. Vulcano. 2008. Simulation-based optimization of virtual nesting controls for network revenue management. *Oper. Res.* **56**(4) 865–880.