A Markov Decision Process with an ADP-based solution for MRI appointment scheduling in Rijnstate

---

Sander Dijkstra

9 maart 2020

UNIVERSITY OF TWENTE.

**Faculty of Electrical Engineering,
Mathematics & Computer Science**

A Markov Decision Process with an ADP-based solution for MRI appointment scheduling in Rijnstate

Sander Dijkstra
M.Sc. Thesis
February 2020

UT Supervisor:
Dr. Ir. Aleida Braaksma

Rijnstate Supervisors:
Dr. Milan Pijl
Dr. Frank Joosten

Graduation committee:
Dr. Ir. Aleida Braaksma (UT)
Prof. Dr. Richard Boucherie (UT)
Dr. Judith Timmer (UT)
Dr. Matthias Walter (UT)
Dr. Milan Pijl (Rijnstate)

Stochastic Operations Research Group
Faculty of Electrical Engineering,
Mathematics and Computer Science
University of Twente
P.O. Box 217
7500 AE Enschede
The Netherlands

Preface

During the past months I have been working on this research with the objective of concluding my study in Applied Mathematics at the University of Twente. For every EC I had to work hard, sometimes because the course was difficult and sometimes because I was giving myself a hard time. I am proud of what I achieved over the last years.
During the research period, I tried to help the Radiology department of Rijnstate to gain insight into the way in which they schedule MRI appointments and I would be honored if I helped you even the smallest bit with this.
I would not have achieved both of these without the help of many people.

First of all I would like to thank my supervisor, Aleida. While I was working on my thesis, we had a lot of very good discussions. Afterwards I was always full of new ideas. Without Aleida's enormous knowledge, enthusiasm and realistic view of the problem at hand, this thesis would not have been what it is, so *tige tank dêrfoar!*

Secondly, I would like to thank my supervisors in Rijnstate, Frank and Milan. I cannot remember a moment when you were not ready to arrange something for me or answer a question. You were always ready to assist me and your enthusiasm for my research provided enormous motivation. Thanks! By the way, I want to say that this enthusiasm came not only from both of you, but from everyone within the department. I received positive reactions from all radiologists, lab technicians and administrative staff and everyone was willing to help me.

Needless to say I would like to thank Richard, Judith and Matthias for their time spent reading and evaluating this work. In particular, I would like to thank Richard. First of all for encouraging me to work on this project. Besides that, also for supervising me throughout my (lengthy) internship period prior to this project. Thank you for your time, effort and patience invested in me.

In the long line of people I want to thank, my parents come first of course. I would like to thank you, Jurjen en Willie, for making it possible for me to study, and for supporting me in everything I did, even though sometimes you may have had no clue what that was exactly. Special thanks go to my brother René, who I also see as my best friend, for the memorable moments we shared in the past years.
Last, but not least I would like to thank all of my friends, both in my home village Oentsjerk and those I met during my study period in Enschede. Special thanks go to Andor, David, Douwe, Hylke, Frank, Pascal, Rico, Rudmer, Sagy, Sven and Symen.

Abstract

Efficient patient scheduling has significant operational, clinical and economical benefits on health care systems by not only increasing the timely access of patients to care but also reducing costs. However, patient scheduling is complex due to, among other aspects, the existence of multiple priority levels, the presence of patient type-resource compatibility constraints, (highly) variable demand and limited capacity. These aspects of patient scheduling make it extremely difficult for a booking agent to manually assess the impact of his/her decisions in order to more efficiently allocate capacity. We present a near-online method to dynamically schedule patients with different access time targets to one of the MRI scanners in hospital Rijnstate in Arnhem, taking into account patient type-resource compatibility constraints and future appointment requests. The goal is to identify effective ways of allocating available service capacity to incoming appointment requests while minimizing the number of patients whose access time exceeds the prespecified, priority-specific target in a cost-effective manner. We formulate this problem as a discounted infinite-horizon Markov Decision Process (MDP). Because the state space is too large for a direct solution, we solve the equivalent linear program through Approximate Dynamic Programming (ADP) to obtain an Approximate Optimal Policy (AOP). Here we use an affine architecture to approximate the value function of the MDP and solve the equivalent linear program through column generation. Using simulation, we compare the performance of the resulting AOP to both easy-to-use rule-based scheduling approaches and approaches based on current patient scheduling practice in Rijnstate for the practical example based on data provided by the Radiology department of Rijnstate. The results indicate that the AOP outperforms the rule-based scheduling approaches in several scenarios. At the same time we realize that, based on the results, the AOP may not deliver the desired result in all scenarios. That is why we also present an extensions of the MDP model.

Keywords: Advanced patient scheduling/Advanced capacity planning, Markov decision process, Approximate dynamic programming, Linear programming, Column generation, Simulation

Samenvatting

Het efficient plannen van patiëntenafspraken heeft aanzienlijke voordelen voor de gezondheidszorg door niet alleen de tijdige toegang van patiënten tot goede zorg te verzorgen, maar ook de kosten ervan te verlagen. Patiëntenplanning is echter complex vanwege het bestaan van meerdere prioriteiten (in termen van acceptable toegangstijden) onder de patiënten, (sterk) variabele vraag en beperkte capaciteit. Daarnaast heeft medische apparatuur ook vaak compatibiliteitsbeperkingen: niet ieder type MRI onderzoek kan gedaan worden op een willekeurige MRI scanner. Voor bepaalde types MRI onderzoeken is een specifiek type MRI scanner nodig (gekenmerkt door de sterkte van het magneetveld dat gegenereerd wordt door de MRI scanner). Deze aspecten van patiëntenplanning maken het voor een planningsmedewerker uiterst moeilijk om beschikbare capaciteit efficiënt toe wijzen aan patiëntgroepen en om de impact van zijn/haar beslissingen handmatig te beoordelen. Het gedane scriptieonderzoek presenteert een methode om patiënten, die verschillende toegangstijdnormen hebben, te plannen op één van de MRI-scanners in ziekenhuis Rijnstate te Arnhem. Hierbij wordt rekening gehouden met de zojuist uitgelegde compatibiliteitsbeperkingen en toekomstige afspraakverzoeken. Het doel is om een effectieve manier te vinden om beschikbare MRI-capaciteit toe te wijzen aan inkomende afspraakverzoeken, terwijl het aantal patiënten waarvan de toegangstijd de vooraf gespecificeerde norm overschrijdt wordt geminimaliseerd. Het doel is om dit op een kosteneffectieve manier te doen. Hiermee wordt bedoeld dat we het liefst zo weining mogelijk overwerktijd voor personeel en MRI-scanner genereren en proberen de bezettingsgraad van de MRI scanners te maximaliseren. We formuleren dit probleem als een *Markov Decision Process (Markov beslissingsprobleem)* (MDP). Omdat de toestandsruimte van dit MDP te groot is voor een directe oplossing, kiezen we een *Approximate Dynamic Programming* (ADP) methode om planningsregels te vinden die bijna-optimaal zijn. Dit wil zeggen: we zoeken planningsregels die mogelijkerwijs wel, maar misschien ook niet, aan de wiskundige criteria voldoen om optimaal te zijn, maar hopelijk wel goed werken in de praktijk. Of ze goed werken in de praktijk testen we met behulp van een computersimulatie. Hierin vergelijken we de prestaties van de gevonden planningsregels met eenvoudig te gebruiken planningsregels (zoals iedere patiënt boeken in het eerstvolgende beschikbare tijdslot) en de momenteel gebruikte aanpak op de afdeling Radiologie van Rijnstate. De resultaten suggereren dat de planningsregels die we gevonden hebben met het ADP algoritme beter presteren dan de eenvoudig te gebruiken planningsregels in verschillende scenario's. Tegelijkertijd realiseren we ons dat, op basis van de resultaten, de *ADP-planningsregels* mogelijk niet in alle scenario's het gewenste resultaat oplevert. Daarom presenteren we ook een uitbreiding van het MDP-model.

Contents

Glossary

The next list describes (technical) terms that will be later used throughout this thesis.

Access time/Indirect waiting time	Time between when a patient requests an appointment and the scheduled appointment time.
Access time window/target	Period within which a patient preferably has (or must) his/her (first) appointment scheduled.
Appointment day	(Future) day that an individual patient is scheduled to receive service.
Appointment time	Start time that an individual patient is scheduled to receive service on the appointment day.
Appointment slot (time slot)	Smallest time window in which one customer can be scheduled.
Appointment system	A system that plans and schedules appointment requests to deliver timely and convenient access to health services for all patients.
Booking horizon/ Appoint. scheduling window	Period how far into the future an appointment can be scheduled.
Cancellation	A situation where a patient cancels his/her appointment far enough in advance to allow for a new appointment to be substituted.
Consultation session	The time period available for serving patients.
(Direct) waiting time	The delay between a patient's appointment time and the start of service. Note that if service starts before the scheduled appointment time, a waiting time of 0 time units is counted.
Inpatient	A patient who stays/lives in hospital while under examination or treatment.
No-show patient	A patient who does not show up for his/her appointment and does not give prior notice.
Offline scheduling approach	Scheduling approach in which appointments are scheduled after all requests have arrived.
Online scheduling approach	Scheduling approach in which patients are scheduled immediately upon the arrival of their request.
Outpatient	A patient who attends the hospital for examination or treatment without staying there overnight.

Outpatient clinic	A health facility that provides care to patients that do not need an overnight stay.
Overtime	The positive difference between the desired completion time of the clinic session and the actual end of the service for the last patient.
Pre-scheduled patients	Patients whose appointment is scheduled in advance of their appointment days.
Regular walk-in patients	Walk-in patients who do not require urgent treatment.
Response time	Time it takes to respond to an appointment request, i.e., the time between the request is known to the booking clerk and the moment the appointment is booked
Same-day patients	Patients whose appointment is scheduled on the same day that they call for an appointment.
Scheduled patients	Patients who make an appointment before arriving at the clinic.
Server idle time	Part of the consultation session that the server is idle due to lack of patient(s).
Urgent walk-in/ Emergency patient	Walk-in patients who need treatment as soon as possible and take priority over other patient types.
Walk-in patients	Patients who arrive at the clinic without an appointment during the consultation session.

1

Introduction

“Managers make resource allocation decisions, but doctors decide what the hospital does with those resources” [8]. This statement is more then fifteen years old, but still considered relevant in some hospitals nowadays. While doctors focus on treating each individual patient as well as they possibly can, managers also focus on optimal usage of resources. Dilemma’s such as what to do if the treatment of a single cancer patient costs 60K Euros, while three other patients suffering from cardiovascular disease can be treated for 20K Euros each? Or a more specific example that shows the difficulty of resource allocation would be the distribution of operating room time between elective and urgent (emergency) patients. It is however very common in hospitals to avoid explicit decisions on such resource allocation and capacity distribution problems and to react on ad-hoc basis to problems that occur, which may consequently result in very undesirable system outcomes, (e.g., cancellations, long (indirect) waiting times, a low utilization rate of expensive medical equipment).

An outpatient clinic is defined as a health facility that provides care to patients that do not need an overnight stay [22]. In these clinics, distribution of available capacity also occurs. For example, there is the question how to share the available capacity between same-day and pre-scheduled appointments? In this outpatient setting, appointment scheduling is an important topic that has gained increasing attention during the past years, as the appointment system (AS) is one of the hospital’s tool to deal with capacity allocation problems and also to establish their quality of service. This quality of service can be increased by providing patients with quick responses to their appointment requests; offering patients choice in the times at which they would prefer to have their appointments, and enabling them to combine multiple appointments on a single day. The appointment scheduling may also affect quality of care. Depending on their condition, patients should receive their first consultation, examination, or treatment within the appropriate access time, as the patient’s condition may deteriorate while waiting. Although in some cases this may have little medical impact, in others, excessive wait times can potentially impact health outcomes [25, 32].

Like outpatient clinics, the Radiology department of hospital Rijnstate operates an AS in which they have to deal with limited capacity and medically acceptable wait-times. In this research we focus the Rijnstate’s AS for MRI examinations. The capacity for MRI is dictated by the available resources, e.g., the number of MRI scanners, radiologists and radiologic technologists (RTs). The department faces the challenge of booking appointments, in real time, when demand comes from three patient groups: inpatient (same-day) patients, emergency patients, and pre-scheduled (out)patients. Based on the group to which a patient belongs and the type of MRI that is requested - which in turn is determined by the underlying diagnostic question - patients may be classified into priority categories with different medically acceptable wait times. For example, some conditions may require immediate diagnostic examination, whereas in other cases it may be medically acceptable to wait up to several days. Because less-urgent patients are booked further into the future, this raises the question for the hospital as to how much MRI capacity to reserve for later-arriving but higher-priority demand? From the patient's perspective it is not only very important that he/she can visit the hospital within the medically acceptable wait time, but also that he/she is quickly notified of his/her appointment.

Three aspects make scheduling of MRI appointments complex: patients can be classified into multiple types; sessions do not necessarily have the same duration; and treatments can typically be delivered on more than one MRI scanner but not on all. These aspects of MRI scheduling, together with the presence of highly variable demand and limited capacity, make it extremely difficult for a booking agent to manually assess the impact of his/her decisions in order to more efficiently allocate capacity. This unintended lack of foresight may result in several inefficiencies that typically translate into unnecessary delays, a non-systematic prioritization of patients, unused appointment slots, and excessive overtime.

This research addresses the question how to design the AS for MRI examinations in Rijnstate such that available capacity is allocated effectively to incoming appointment requests while minimizing the number of patients whose access time exceeds the pre-specified access time target. To this end, we develop a discounted infinite-horizon Markov Decision Process (MDP). This MDP provides a near-online, dynamic method for advance capacity planning involving multiple resources (MRI scanners) with different capabilities and taking into account future appointment requests. Due to the curse of dimensionality, the proposed MDP cannot be solved analytically even for small instances of the problem. For this reason, we use an Approximate Dynamic Programming (ADP) approach to approximately solve it. Here we use an affine architecture to approximate the value function of the MDP and solve an equivalent linear programming model through column generation to obtain an approximate optimal policy.

1.1 Structure of this thesis

In the next chapter, Chapter 2, we elaborate on the AS currently operated by the Radiology department of Rijnstate. We explain how this AS works and what dissatisfaction it causes. In this chapter we also explain what the requirements are for an alternative AS and what options/freedoms (i.e., which decision variables) we have when designing the alternative AS. Based on these design options, we will review available literature in Chapter 3. We have structured our literature research in such a way that we start from several literature reviews on the subject of (outpatient) appointment scheduling and to find relevant research articles that can support us in designing the AS for MRI scans in Rijnstate, we perform a backward and forward search on these literature reviews. In Chapter 4 we present the near-online, dynamic method for advance capacity planning. We formulate this planning method as a discounted infinite-horizon Markov Decision Process (MDP) by providing the decision epochs, state space, action sets, transition probabilities and costs in Section 4.1. In Section 4.2 we present the (ADP) approach where we solve an equivalent linear programming model through column generation to obtain an approximate optimal policy. In Chapter 6 we evaluate the performance of the resulting approximate optimal policy for a practical scale case study at the Radiology department of Rijnstate using simulation. The performance results are compared to the performance of benchmark policies that are commonly used in practice as well as to the policy that comes closest to Rijnstate's currently used AS. The input parameters for this practical scale case study are determined in Chapter 5 using historical data. In Chapter 5 we also report on the results of a data analysis we performed on a data set provided by the Radiology department of Rijnstate. The purpose of this data analysis is to provide a quantitative (approximate) summary of the current performances of the department. In the last chapter, Chapter **??** we state our main conclusions, but before we elaborate on several further research directions and suggest a possible extensions to our MDP scheduling model in more depth in Chapter 7.

2

The Radiology department of Rijnstate

Rijnstate is a general hospital in Gelderland, the Netherlands. The hospital's headquarters are in Arnhem and additional sites are in Velp, Zevenaar and at another location in Arnhem-South. Rijnstate performs approximately 500,000 outpatient consultations (including follow-up appointments) and 63,000 admissions per year - both daycare and clinical - and it has 766 registered beds [31]. Furthermore, Rijnstate's catchment area is approximately 450,000 inhabitants and, spread over the five locations, almost 5,000 people work at the hospital [31]. For the catchment area, Rijnstate serves as a general hospital, providing regular care. In addition to their position as general hospital, Rijnstate's main location in Arnhem has assigned a number of top-clinical functions. This means that as one of the 26 large training hospitals in the Netherlands, Rijnstate provides some medical treatments and services that are only allocated to a limited number of hospitals in view of the high costs and the required expertise[1]. Finally, an average of 80 large medical-scientific studies with the intention to identify causes of diseases and finding better treatments start every year in Rijnstate[2].

The Radiology department of Rijnstate accommodates the medical specialty that uses medical imaging techniques to diagnose and treat diseases within the human body. All sorts of diagnostic examinations are carried out on various diagnostic facilities to support other specialists in order to increase the quality of diagnosis and treatment. The department itself also performs some treatments using the diagnostic facilities, mainly under the name of interventional radiology, whereby, for example, constrictions or blockages of a blood vessel are treated with stents. Physical exams are carried out using X-rays (regular X-rays, mammography, fluoroscope or CT scan), sound waves (sonograms) and magnetic fields (MRI).

The staff of the Radiology department consists of radiologists, radiologic technologists (RTs; sometimes also called radiographers), physician assistants (PAs), administrative and support staff and junior radiologists, junior PAs and junior RTs in training. The RTs carry out most of the diagnostic examinations and thus make the majority of

[1] See `http://www.rijnstate.nl/over-rijnstate/waar-staan-we-voor/topklinische-zorg/`

[2] See `http://www.rijnstate.nl/over-rijnstate/waar-staan-we-voor/wetenschap-en-innovatie/over-onderzoek/`

the medical images. After the RT has finished the diagnostic examination, the images taken are sent to a radiologist. The radiologist assesses the images and writes a radiological report. In this report, the radiologist describes his/her findings and based on the images he/she tries to answer the diagnostic question as good as possible. Afterwards, the report is made available to other physicians in the patient's electronic health record.

One of the goals that Rijnstate has set itself for the coming years is to further optimize the arrangement of the care activities for the patients [31]. With this goal in mind, the Radiology department of Rijnstate is committed to improve the appointment system (AS) for MRI examinations. In this chapter we describe how the currently used MRI AS looks like, how this AS causes dissatisfaction among the hospital and its patient, what the requirements for a new AS are, and what design options we have for this new system. Based on the design options, we will review available literature in the next chapter.

2.1 The current MRI appointment system

MRI (Magnetic Resonance Imaging) uses strong magnetic fields to align atomic nuclei within body tissues, then uses a radio signal to disturb the axis of rotation of these nuclei and observes the radio frequency signal generated as the nuclei return to their baseline states. The radio signals are collected by small antennae placed near the area of interest. An MRI examination is usually done by two radiologic technologist. Some examinations needs to be done by a radiologist or PA. Because a patient will receive an injection of a contrast medium for some MRI examinations, a radiologist must always be present in the Radiology department when MRI examinations are done.

In the Radiology department there are three non-identical MRI scanners. The technical difference is in the strength of the magnetic fields. The first scanner generates a magnetic field of 3 Tesla (SI symbol T) and both others one of 1.5 T. The difference in use that results from this is that a selection of examinations can only be done on one of the MRI scanners. The regular times at which outpatient MRI is performed during working days is between 8 a.m. and 6 p.m. Because the demand for MRI scans is high, this period is regularly extended to 9 p.m. on Tuesdays, Wednesdays and Thursdays. During the period that we had insight into the MRI agendas, that is from 2 January 2019 to 31 August 2019, this happened on average once every two weeks. Sometimes examinations are also conducted on Saturdays. During the months just mentioned, this happened on average once a month, usually on the second Saturday of the month.

In the appointment scheduling approach for MRI scans, patients always need an appointment, except emergency patients of course. Based on the symptoms and the clinical picture of the emergency patient, the radiologist determines the time-frame within which an MRI must be made. This may imply that the emergency patient gets the highest priority in queue for a scanner that is suitable for his type of MRI examination and is examined as soon as the current service ends. However, if there is an idle period in today's planning it is also possible that the emergency patient is examined during this period as long as it is within the time-frame determined by the radiologist. The currently operating AS is based on a *blueprint agenda*. This means that it is determined in advance when what type of examination (for example a brain MRI) can be done.

Time	MRI examinations that can be done according to the blueprint agenda
08:00 - 08:30	3190, 3190D
08:30 - 08:30	1390, 1390D, 1390A, 1390E, 2090, 2290, 2990, 3090T, 3190, 3190D, 3190PB, 3290, 3390, 3390D, 3390P, 3490, 3690, 9090 L/R, 9491 L/R
09:30 - 10:10	NEURO (1390 - 1 day before)
10:10 - 10:30	TIA/TIAS (1390 - 1 day before)
10:30 - 10:50	3190, 3190D, 3390, 3390D, 3390P, 3490
10:50 - 12:30	*Emergency patients or Same-day patients*
12:30 - 13:00	2290, 2990, 3090T, 3690, 7090, 7490, 7690, 8490, 91900T, MWEKE
13:00 - 13:30	3190, 3190D, 3390, 3390D, 3390P, 3490
13:30 - 17:00	5190/5192, 5191/5191R
17:00 - 18:00	5190/5192, 5191/5191R

Figure 2.1: Example of a blueprint agenda. This blueprint agenda applies to the 3T MRI scanner (the Radiology department of Rijnstate labels it as MRI scanner 2) and was used on Thursdays during both weeks with odd week numbers and even weeks from 2 January 2019 to 26 April 2019.

If a request comes for a brain scan, the booking agent knows exactly where this can be scheduled and where not and the only choice that must be made is to book the patient into (one of) the suitable appointment slots. Figure 2.1 shows an example of a blueprint agenda of one the MRI scanners. The coding used in Figure 2.1 for the various MRI examinations comes from the currently used computer system *HIX*. Table A.1 in Appendix A contains all the HIX codes and the corresponding MRI examinations. This table also shows which MRI scanner/scanners is/are suitable for each MRI examination. As you may have argued, these slots are defined with the intention that emergency/same-day patients will not have to be squeezed into the scheduled program in the manner that they get highest priority in queue for a scanner that is suitable for his type of MRI scan.

As can be seen in Figure 2.1, the blueprint agenda can be specified up to different levels. Some appointment slots are reserved for specific HIX codes and others for specific classes of patients, for example the emergency/same-day slots. In these slots all types of MRI scans that can feasibly be done on the specific MRI scanner can be booked, as long as it is for an emergency or same-day patient. Same-day patients are usually inpatients for whom an MRI scan is requested after an examination or a morning round at the nursing ward, done by a physician.
Furthermore, as can also be seen in Figure 2.1, some of the HIX codes have the suffix *D*, for example *1390D*. Such a HIX code corresponds to the type of MRI scan with HIX code *1390*, but is booked decentrally. Decentralized appointment scheduling enables other outpatient clinics in Rijnstate to autonomously book appointments for their patients in the agenda of one of the MRI scanners. For example, the Neurology department is authorized to schedule MRI examinations of the brain at free appointment slots

for which the blueprint agenda dictates HIX code *1390D* without the involvement of a booking agent of the Radiology department.
Finally, Figure 2.1 also illustrates the protection of appointment slots for *one-stop-shop* patients. For example, the slots reserved for MRI scans with HIX code *NEURO*. This code corresponds to a brain scan identical to those specified by the HIX codes *1390* or *1390D*. However, only one-stop-patients can be scheduled during the time slots mentioned. One-stop-shop means that patients who have to undergo various examinations, for example blood tests, an electrocardiogram, and also an MRI, can undergo all of this during the same day and therefore only have to visit the hospital once. The MRI blueprint agenda protects specific time slots for these patients up to one or a few days in advance. We take the following example to illustrate this. Assume all appointment slots on Thursday morning are reserved for brain MRIs for one-stop-shop patients (that means Thursday morning is reserved for HIX code *NEURO*) and the slots are protected up to one day in advance. Then, up to Wednesday, booking agents of the Radiology department can only schedule patients for which an MRI scan with the *NEURO* code is requested on Thursday mornings. From Wednesday, the blueprint for Thursday morning is updated and MRI examinations specified by this updated blueprint can also be scheduled on Thursday morning. In the example of the one-stop-shop brain scan, the updated blueprint now allows to schedule patients for which an MRI scan with either the code *NEURO* or *1390* is requested.

There are several incentives for a blueprint agenda. The first is that there are time slots for emergency/same-day patients with a high need for an MRI examination. Here, however, the time during the day at which these blocks are placed is very important for their success. Another reason for a blueprint agenda is to create convenience for other departments within the hospital: some of them can book appointments decentrally and others have the option of creating a one-stop shop for their patients. For both aforementioned incentives, access times also play an important role: for the emergency/same-day patients their access time target is within the current day. For the one-stop-shop patients different access time targets have been agreed with the departments involved, but the idea of the protected time slots in the blueprint calendar is that these access time targets can be achieved for a larger group of patients that if the blueprint agenda was to be absent. Finally, the blueprint agenda clusters MRI exams with the same MRI scanner settings to save set-up periods between examinations.

Since the blueprint agenda determines the booking options for each MRI examination request, the only choice that must be made is to book the patient into one or more the suitable appointment slots. For patients who appear at the department's desk with an MRI appointment request, this is done in an online fashion as they immediately get an appointment. For requests that arrive otherwise (via telephone or online in HIX) the moment of booking has not been established according to a strict policy. Sometimes it happens online, but more often not. Therefore, the current AS is not an fairly online appointment system. However, it is also not a near-online system in the sense that it has fixed decision moments during the day. It is totally up to the booking agent when an appointment request is provided with an answer. The currently used assignment rule is to book a patient into the first suitable available slot. Depending on a patient's condition, it may be decided, after consultation between a radiologist and the treating physician, to deviate from the blueprint agenda if the next available suitable time slot is too far into the future. In that case, a patient is either examined during a block that is actually reserved for another type of MRI scan, or the patient is examined after the regular program, i.e, in overtime.

2.2 Requirements to an alternative appointment system

The currently used AS for MRI scans is a fairly simplistic design: a static blueprint calendar combined with an assignment rule and the flexibility to release the reservation of appointment slots in the blueprint under certain circumstances. According to staff members of the Radiology department this design of AS results in long access times for patients. In addition, as may be clear from the introduction, the problem is that the fraction of patients who are not examined within the desired access time window is considered too large. With an alternative MRI AS, the Radiology department of Rijnstate aims to examine more patients within their medically desired access time window together with maximizing the MRI scanners' utilization rates and minimizing the overtime. Additionally, the AS must provide patients with quick responses to their appointment requests. With a quick response it is meant that patients may experience maximum of one day's delay before receiving a response to their appointment request.

The radiologists of the Radiology department suggest dividing patients into four priority classes. The first class consists of same-day patients and should be examined today. The access time targets for the other patient classes are three days, one week and two weeks, respectively. This access time applies to the number of calendar days between the appointment request and the MRI scan and only future days are counted (the remaining part of today does not count as day 1).

All together, this already gives us the following requirements to the AS: the model must include multiple MRI scanners, multiple patient types and patient type-MRI scanner compatibility constraints. Patient types are defined in terms of access time target and capacity requirement. In addition, it must be possible to define different working hours for the various MRI scanners.
A final, hard requirement for the new AS is that, if it is not an online AS, there are at least two booking moments per day. We will explain why this is demanded. Part of the MRI requests comes from inpatients. The wish with these requests is that they will be seen the same day. If we define a booking moment at the end of the morning or early afternoon, we have the option to book same-day inpatient demand during the afternoon. If we also define a booking moment at the end of the afternoon, we can try to book the inpatient demand that is submitted during the afternoon in the remaining part of the working day. The same applies here for emergency requests coming from other outpatient clinics. The other outpatient clinics close at five in the afternoon, while the MRI scanners then operate for another hour. With a second decision moment at five o'clock in the afternoon we can, if necessary, try to book this urgent demand in the last hour of the day. In addition, if we, for example, make only booking decision in the morning, it is difficult to schedule an outpatient today. However, if we had the possibility to book this patient the previous afternoon, we might have the opportunity to book him/her today.

Decisions made to design ASs can be subdivided into three hierarchical levels, as it requires coordinated long-term, medium-term and short-term decisions. For the hierarchical levels, [20] applies the well known breakdown of *strategic, tactical and* *operational*. Strategic (or design) decisions are the long-term decisions that determine the main structure of an AS. Examples of strategic decisions include the number of servers/resources and the type of scheduling (offline, online or near-online). Tactical decisions are medium-term decisions related to how patients as a whole are scheduled, or how groups of patients are processed. Examples are allocation of capacity to different patient groups (probably with different priorities) or the length of appointment intervals. Operational decisions are short-term and are concerned with efficiently scheduling individual patients. Examples of operational decisions are the appointment day and the appointment time.

In the MRI AS to be designed, we do not have the flexibility to make strategic decisions. At the tactical level, our most influential decision is the allocation of available MRI-capacity over the different patient types and whether or not we do fix some appointment slots as emergency slots, or one-stop-shop appointment slots, or slots that could be booked decentrally. However, this is not a hard requirement for the new AS. If we want, we can undo the distinction between one-stop-shop patients, decentrally booked patients and *regular* patients. In that case, every patient submits a request to the booking clerks of the Radiology department and can be scheduled during every moment of the booking horizon.

The operational decisions in our AS to be designed include the assignment of appointment day, appointment time and the assignment of a patient to a specific MRI-scanner.

3

Literature study

This chapter provides the results from our literature study on operations research (OR) related literature that can support us in designing the appointment system (AS) for MRI scans in Rijnstate. In OR, appointment scheduling problems are an attractive research area, having been studied for more than half a century (since the seminal paper [2] by Bailey). The field of operations research provides numerous methodologies and solution techniques to simultaneously reduce costs and improve access to healthcare services.

In Chapter 2 we explained that the currently MRI AS in Rijnstate is based on a blueprint calendar. This blueprint calendar reserves time slots for specific patient classes, specific MRI examinations and the combination of both. An obvious design for an AS would again be a blueprint calendar, only a better functioning one. Such an AS design would include the tactical decision of defining a blueprint calendar combined with an assignment rule or algorithm at operational level. We reviewed some research articles dealing with this approach after which we conclude there is another field within OR that probably suits better for the problem at hand: *advanced patient scheduling*. The large body of literature associated with patient scheduling can broadly be divided into two streams: appointment (or allocation) scheduling and advance scheduling. Appointment scheduling refers to the assignment of specific appointment times and resources to patients but only once all patients for a given service day have been identified. Advance scheduling, on the other hand, refers to the allocation of future service capacity to demand as it arrives. The model we present in the next chapter fits within advance patient scheduling.
Before we elaborate on the related advance patient scheduling literature, we first turn to literature reviews on the subject of appointment scheduling. These literature reviews formed the starting point for our literature study.

3.1 Literature reviews on appointment scheduling

For literature reviews on appointment scheduling problems we refer to [1, 9, 18]. For a review of literature on ASs in which each patient needs multiple appointments, we refer to [24]. For an overview of the literature of the field of appointment scheduling that is not restricted to healthcare applications, we refer to [7, Chapter 2].

The focus of [1] is on post-2003 articles which provide optimization-based decision tools for AS decision makers. Literature is classified by the level of decision making: strategic, tactical and operational; and is evaluated from four perspectives: problem settings, environmental factors, modeling approaches and solution methods. The main goal of Cayirli and Veral in [9] is to review pre-2003 papers based on formulations and modeling considerations for ASs. In [18], Gupta and Denton focus on describing the most common types of healthcare appointment systems. For an literature review of articles from the field of OR that address the typical decisions to be made in resource capacity planning and control in healthcare, we refer to [21]. This review does not only include capacity allocation problems for outpatient clinics - or similar departments like primary care or diagnostic facilities such as the Radiology department of Rijnstate - but also discusses capacity planning and control in other services in healthcare such as surgical care services, emergency care services and home care services. From the perspective of resource capacity planning and control, different services may face similar questions.

To find relevant research articles that can support us in designing the AS for MRI scans in Rijnstate, we performed a backward and forward search on relevant articles cited in [1] and [7, Chapter 2]. Whether or not an article cited in either [1] or [7, Chapter 2] is relevant, is determined by the following criteria: the objective must be to minimize the access time or to maximize the number of patients seen within their pre-specified access time window; the model must include (or could (easily) be extended to) multiple patient types; the model must include (or could (easily) be extended to) multiple, heterogeneous resources and is able to deal with patient type-resource compatibility constraints.

3.2 Literature on capacity allocation

In [10], Creemers, Beliën and Lambrecht face multiple patient classes and propose a model for assigning server time slots to these classes that minimizes the total expected weighted waiting time of a patient (where different patient classes may be assigned different weights). They use a bulk service queueing model to obtain the expected waiting time of a patient of a particular class, given a feasible allocation of service time slots and use the output of this bulk service queueing model as the input of an optimization procedure.
This distribution of service time slots across various patient types results in a static blueprint for the complete booking horizon, something that offers little flexibility. To have more flexibility, in [37], Vermeulen et. al present an adaptive approach to optimize the allocation of CT scanners' capacity to different patient groups. It is adaptive in the sense that it takes into account the current and expected future situation. Upon his/her appointment request, a patient is assigned to a time slot within his/her access time window, randomly selected from all the free time slots that are suitable for the type of service. If there are no such free time slots, the approach shifts capacity between the different patients groups or it temporarily increases the number of appointment slots by extending opening hours. The decisions made in their approach are rule-based and simulation shows the impact of the decisions made.

Another article that deals with blueprint alteration on the day of service is [19]. Here, the dynamic uncertainty that arises from requests for appointments that arrive in real time and uncertainty due to last minute scheduling changes is addressed. The authors propose dynamic template scheduling for chemotherapy scheduling: a technique that combines proactive and online optimization using a blueprint calendar. First, a static blueprint agenda is created using a deterministic optimization model and a sample set of appointments. As requests for appointments arrive, this blueprint calendar is used to schedule them. When a request arrives that does not fit the template, the blueprint calendar is updated online using the proposed optimization model and a revised sample set of appointments. The goal is to minimize the possibly generated overtime that is needed to schedule this/these request(s) that does not fit in the blueprint agenda but must be served today.

If we want to keep time slots that are particularly blocked for patients with decentralized bookings or one-stop-shop patients, an important question is: when will we cancel this lock and will the slots be released for other types of patients. [16] addresses the service capacity reservation for a given class of customers. The reservation process is characterized by *contracted time slots* (CTS), reserved for the class of customers (in our setting the one-stop-shop or decentrally scheduled patients), and two advance cancellation modes to cancel CTSs either one period or two periods in advance. The optimal control under a given contract is formulated as an average cost Markov Decision Process (MDP) in order to minimize customer access times (of all classes), unused capacity and cancellation rate. Numerical results show that two-period advance CTS cancellation can significantly improve the contract-based solution.

The aforementioned articles cannot be directly labeled as advance patient scheduling and are not closely related to the approach we have chosen, as in our approach we say goodbye to a blueprint calendar and the specially blocked appointment slots for one-stop-shop patients. Nevertheless, we have come across the above articles in our search for relevant literature. Because our approach is not the only one and the Radiology department of Rijnstate might want to study the potency of a different approach than ours in the future, we have included the articles in our literature review.

3.3 Literature on advance patient scheduling

Advance scheduling problems typically assume that patients can be classified into multiple types according to their capacity requirements and urgency; resources have fixed regular capacity and that there exists the possibility of using overtime or an alternative source of surge capacity (see [27] for a more elaborate analysis of surge capacity and its usage). The aim is to identify effective ways of allocating available service capacity to incoming appointment requests while either maximizing the service level, i.e., the number of patients booked within the prespecified access time windows in a cost-effective manner or else maximizing revenue or throughput. Application areas include the scheduling of diagnostic tests such as MRIs [35] or CT scans [28] as well as radiation therapy treatments [34]. Papers in the area of advance scheduling mostly use dynamic programming, or approximate dynamic programming, due to the sequential nature of the scheduling decisions.

We begin with the advance scheduling model provided in [28]. In [28], Patrick, Puterman, and Queyranne present an infinite-horizon MDP formulation to dynamically allocate available daily CT scan capacity to incoming demand to achieve wait-time targets in a cost-effective manner. Because their state space is too large for a direct solution, approximate dynamic programming is used to find an approximate optimal policy for the MDP.

In [14], Erdelyi and Topaloglu present a model to dynamically allocate capacity to jobs of different priority by using stochastic approximation methods. In their paper they focus on a class of policies that are characterized by a set of protection levels. The role of these protection levels is to *protect* a portion of the daily capacity from the lower priority jobs so as to make it available for the future higher priority jobs.
In [15], Erdelyi and Topaloglu, present a general dynamic capacity allocation problem without an underlying healthcare related problem. There is a fixed amount of daily processing capacity. On each day, jobs of different priorities arrive randomly and a decision has to made about which jobs should be scheduled on which days. Waiting jobs incur a holding cost that is a function of their priority levels. The objective is to minimize the total expected cost over a finite planning horizon. The problem is formulated as a dynamic program, but this formulation is computationally difficult. Hence an approximate dynamic programming approach is used that decomposes the original formulation. Their results show that the found policy performs significantly better than a variety of benchmark strategies.

In both [14] and [28], the authors assume that each entity to be served consumes only one unit of capacity. In [34], Sauré et al. extent the model from [28] by adding multiple appointment requests that can have various capacity requirements. In their model, a patient can request multiple appointments for a single day and/or a series of appointments on consecutive days. The formulated infinite-horizon MDP for (dynamically) capacity allocation to various cancer treatments in radiation therapy units becomes intractable for reasonable instances and an approximate optimal policy is found via an equivalent linear programming model, which is solved through column generation.
In [6, Chapter 6] the model of Sauré et. al. in [34] is extended to include multiple servers and patient type-server compatibility constraints. Because the number patient types has tripled compared to [34], adding another approximation step in the solution approach is needed. The MDP is rewritten as a set of weakly coupled sub-MDPs and then Lagrangian relaxation is applied to the linking constraint. Afterwards, an affine value function approximation is used to solve an equivalent linear programming model through column generation to obtain an approximate optimal policy for the original MDP.

In [28], Patrick, Puterman, and Queyranne aim to schedule patients in a particular urgency class prior to a specific target date and the system is only penalized for lateness. In [17], Cocgun and Puterman study a scheduling problem in which arriving patients require appointments at specific future days within a treatment specific time window. In this paper, the system is penalized when appointments are either early or late. By varying the relative magnitude of penalty costs to diversion costs, this paper allows tolerance limits to be relaxed. This is relevant to manufacturing settings where time windows are more flexible. Cocgun and Puterman also model the problem as an infinite-horizon discounted MDP and find an approximate optimal policy via an equivalent linear programming model, which is solved through column generation.

In advance scheduling problems, the assumption of deterministic service times is largely made for convenience and in the general hope that, over time, average service times will work fairly well as an approximation. The value of this simplifying assumption is that the calculation of the performance metrics does not depend on the sequencing of patients but only on their service times. It allows overtime or idle time to be easily calculated as the number of appointments booked on a given day times the appointment length minus the regular capacity. In [33], the authors adapt the MDP from [28] to incorporate stochastic service times. They describe an enhanced version of the former MDP that can be used to incorporate patient classes differentiated by both priority and resource consumption as well as stochastic service times. They do not include multiple servers and patient type-server compatibility constraints. The calculation of the overtime or idle time is now based on a method described in [3]. The formulated MDP is again computationally intractable and an approximate optimal policy is found via an equivalent linear programming model, which is solved through column generation.

Finally, we would like to mention the research paper of Parizi and Ghate, [26]. They study an advance scheduling problem where appointment requests dynamically arrive over time and can either be rejected or booked for future slots. Furthermore, customers are heterogeneous in all problem parameters and may cancel an appointment or do not show up. The booking agent may overbook appointments to mitigate the detrimental effects of cancellations and no-shows. Parizi and Ghate provide a MDP formulation of this problem where the system receives a reward for providing service and incur costs for rejecting request, appointment delays and overtime. This MDP is intractable to an exact solution and has a weakly coupled structure (similar as the model in [6, Chapter 6]) that enables to apply the ADP method with Lagrangian relaxation.

The problem we face at the Radiology department of Rijnstate is closely related to that in [34] and [6, Chapter 6]. We have also the need to allocate finite capacity to different patient types with different access time targets; have multiple, heterogeneous servers; and patient type-server compatibility constraints. Our research expands both the MPDs from [34] and [6, Chapter 6] in the number of decision epochs each day incorporates. The direct cost structure in our MDP comes closest to the deterministic version of the MDP from [33].
Where the models in [34] and [6, Chapter 6] allow only to allocate available capacity once a day, this is possible twice a day in our model. This modification to the model allows more flexibility in responding to same-day (emergency) patients. Where emergency patients in [34] and [6, Chapter 6] were ignored (the model was not capable of preserving capacity for them), our model presented the following chapter does allow it to protect capacity for them. Although going from one decision epoch each day to two, we use the linear programming approach to ADP to find an approximate optimal policy. The obtained linear program is also solved via column generation.
The extension of our model from Chapter 4 that we present in Chapter 7 can be seen as the multiple server version of [33] with patient type-server compatibility restrictions.

4

Near-online multipriority patient scheduling with resource compatibility restrictions

In this chapter, we present a near-online, dynamic method for advance capacity planning involving multiple resources (MRI scanners) with different capabilities and taking into account future appointment requests. The goal is to identify effective ways of allocating available service capacity to incoming appointment requests while minimizing the number of patients whose access time exceeds the prespecified, priority-specific target in a cost-effective manner. Here greater weight is given to any late bookings of higher-priority demand. We formulate this planning method as a discounted infinite-horizon Markov Decision Process (MDP) by providing the decision epochs, state space, action sets, transition probabilities and costs in Section 4.1.

Due to the curse of dimensionality, the proposed MDP cannot be solved analytically even for small instances of the problem. For this reason, we develop in Section 4.2 an Approximate Dynamic Programming (ADP) approach to approximately solve it. Here we use an affine architecture to approximate the value function of the MDP and solve an equivalent linear programming model through column generation to obtain an approximate optimal policy.
In Chapter 6 we evaluate the performance of the resulting approximate optimal policy for a practical scale case study at the Radiology department of Rijnstate using simulation. The input parameters for this practical scale case study are determined in Chapter 5 using historical data.

4.1 The MDP formulation

In this section, we formulate a discounted infinite-horizon MDP model by providing the decision epochs, state space, action sets, transition probabilities and costs. Table 4.1 summarizes the notation used. Throughout this chapter we denote a vector by bolding it, such as $\mathbf{s}$.

Table 4.1: Notation of parameters used in the MDP model.

Set/Parameter(s)	Description
$\widetilde{N}$	length of the booking horizon in days
$\mathcal{N} = \{0, 1, \ldots, N\}$	set of sessions in the booking horizon (indexed by n).
$\mathcal{T} = \{1, \ldots, T\}$	set of patient types (indexed by t)
$\mathcal{T}' \subset \mathcal{T}$	set of all outpatient types
$AT(k,t)$	access time target for type t patients at decision epoch k ($k = 1, 2$), expressed in number of sessions (that is $AT(t) \in \mathcal{N}$)
$\mathcal{M} = \{1, \ldots, M\}$	set of MRI scanners (indexed by m)
$\mathcal{M}(t)$	set of MRI scanners suitable for type t (defined as a subset of $\mathcal{M}$, for all $t \in \mathcal{T}$)
$\mathcal{T}(m)$	set of patient types that can be served by MRI scanner m (defined as a subset of $\mathcal{T}$, for all $m \in \mathcal{M}$)
$k \in \{1, 2\}$	label/number corresponding to the first and second decision epoch on each day, respectively
$C_{mn}^{R,k}$	regular capacity of MRI-scanner m during session n expressed in number of time slots as observed from decision epoch k in a day ($k \in \{1, 2\}$)
$C_{mn}^{OT,k}$	overtime capacity of MRI-scanner m during session n expressed in number of time slots as observed from decision epoch k in a day ($k \in \{1, 2\}$)
d_t	duration of a type t MRI-scan expressed in number of appointment slots
x_{mn}	number of appointment slots booked/occupied on MRI-scanner m during session n of the booking horizon
y_t	number of appointment requests of type t waiting to be scheduled
Q_t^1, Q_t^2	maximum number of type t appointment requests that can be observed at the first and second decision epoch of next day, respectively

4.1.1 The decision epochs and the booking horizon

In the model we describe in this chapter we divide each day into three sessions: a morning session, an afternoon session and an evening session. In the general case, where we consider the hospital to have M MRI scanners (indexed by m), we assume that the start times of the morning, afternoon and evening session are identical for each MRI scanner and for each day, as the decision epochs in the MDP correspond to the begin of the afternoon and evening session. The exact duration of the sessions are hospital-specific and in particular the duration of the evening session can vary from MRI scanner to MRI scanner. However, we assume the duration of the evening session of MRI scanner m to be identical for each day. In Chapter 6 we describe the duration of the various sessions in Rijnstate. How the sessions' duration are formulated in the general model setting will become clear in the remainder of this section.

Throughout each day, requests for MRI examinations are submitted to the Radiology department. At the two decision epochs, which correspond to the begin of the afternoon and evening session (illustrated in Figure 4.1), the booking agent's task is to book these appointment requests into a future session. Usually, an outpatient clinic considers a booking horizon of $\widetilde{N}$ future days. As observed from the first decision epoch of the day, the current day has two remaining sessions, which means that an $\widetilde{N}$-day booking horizon corresponds to $3\widetilde{N} + 2$ sessions. As observed from the second decision epoch of the day, an $\widetilde{N}$-day booking horizon corresponds to $3\widetilde{N} + 1$ sessions since the current day has only one session left. We choose to define a booking horizon of $N := 3\widetilde{N} + 2$ sessions. This implies that at the start of an afternoon session on a specific day, the booking agent can book appointments into the afternoon and evening session of the current day and into the morning, afternoon and evening sessions of at most $\widetilde{N}$ days in advance. At the start of an evening session, appointments can be scheduled into the upcoming evening session, into the morning, afternoon and evening sessions of the next $\widetilde{N}$ days and into the morning session $\widetilde{N} + 1$ days from now. We denote the booking horizon by the set $\mathcal{N} := \{1, 2, \ldots, N(= 3\widetilde{N} + 2)\}$ (indexed by n).

Our model is complicated by the fact that the horizon is not static, but rolling. Thus, session n of the booking horizon at the current decision epoch becomes session $n - 1$ at the subsequent decision epoch. This implies that when we *jump* from the first decision epoch on a day (shown in Figure 4.1a) to the second one (shown in Figure 4.1b), the *newly added* last session of the booking horizon becomes a morning session and no appointments have been booked into this session. Similarly, if we *jump* from a second decision epoch on a day to the first decision moment on the successive day, i.e., from the situation shown in Figure 4.1b to the situation shown in Figure 4.1a, no appointments have been booked into the last two (new) sessions of the booking horizon. These sessions correspond to an afternoon and evening session, respectively.

4.1.2 Patient types and MRI scanner compatibility

In the MDP model we define patient types by the set $\mathcal{T} = \{1, \ldots, T\}$ (indexed by t) based on their clinical status (in- or outpatient), the duration of the MRI scan that they request, the compatibility of each MRI scanner with the requested MRI scan and their access time target. Since the M MRI scanners are represented by the set $\mathcal{M} = \{1, \ldots, M\}$, we can define the set $\mathcal{M}(t) \subseteq \mathcal{M}$ as the set of MRI scanners that can feasibly serve a patient of type t. The set $\mathcal{M}(t)$ is thus defined for all $t \in \mathcal{T}$. Conversely, $\mathcal{T}(m) \subseteq \mathcal{T}$, defined for all $m \in \mathcal{M}$, is the set of patient types t that can feasibly receive service on MRI scanner m. Furthermore, for a reason that will become clear in Section 4.1.4 on the possible actions, we define the set $\mathcal{T}' \subset \mathcal{T}$ to contain all outpatient types.

For all patient types $t \in \mathcal{T}$ we agree on an access time target (that is the medically acceptable (indirect) waiting time) of $AT(t)$ days. For the Radiology department it is thus the goal to examine a type t patient within $AT(t)$ days from the moment on the request was received.

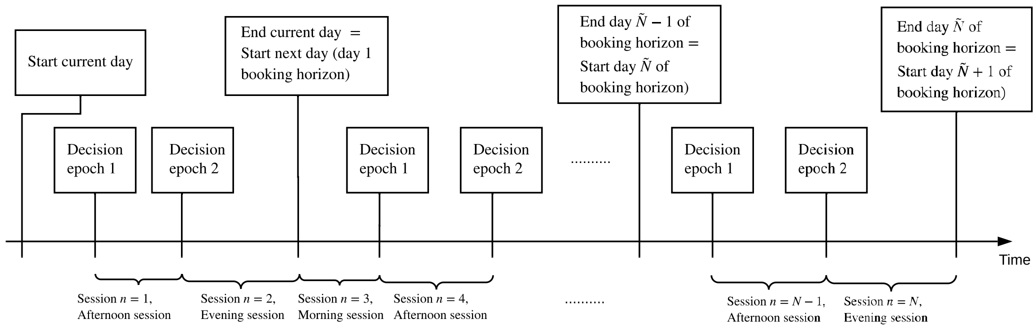

(a) Booking horizon and as observed from the first decision epoch (corresponding to the start of the afternoon session) on a given day.

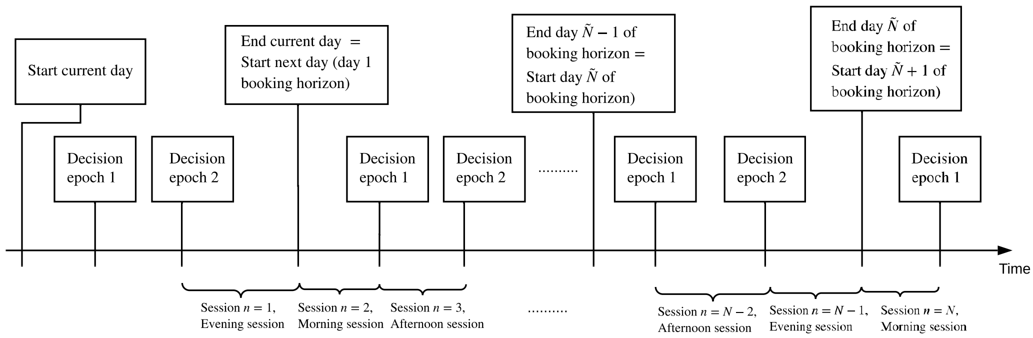

(b) Booking horizon as observed from the second decision epoch (corresponding to the start of the evening session) on a given day.

Figure 4.1: The booking horizon as observed from (a) the first decision epoch on a day and (b) the second decision epoch on a day. The dynamics of the MDP is going from (a) to (b) and then back to (a).

Since our booking horizon is defined in terms of sessions, we need to translate the access times from days into sessions. For example, take an appointment request with original priority to be served within three days. If this request is known at the start of the afternoon session, i.e., at the first decision epoch of a day, it translates to a priority to be scanned within eight sessions. On the other hand, if the request is submitted during the afternoon session and is observed at the second decision epoch of a day, it translates to a priority to be scanned within seven sessions. One way to deal with this, is to define the access time target $AT(t)$ for patient type t as observed at the day's second decision epoch. For the precedent example, this means that the request that is known at the first decision is also prioritized as to be scanned within seven sessions. Another way is to define both as a different patient types. That is, we define patient type t with access time target $AT(t) = 8$ sessions and type t' with $AT(t') = 7$ sessions. Subsequently, type t MRI appointment requests only arrive at the first decision epoch on a day and type t' only at the second. As follows from the state vector definition in Section 4.1.3, this is detrimental to the size of the state space and therefore we explain a third way to deal with different access times at the two decision epochs on a day that we will use in the sequel.
Note that besides different access times at the two decision epochs, type t patients do have the same characteristics as type t' patients.

Hence, for patient type t, we define an access time target $AT(1,t)$ for appointment requests that need to be booked at the first decision epoch on a day and an access time target $AT(2,t)$ for appointment requests that need to be booked at the second decision epoch on a day. Now we only need to define the patient type t and we can easily translate the access time expressed in days into the access time expressed in sessions. Note that, for all $t \in \mathcal{T}$, we have $AT(1,t) = AT(2,t) + 1$ and that we might have $AT(k,t) = AT(k,t')$, for $k = 1,2$, $t \neq t'$, $t,t' \in \mathcal{T}$.

4.1.3 The state space

At both decision epochs of a day, the number of time slots already booked on each MRI scanner during all sessions of the booking horizon is known, as well as the number of appointment requests from each patient type to be scheduled at the current decision epoch. Thus, a typical state of the system, denoted by $\mathbf{s}$, takes the form

$$\mathbf{s} = (k, \mathbf{x}, \mathbf{y}) = (k, x_{mn}, y_t)_{m \in \mathcal{M}, n \in \mathcal{N}; t \in \mathcal{T}},$$

where x_{mn} is the number of slots already booked on MRI scanner m during session n of the booking horizon and y_t is the number of type t patients waiting to be booked. In this state vector, k can take the value 1 or 2, corresponding to the decision epoch during the day.

At each decision epoch, the decision to be made for each appointment request is the assignment to a session and to an MRI scanner, or *diversion* to an alternative capacity source at an additional cost. This is often referred to as *surge* capacity (see [27]). The surge capacity in our model is overtime - alternatively it could be outsourcing. The overtime is usually at the end of a day and limited to a number of overtime units. However, we consider a system in which every session has limited overtime capacity and motivate this with the following example. Assume the following: on a fixed day, both the afternoon and evening session are fully booked and in the morning session only one appointment slot of ten minutes is empty. Futhermore, an appointment request with a duration of twenty minutes that must be booked into the fixed day is waiting to be booked. Then, if we do not allow overtime during the morning session, we know for sure that the MRI scanner will remain idle during the empty appointment slot in the morning, which results in an idle period of ten minutes. We will book the appointment request in the overtime of the fixed day, so after the regular program in the evening session. This results in twenty minutes of overtime on this day. On the other hand, if we allow overtime in the morning session, say an overtime of one 10-minute slot, we can book the appointment partly in the regular time of the morning session and partly in the overtime of the morning session. The result: it only adds ten minutes to the overtime of the day and we have no idle time. The disadvantage of this decision could be an increase in patients' waiting time.

In state $\mathbf{s}_1 = (1, \mathbf{x}, \mathbf{y})$, that is at the first decision epoch of a given day, sesion $n = 1$ corresponds to an afternoon session (see Figure 4.1a). During this afternoon session, MRI scanner m has regular capacity of $C_{m1}^{R,1}$ time slots and overtime capacity of $C_{m1}^{OT,1}$ time slots. At the next decision epoch of the given day, we observe some state $\mathbf{s}_2 = (2, \mathbf{x}', \mathbf{y}')$ and now session $n = 1$ corresponds to an evening session (see Figure 4.1b).

During this evening session, MRI scanner m has regular capacity of $C_{m1}^{R,2}$ time slots and overtime capacity of $C_{m1}^{OT,2}$ time slots. Since evening sessions usually have different capacity than afternoon sessions, we have that $C_{m1}^{R,2} \neq C_{m1}^{R,1}$ and $C_{m1}^{OT,2} \neq C_{m1}^{OT,1}$. However, since the evening session of the given day corresponds to session $n = 2$ as observed from the first decision epoch of the given day, we have $C_{m1}^{R,2} = C_{m2}^{R,1}$ and $C_{m1}^{OT,2} = C_{m2}^{OT,1}$ for all $m \in \mathcal{M}$. In general we have

$$
\begin{aligned}
C_{mn}^{R,2} &= C_{m,n+1}^{R,1}, \quad \text{for all } m \in \mathcal{M} \text{ and for all } n = 1, 2, \ldots, N-1, \\
C_{mn}^{OT,2} &= C_{m,n+1}^{OT,1}, \quad \text{for all } m \in \mathcal{M} \text{ and for all } n = 1, 2, \ldots, N-1, \\
C_{mN}^{R,2} &= C_{m3}^{R,1}, \qquad \text{for all } m \in \mathcal{M}, \\
C_{mN}^{OT,2} &= C_{m3}^{OT,1}, \qquad \text{for all } m \in \mathcal{M}.
\end{aligned}
$$

The last and second to last equation say, respectively, that the regular and overtime capacity of the last session of the booking horizon, i.e., session N, as observed from the second decision epoch on a day is equal to the regular and overtime capacity of session 3 as observed from the first decision epoch on a day. This is because session N as observed from the second decision epoch on a day is always a morning session and the first morning session as observed from the first decision epoch on a day corresponds to session 3 of the booking horizon.

For $k \in \{1, 2\}$, we define the sets

$$
\mathcal{S}_k = \left\{ (k, \mathbf{x}, \mathbf{y}) \;\middle|\; \begin{array}{ll} x_{mn} \leq C_{mn}^{R,k} + C_{mn}^{OT,k}, & \text{for all } m \in \mathcal{M} \text{ and } n \in \mathcal{N}; \\ y_t \leq Q_t^k, & \text{for all } t \in \mathcal{T}; \\ (\mathbf{x}, \mathbf{y}) \in \mathbb{N}_0^{MN} \times \mathbb{N}_0^T & \end{array} \right\}. \tag{4.1}
$$

Here we adopted the notation $\mathbb{N}_0 = \{0, 1, 2, 3, \ldots\}$ and let $\mathbb{N}_0^n$ be its n-dimensional extension. Q_t^k is the maximum number of MRI appointment requests of type t present in the system (and thus waiting to be booked) at decision epoch k of a day. That is, Q_t^1 is the maximum number of type t appointment requests arrived between the two decision epochs in the same day and Q_t^2 is the maximum number of type t appointment requests arrived between the second decision epoch in a fixed day and its immediate successor, which is the first decision epoch in the following 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. Note here that we implicitly assume that the number of appointment requests observed at a decision epoch only concerns newly arrived requests. It is therefore not possible to postpone booking decisions and all requests that are observed at the current decision epoch will be scheduled at the current decision epoch.

Because we experience a rolling booking horizon, in state $\mathbf{s}_1 = (1, \mathbf{x}, \mathbf{y})$ there are no appointments scheduled during sessions $N-1$ and N, that is, $x_{m,N-1} = x_{mN} = 0$, for all $m \in \mathcal{M}$. Similarly, in state $\mathbf{s}_2 = (2, \mathbf{x}, \mathbf{y})$ there are no bookings in session N, that is $x_{mN} = 0$, for all $m \in \mathcal{M}$.

The state space of the MDP is given by

$$\mathcal{S} = \mathcal{S}_1 \cup \mathcal{S}_2. \tag{4.2}$$

4.1.4 The action sets

At each decision epoch, the booking agent's task is to decide on which MRI scanner and during which session to schedule each of the patients waiting to be booked. Thus, a vector of possible actions can be written as $\mathbf{a} = (a_{tmn})_{t\in\mathcal{T},m\in\mathcal{M},n\in\mathcal{N}}$, where a_{tmn} is the number of type t patients to book on MRI scanner m during session n. Note that, once a patient is assigned to a session, a second level of scheduling is needed which assigns patients to specific appointment times. That goes beyond the scope of this MDP model.

Input parameters of the decision are y_t, the number of type t patients to be booked, and x_{mn}, the number of time slots already occupied on MRI scanner m during session n of the booking horizon. We experience a rolling horizon and have two decision epochs in each day and hence we face different constraints for action $\mathbf{a}$ to be valid in states $\mathbf{s}_1 = (1, \mathbf{x}, \mathbf{y})$ and $\mathbf{s}_2 = (2, \mathbf{x}, \mathbf{y})$. To be valid, any action in $\mathbf{s}_1$ must satisfy the following constraints:

$$\sum_{n\in\mathcal{N}} \sum_{m\in\mathcal{M}(t)} a_{tmn} = y_t, \qquad \text{for all } t \in \mathcal{T}, \tag{4.3}$$

$$x_{mn} + \sum_{t\in\mathcal{T}(m)} d_t a_{tmn} \le C^{R,1}_{mn} + C^{OT,1}_{mn}, \qquad \text{for all } m \in \mathcal{M} \text{ and } n \in \mathcal{N}, \tag{4.4}$$

$$a_{tmn} = 0 \qquad \text{for all } t \in \mathcal{T}, m \notin \mathcal{M}(t) \text{ and } n \in \mathcal{N}, \tag{4.5}$$

$$a_{tm1} = a_{tm2} = 0 \qquad \text{for all } t \in \mathcal{T}' \text{ and for all } m \in \mathcal{M}(t), \tag{4.6}$$

$$a_{tmn} \in \mathbb{N}_0^{TMN} \tag{4.7}$$

Constraint (4.3) requires the number of bookings for each patient type to be equal to the number of requests waiting to be booked. As we discussed earlier, all the demand waiting to be booked is newly arrived demand and we do not have the possiblity the postpone the booking decision. In Constraint (4.4), d_t describes the duration of a type t scan in number of time slots. Constraint (4.4) therefore restricts the total number of slots booked on MRI scanner m during session n of the booking horizon to be less than or equal to the number of available regular slots plus the overtime slots. Note that in this constraint we only need to sum over $\mathcal{T}(m)$, i.e., we only sum over those patient types who can be served by MRI scanner m. Constraint (4.5) ensures patients are only booked on an MRI scanner suitable for their type of MRI scan. Recall that in Section 4.1.2 we defined the set $\mathcal{T}'$ to contain all outpatient patient types. We do not allow outpatient demand that does not have the priority to be scanned today to be booked into one of today's remaining sessions, because these patients are not necessarily already present in the hospital and because we cannot control the appointment time with this MDP model, we cannot be sure that the patients will be in the hospital on time. Hence we have Constraint (7.6). Finally, all action variables are integer and non-negative (captured by Constraint (4.7)).

Similarly, to be valid, any action in $\mathbf{s}_2$ must satisfy the following constraints:

$$\sum_{n \in \mathcal{N}} \sum_{m \in \mathcal{M}(t)} a_{tmn} = y_t, \qquad \text{for all } t \in \mathcal{T}, \tag{4.8}$$

$$x_{mn} + \sum_{t \in \mathcal{T}(m)} d_t a_{tmn} \leq C_{mn}^{R,2} + C_{mn}^{OT,2}, \qquad \text{for all } m \in \mathcal{M} \text{ and } n \in \mathcal{N}, \tag{4.9}$$

$$a_{tmn} = 0 \qquad \text{for all } t \in \mathcal{T}, m \notin \mathcal{M}(t) \text{ and } n \in \mathcal{N}, \tag{4.10}$$

$$a_{tm1} = 0 \qquad \text{for all } t \in \mathcal{T}' \text{ and for all } m \in \mathcal{M}(t), \tag{4.11}$$

$$a_{tmn} \in \mathbb{N}_0^{TMN} \tag{4.12}$$

We define the action sets $\mathcal{A}(\mathbf{s}_1)$, for any given state $\mathbf{s}_1 = (1, \mathbf{x}, \mathbf{y}) \in \mathcal{S}_1$, as the set of actions $\mathbf{a}$ satisfying Equations (4.3) to (4.7). We define the action sets $\mathcal{A}(\mathbf{s}_2)$, for any given state $\mathbf{s}_2 = (2, \mathbf{x}, \mathbf{y}) \in \mathcal{S}_2$, as the set of actions $\mathbf{a}$ satisfying Equations (4.8) to (4.12). In general we let $\mathcal{A}(s)$ denote the set of all feasible actions in state $\mathbf{s} \in \mathcal{S}$.

4.1.5 The transition probabilities

First consider the situation in which the system is at the first decision epoch of a given day. That is, the system is in some state $\mathbf{s}_1 = (1, \mathbf{x}, \mathbf{y}) \in \mathcal{S}_1$. Once a decision is made in this state, the system moves to a state corresponding to the second decision epoch of the given day, i.e., the system moves to some state $\mathbf{s}_2 = (2, \mathbf{x}', \mathbf{y}') \in \mathcal{S}_2$. The only source of uncertainty in this transition from $\mathbf{s}_1$ to $\mathbf{s}_2$ is the number of new appointment requests of each patient type. The other parameters are updated based on the capacity allocation to booked patients. Thus, as a result of choosing booking action $\mathbf{a}$ in state $\mathbf{s}_1 = (1, \mathbf{x}, \mathbf{y})$, and having y'_t new requests of type t, the state of the system the next decision epoch, denoted by $\mathbf{s}_2 = (2, \mathbf{x}', \mathbf{y}')$, will be determined by the following probability distribution:

$$p(\mathbf{s}_2 | \mathbf{s}_1, \mathbf{a}) = \begin{cases} p_2(\mathbf{y}') = \prod_{t \in \mathcal{T}} p_2\left(y'_t\right), & \text{if } \mathbf{s}_2 = (2, \mathbf{x}', \mathbf{y}') \text{ satisfies Eq. (4.14),} \\ 0 & \text{otherwise.} \end{cases} \tag{4.13}$$

$$x'_{mn} = \begin{cases} x_{m,n+1} + \sum_{t \in \mathcal{T}(m)} d_t a_{tm,n+1}, & \text{for all } m \in \mathcal{M} \text{ and } n = 1, 2, \ldots, N-1, \\ 0, & \text{for all } m \in \mathcal{M} \text{ and } n = N. \end{cases} \tag{4.14}$$

We assume that demand for each patient type is independent and that the each day's demand at decision epoch 1 is independent as well. Because demand arises from multiple independent sources (the nursing wards, other outpatient clinics within the hospital and specialists in the region serviced by the department), independence between patient types seems a reasonable assumption. Note that we denote the probability distribution of new appointment requests between the first and second decision epoch on the same day by p_2, because this implies that we can write the expected number of appointment requests observed at the second decision epoch of the day as $\mathbb{E}_{p_2}[\mathbf{y}]$, which seems more logical than $\mathbb{E}_{p_1}[\mathbf{y}]$.

Similarly, if the system is at the second decision epoch on a given day, the system's next state is at the first decision epoch in the successive day. Thus, as a result of choosing booking action $\mathbf{a}$ in state $\mathbf{s}_2 = (2, \mathbf{x}, \mathbf{y})$, and having y'_t new requests of type t, the state of the system the next decision epoch, denoted by $\mathbf{s}_1 = (1, \mathbf{x}', \mathbf{y}')$, will be determined by the following probability distribution:

$$p(\mathbf{s}_1|\mathbf{s}_2, \mathbf{a}) = \begin{cases} p_1(\mathbf{y}') = \prod_{t\in\mathcal{T}} p_1\left(y'_t\right), & \text{if } \mathbf{s}_1 = (1, \mathbf{x}', \mathbf{y}') \text{ satisfies Eq. (4.16)}, \\ 0 & \text{otherwise.} \end{cases} \tag{4.15}$$

$$x'_{mn} = \begin{cases} x_{m,n+2} + \sum_{t\in\mathcal{T}(m)} d_t a_{tm,n+2}, & \text{for all } m \in \mathcal{M} \text{ and } n = 1, 2, \ldots, N-2, \\ 0, & \text{for all } m \in \mathcal{M} \text{ and } n = N-1, N. \end{cases} \tag{4.16}$$

In general we expect $p_1 \neq p_2$.

4.1.6 The direct costs

The direct costs associated with state-action pair $\mathbf{s} \in \mathcal{S}$, $\mathbf{a} \in \mathcal{A}(\mathbf{s})$ derives from two sources: (i) a cost associated with booking patients beyond their priority-specific access time targets; and (ii) a cost associated with the use of overtime.

$$c(\mathbf{s}, \mathbf{a}) = c(k, \mathbf{x}, \mathbf{y}, \mathbf{a}) = \sum_{n\in\mathcal{N}} \sum_{t\in\mathcal{T}} f^{AT}(k, t, n) \left[\sum_{m\in\mathcal{M}(t)} a_{tmn}\right] + f^{AS}(\mathbf{s}, \mathbf{a}). \tag{4.17}$$

In Equation (4.17), $f^{AT}(k, t, n)$ is the cost associated with booking a type t patient into session n of the booking horizon, at decision epoch k of a day. The cost associated with overtime is $f^{AS}(\mathbf{s}, \mathbf{a})$ (referring to the appointment schedule).

Recall that $AT(k, t)$ represents the access time target for a type t patient at decision epoch k of day. A suitable manner to determine the $f^{AT}(k, t, n)$, is to make use of the penalties g_{tn}, which are a penalty for each additional session of wait n before a type t patient can be served. It is clearly reasonable to assume that $f^{AT}(t, n)$ should be zero if $n \leq AT(t)$. Thus, a suitable form for the booking cost $f^{AT}(t, n)$ is

$$f^{AT}(k, t, n) = \begin{cases} \sum_{j=1}^{n-AT(k,t)} \gamma^{j-1} g_{tn} & \text{for all } n > AT(t) \\ 0 & \text{otherwise.} \end{cases}$$

Here γ is the discount factor per session.

The overtime costs at the first decision epoch of the day are defined as

$$f^{AS}(\mathbf{s}_1, \mathbf{a}) = \sum_{m \in \mathcal{M}} \sum_{n \in \mathcal{N}} h \left[x_{mn} + \sum_{t \in \mathcal{T}(m)} d_t a_{tmn} - C^{R,1}_{mn} \right]^+ - \sum_{m \in \mathcal{M}} \sum_{n \in \mathcal{N}} h \left[x_{mn} - C^{R,1}_{mn} \right]^+,$$

where h is the overtime cost per time slot and $[z]^+ = \max\{0, z\}$. Note that the overtime cost h is to be assumed identical for all MRI scanners $m \in \mathcal{M}$.
Note that the system only incurs overtime costs for overtime slots that are booked at the current decision epoch. For the overtime slots that were already booked, costs will not be charged again, since this has already happened at an earlier decision epoch. Finally, note that this structure is not the easiest way to penalize the system for overtime, but we believe it is a manner to keep the size of the state and action vector small.

Similarly, the overtime costs at the second decision epoch of the day are defined as

$$f^{AS}(\mathbf{s}_2, \mathbf{a}) = \sum_{m \in \mathcal{M}} \sum_{n \in \mathcal{N}} h \left[x_{mn} + \sum_{t \in \mathcal{T}(m)} d_t a_{tmn} - C^{R,2}_{mn} \right]^+ - \sum_{m \in \mathcal{M}} \sum_{n \in \mathcal{N}} h \left[x_{mn} - C^{R,2}_{mn} \right]^+.$$

4.1.7 The Bellman equations

The value function $v^{\pi}_{\gamma}(\mathbf{s})$ of the MDP specifies the total expected discounted cost (discount factor γ) over the infinite horizon for state $\mathbf{s} \in \mathcal{S}$ under policy π. Of course, we are not so much interested in determining the value function for a given policy π as in finding an optimal policy π^*. To identify such a policy we need to solve the following optimality equations for $v(\mathbf{s})$ - also known as the Bellman equations:

$$v(\mathbf{s}) = v(k, \mathbf{x}, \mathbf{y}) = \min_{\mathbf{a} \in \mathcal{A}(\mathbf{s})} \left\{ c(\mathbf{s}, \mathbf{a}) + \gamma \sum_{\mathbf{s}' \in \mathcal{S}} p(\mathbf{s}'|\mathbf{s}, \mathbf{a}) v(\mathbf{s}') \right\}, \quad \text{for all } \mathbf{s} \in \mathcal{S}, \tag{4.18}$$

where $0 \leq \gamma < 1$ is the discount factor per session. Note that, depending on the decision epoch of the day the system is *jumping* to, k', we always must have that

$$p(\mathbf{s}'|\mathbf{s}, \mathbf{a}) = p((k', \mathbf{x}', \mathbf{y}')|(k, \mathbf{x}, \mathbf{y}), \mathbf{a}) = p_{k'}(\mathbf{y}'),$$

where $\mathbf{y}' \in \mathcal{Y}_k$ and

$$\mathcal{Y}_k = \left\{ \mathbf{y} = (y_1, y_2, \ldots, y_T) \;\middle|\; \mathbf{y} \in \mathbb{N}_0^T,\, y_t \leq Q_t^k, \text{ for all } t = 1, 2, \ldots, T \right\}.$$

From its definition given by (4.2) we will determine the size of the state space, $|\mathcal{S}|$, in Section 4.1.8. However, we can already conclude that, although $|\mathcal{S}|$ may be very large, the state space is finite. Also, from the action set definitions given by the Eqs. (4.3)-(4.7) and Eqs (4.8)-(4.12) it follows that the action sets $\mathcal{A}(\mathbf{s})$ are finite for each $\mathbf{s} \in \mathcal{S}$. The direct costs $c(\mathbf{s}, \mathbf{a})$, defined in (4.17), are bounded (for some Z we have that $|c(\mathbf{s}, \mathbf{a})| \leq Z < \infty$, for all $\mathbf{s} \in \mathcal{S}$ and $\mathbf{a} \in \mathcal{A}(\mathbf{s})$) and like the transition probabilities $p(\mathbf{s}'|\mathbf{s}, \mathbf{a})$ stationary, i.e., they do not vary from decision epoch to decision epoch. Hence, it follows from the following theorem that there exists a unique solution to the optimality equations (4.18).

Theorem 4.1 (*Theorem 6.2.5 of [30]*)**.** Suppose $0 \leq \gamma < 1$, $\mathcal{S}$ is finite and $c(\mathbf{s}, \mathbf{a})$ is bounded for all $\mathbf{s} \in \mathcal{S}$, $\mathbf{a} \in \mathcal{A}(\mathbf{s})$. Then there exists an unique solution $\mathbf{v}^* = \{v^*(\mathbf{s})\}_{\mathbf{s}\in\mathcal{S}}$ to the optimality equations (4.18).

Proof. See the proof of Theorem 6.2.5 on p.151 of [30]. ∎

Existence of a deterministic stationary policy which is optimal

Under the assumptions prior to Theorem 4.1 it follows from [30, Theorem 6.2.10, p. 151] that there exists an optimal deterministic stationary policy $\pi^* = (d, d, \dots) = d^\infty$ and that the corresponding value function $v_\gamma^{d^\infty}(\mathbf{s})$ solves the optimality equations (4.18) and thus equals $v^*(\mathbf{s})$ from Theorem 4.1. Hence the optimal decision rule d is the function $d : \mathcal{S} \to \mathcal{A}(\mathbf{s})$, given by

$$d(\mathbf{s}) = a^*_{\mathbf{s}} \in \underset{\mathbf{a}\in\mathcal{A}(\mathbf{s})}{\arg\min} \left\{ c(\mathbf{s}, \mathbf{a}) + \gamma \sum_{\mathbf{s}'\in\mathcal{S}} p(\mathbf{s}'|\mathbf{s}, \mathbf{a}) v^*(\mathbf{s}') \right\}.$$

4.1.8 The dimensions of the state space and actions sets

We consider two types of state variables: those that are feasible at the first decision epochs on a day: $\mathbf{s}_1 = (1, \mathbf{x}, \mathbf{y}) \in \mathcal{S}_1$ and those feasible at the second decision epoch $\mathbf{s}_2 = (2, \mathbf{x}, \mathbf{y}) \in \mathcal{S}_2$. Recall that $\mathcal{S}_1$ and $\mathcal{S}_2$ are defined in (7.1). Both types of state variables have $(1 + M \times N + T)$ dimensions. For $\mathbf{s}_1 = (1, \mathbf{x}, \mathbf{y}) \in \mathcal{S}_1$, x_{mn} can take values $0, 1, 2, \dots, C^{R,1}_{mn} + C^{OT,1}_{mn}$, for all $m \in \mathcal{M}$ and $n = 1, 2, \dots, N-2$ and y_t can take values $0, 1, 2, \dots, Q^1_t$. $x_{m,N-1}$ and x_{mN} are both restricted to be zero for all $m \in \mathcal{M}$. Thus, the number of feasible states at the first decision epoch of a day, denoted by $|\mathcal{S}_1|$, equals

$$|\mathcal{S}_1| = \prod_{m=1}^{M} \prod_{n=1}^{N-2} \left(1 + C^{R,1}_{mn} + C^{OT,1}_{mn}\right) \times \prod_{t=1}^{T} \left(1 + Q^1_t\right).$$

Similarly, for $\mathbf{s}_2 = (2, \mathbf{x}, \mathbf{y}) \in \mathcal{S}_2$, x_{mn} can take values $0, 1, 2, \dots, C^{R,2}_{mn} + C^{OT,2}_{mn}$, for all $m \in \mathcal{M}$ and $n = 1, 2, \dots, N-1$ and y_t can take values $0, 1, 2, \dots, Q^2_t$. x_{mN} is restricted to be zero for all $m \in \mathcal{M}$. Thus, analogous to $|\mathcal{S}_1|$, we get

$$|\mathcal{S}_2| = \prod_{m=1}^{M} \prod_{n=1}^{N-1} \left(1 + C^{R,2}_{mn} + C^{OT,2}_{mn}\right) \times \prod_{t=1}^{T} \left(1 + Q^2_t\right).$$

From its definition, provided by Eq. (4.2), it follows that

$$\begin{aligned} |\mathcal{S}| = &\prod_{m=1}^{M} \prod_{n=1}^{N-2} \left(1 + C^{R,1}_{mn} + C^{OT,1}_{mn}\right) \times \prod_{t=1}^{T} \left(1 + Q^1_t\right) \\ &+ \prod_{m=1}^{M} \prod_{n=1}^{N-1} \left(1 + C^{R,2}_{mn} + C^{OT,2}_{mn}\right) \times \prod_{t=1}^{T} \left(1 + Q^2_t\right). \end{aligned}$$

The action variable $\mathbf{a}$ has $(T \times M \times N)$ dimensions. If $|\mathcal{M}(t)|$ denotes the cardinality of the set $\mathcal{M}(t)$ of MRI scanners which are suitable for type t patients, we might have up to

$$|\mathcal{A}(\mathbf{s}_1)| = \prod_{t=1}^{T} \left(1 + Q_t^1\right)^{N \times |\mathcal{M}(t)|},$$

different (not necessarily feasible) actions for states $\mathbf{s}_1 = (1, \mathbf{x}, \mathbf{y}) \in \mathcal{S}_1$. Similarly, for states $\mathbf{s}_2 = (2, \mathbf{x}, \mathbf{y}) \in \mathcal{S}_2$, we might have up to

$$|\mathcal{A}(\mathbf{s}_2)| = \prod_{t=1}^{T} \left(1 + Q_t^2\right)^{N \times |\mathcal{M}(t)|},$$

different (not necessarily feasible) actions.

The challenge is that even for very small instances the size of the state space and the size of the corresponding action sets become intractable, which makes a direct solution to (4.18) impossible. This is the widely known *curse of dimensionality* of dynamic programming and is the most often-cited reason why algorithms such as value iteration will not find an optimal policy in reasonable time [29].

4.2 The solution approach

In order to deal with an intractable number of states and actions, we first transform our MDP model into its equivalent linear programming form. The linear programming approach to discounted infinite-horizon MDPs, initially presented in [11], is based on writing the optimality equations in (4.18) as follows:

$$\max_{\mathbf{v}} \quad \sum_{\mathbf{s} \in \mathcal{S}} \eta(\mathbf{s}) v(\mathbf{s}), \tag{4.19a}$$

$$\text{s.t.} \quad c(\mathbf{s}, \mathbf{a}) + \gamma \sum_{\mathbf{s}' \in \mathcal{S}} p(\mathbf{s}'|\mathbf{s}, \mathbf{a}) v(\mathbf{s}') \geq v(\mathbf{s}), \quad \text{for all } \mathbf{s} \in \mathcal{S}, \mathbf{a} \in \mathcal{A}(\mathbf{s}), \tag{4.19b}$$

$$\mathbf{v} \in \mathbb{R}^{|\mathcal{S}|}. \tag{4.19c}$$

The value of $\eta(\mathbf{s})$ represents the weight of state $\mathbf{s} \in \mathcal{S}$ in the objective function. The solution to this equivalent linear programming model (4.19a)-(4.19c) is the same as the solution to the optimality equations (4.18) when $\eta(\mathbf{s})$ is strictly positive for all $\mathbf{s} \in \mathcal{S}$ [13, 23]. For the sake of completeness, we indicate that $\mathbb{R}^n$ is the n-dimensional real space and consists of all n-tuples of real numbers, i.e., an element of $\mathbb{R}^n$ is a vector $\mathbf{z} = (z_1, z_2, \ldots, z_n)$, where each $z_i \in \mathbb{R}$ is a real number.

The equivalent linear program (LP), however, does not avoid the curse of dimensionality. The model in (4.19a)-(4.19c) has one variable for every state $\mathbf{s} \in \mathcal{S}$ and one constraint for every feasible state-action pair $(\mathbf{s}, \mathbf{a}), \mathbf{s} \in \mathcal{S}, \mathbf{a} \in \mathcal{A}(\mathbf{s})$, making its solution impossible. Fortunately, a whole field of potential methods for dealing with the curse of dimensionality, called Approximate Dynamic Programming (ADP), has been developed in the last decade [4, 29]. The approach within this field that we consider consists of using an approximation architecture to represent the value function in the MDP formulation or, equivalently, the variables in the equivalent linear programming model. This *approximate linear programming approach to ADP* was initially introduced in [36].

In the approximate linear programming (ALP) approach, we approximate the value function
$v(\mathbf{s}), \mathbf{s} \in \mathcal{S}$ in the LP formulation by a linear superposition of fitting/basis functions $\{\phi_f : f \in \mathcal{F}\}$:

$$v(\mathbf{s}) \approx \sum_{f \in \mathcal{F}} \theta_f \phi_f(\mathbf{s}), \quad \mathbf{s} \in \mathcal{S}, \tag{4.20}$$

where the fitting functions $\{\phi_f : f \in \mathcal{F}\}$ are given and the coefficients/weights
$\{\theta_f \in \mathbb{R} : f \in \mathcal{F}\}$ must be chosen to give a good fit [29, p. 411]. The incentive for this approximation approach is the reduction in problem dimensionality from $|\mathcal{S}|$ to $|\mathcal{F}|$, with significant savings in computation time if $|\mathcal{F}|$ is much smaller than the number of states. A successful fit occurs if one can find a small set of fitting functions such that one obtains both a good approximation to the value function and a good suboptimal policy from this approximation to the value function [36]. The goal to find a good suboptimal policy from the approximation to the value function goes hand in hand with the quality of the approximation to the value function. After this approximation is found, it is inserted into the optimality equations (4.18). To find the approximate optimal action, these optimality equations, with the approximated value function at the place of the real value function, are solved for $\mathbf{a} \in \mathcal{A}(\mathbf{s})$. Hence, a bad quality approximation to the value function leads to a bad quality approximate optimal action. The fitting functions can be chosen in essentially any appropriate manner: polynomials, splines [12], etc.

To solve the MRI appointment scheduling problem we chose an affine approximation to $v(\mathbf{s}), \mathbf{s} \in \mathcal{S}$ in Section 4.2.1. This affine approximation results in an ALP, whose dual we solve using column generation. The column generation algorithm is explained in Section 4.2.2. In Section 4.2.3 we use the solution of the ALP found by the column generation algorithm to identify an approximate optimal policy. The affine approximation we introduce in Section 4.2.1 does not take the structure as in (4.20), because this structure is less convenient than the one introduces in Section 4.2.1. However, the structure in (4.20) allows us to prove that the LP and its dual that we eventually get in Section 4.2.1 both have finite optima. After specifying the affine approximation adopting the structure of (4.20) and the resulting LP, we formally state this in Theorem 4.3 and conclude this section with its proof. This theorem and its proof, together with the formulation of the ALP in Section 4.2.1 are technical and include many mathematical manipulations. If, as a reader, you want to skip this technical part for whatever reason, that is possible. Then read on from the ALP (4.27a)-(4.27d) on page 35. This is the final LP whose solution will lead to an approximate optimal policy.

Define the index set as

$$\mathcal{F} = \{1, 2, (m,n), (\tilde{m}, \tilde{n}), t, \tilde{t} \mid m, \tilde{m} = 1, \ldots, M, \;\; n, \tilde{n} = 1, \ldots, N, \;\; t, \tilde{t} = 1, \ldots, T\}.$$

That is, we use $2(1 + M \times N + T)$ fitting functions. We define the fitting functions as

$$\phi_1(k, \mathbf{x}, \mathbf{y}) = \begin{cases} 1 & \text{if } k = 1, \\ 0 & \text{if } k = 2, \end{cases}, \qquad \phi_2(k, \mathbf{x}, \mathbf{y}) = \begin{cases} 0 & \text{if } k = 1, \\ 1 & \text{if } k = 2, \end{cases}$$

$$\phi_{(m,n)}(k,\mathbf{x},\mathbf{y}) = \begin{cases} x_{mn} & \text{if } k=1, \\ 0 & \text{if } k=2, \end{cases}, \quad \phi_{(\tilde{m},\tilde{n})}(k,\mathbf{x},\mathbf{y}) = \begin{cases} 0 & \text{if } k=1, \\ x_{mn} & \text{if } k=2, \end{cases} \tag{4.21}$$

$$\phi_t(k,\mathbf{x},\mathbf{y}) = \begin{cases} y_t & \text{if } k=1, \\ 0 & \text{if } k=2, \end{cases}, \qquad \phi_{\tilde{t}}(k,\mathbf{x},\mathbf{y}) = \begin{cases} 0 & \text{if } k=1, \\ y_t & \text{if } k=2, \end{cases}$$

Finally, given the coefficients $\{\theta_f \in \mathbb{R} : f \in \mathcal{F}\}$, we introduce the functions $\omega(\mathbf{s},\boldsymbol{\theta}) = \sum_{f\in\mathcal{F}} \theta_f \phi_f(\mathbf{s})$, for all $\mathbf{s} \in \mathcal{S}$, because this turns out to be convenient later.

If we now substitute the affine approximation into (4.19a)-(4.19c) and rearrange terms, we get

$$\max_{\boldsymbol{\theta}} \quad \sum_{\mathbf{s}\in\mathcal{S}} \eta(\mathbf{s}) \left[\sum_{f\in\mathcal{F}} \theta_f \phi_f(\mathbf{s}) \right], \tag{4.22a}$$

$$\text{s.t.} \quad \sum_{f\in\mathcal{F}} \theta_f \left[\phi_f(\mathbf{s}) - \gamma \sum_{\mathbf{s}'\in\mathcal{S}} p(\mathbf{s}'|\mathbf{s},\mathbf{a}) \phi_f(\mathbf{s}') \right] \le c(\mathbf{s},\mathbf{a}), \quad \text{for all } \mathbf{s}\in\mathcal{S}, \mathbf{a}\in\mathcal{A}(\mathbf{s}), \tag{4.22b}$$

$$\boldsymbol{\theta} \in \mathbb{R}^{|\mathcal{F}|}. \tag{4.22c}$$

We define $\mathcal{V}$ as the set of bounded, real valued functions on the state space $\mathcal{S}$ with componentwise partial order and norm $\|v\| = \sup_{\mathbf{s}\in\mathcal{S}} |v(\mathbf{s})|$. That is, $v \in \mathcal{V}$ if $v : \mathcal{S} \to \mathbb{R}$ and there exists a K_v such that $|v(\mathbf{s})| = |v(k,\mathbf{x},\mathbf{y})| \le K_v$, for all $\mathbf{s} \in \mathcal{S}$. With partial order it is meant that for $u, v \in \mathcal{V}$, $u \ge v$ if $u(\mathbf{s}) \ge v(\mathbf{s})$, for all $\mathbf{s} \in \mathcal{S}$. With this definition in mind we point out the following lemma.

Lemma 4.2 (*Theorem 6.2.2 of [30]*)**.** Suppose there exists a $v \in \mathcal{V}$, for which

$$v(\mathbf{s}) \le \min_{\mathbf{a}\in\mathcal{A}(\mathbf{s})} \left\{ c(\mathbf{s},\mathbf{a}) + \gamma \sum_{\mathbf{s}'\in\mathcal{S}} p(\mathbf{s}'|\mathbf{s},\mathbf{a}) v(\mathbf{s}') \right\}, \text{ for all } \mathbf{s}\in\mathcal{S},$$

then $v(\mathbf{s}) \le v^*(\mathbf{s})$, for all $\mathbf{s} \in \mathcal{S}$.

Here $v^*(\mathbf{s})$ is unique solution to the optimality equations (4.18) as defined in Theorem 4.1.

Proof. See the proof of Theorem 6.2.2 on p.148 of [30] ■

Theorem 4.3. The LP (4.22a)–(4.22c) and its dual both have finite optima for the fitting functions defined in (4.21).

Proof. Strong duality tells us that if a linear programming problem has an optimal solution, so does its dual, and the respective optimal costs are equal (see, e.g., [5, Theorem 4.4]). Thus, it suffices to show that the LP (4.22a)-(4.22c) is feasible and bounded, because then the LP satisfies strong duality and the statement to prove follows from strong duality.

To find a feasible solution $\boldsymbol{\theta}$, we set $\theta_f = 0$, for all $f \in \mathcal{F} \backslash \{1, 2\}$. Then, for $\mathbf{s}_1 \in \mathcal{S}_1$, $\mathbf{a} \in \mathcal{A}(\mathbf{s}_1)$, it follows from the transition probabilities (4.13) that Constraint (4.22b) becomes

$$\theta_1 \left[1 - \gamma \sum_{\mathbf{s}_2 \in \mathcal{S}_2} p(\mathbf{s}_2 | \mathbf{s}_1, \mathbf{a}) \right] \leq c(\mathbf{s}_1, \mathbf{a}),$$

$$\theta_1 (1 - \gamma) \leq c(\mathbf{s}_1, \mathbf{a})$$

Similarly, for $\mathbf{s}_2 \in \mathcal{S}_2$, $\mathbf{a} \in \mathcal{A}(\mathbf{s}_2)$, it follows from the transition probabilities (4.15) that Constraint (4.22b) becomes

$$\theta_2 (1 - \gamma) \leq c(\mathbf{s}_2, \mathbf{a}).$$

Hence a feasible solution of (4.22a)-(4.22c) is given by

$$\theta_1 \leq \min_{\substack{\mathbf{s}_1 \in \mathcal{S}_1, \\ \mathbf{a} \in \mathcal{A}(\mathbf{s}_1)}} \frac{c(\mathbf{s}_1, \mathbf{a})}{(1 - \gamma)}, \quad \theta_2 \leq \min_{\substack{\mathbf{s}_2 \in \mathcal{S}_2, \\ \mathbf{a} \in \mathcal{A}(\mathbf{s}_2)}} \frac{c(\mathbf{s}_2, \mathbf{a})}{(1 - \gamma)} \quad \text{and} \quad \theta_f = 0, \text{ for all other } f \in \mathcal{F}.$$

For any feasible $\boldsymbol{\theta}$, it follows from (4.22b) that

$$\omega(\mathbf{s}, \boldsymbol{\theta}) \leq c(\mathbf{s}, \mathbf{a}) + \gamma \sum_{\mathbf{s}' \in \mathcal{S}} p(\mathbf{s}' | \mathbf{s}, \mathbf{a}) \omega(\mathbf{s}', \boldsymbol{\theta}), \qquad \text{for all } \mathbf{s} \in \mathcal{S}, \, \mathbf{a} \in \mathcal{A}(\mathbf{s})$$

and thus

$$\omega(\mathbf{s}, \boldsymbol{\theta}) \leq \min_{\mathbf{a} \in \mathcal{A}(\mathbf{s})} \left\{ c(\mathbf{s}, \mathbf{a}) + \gamma \sum_{\mathbf{s}' \in \mathcal{S}} p(\mathbf{s}' | \mathbf{s}, \mathbf{a}) \omega(\mathbf{s}', \boldsymbol{\theta}) \right\}, \quad \text{for all } \mathbf{s} \in \mathcal{S}. \tag{4.23}$$

For fixed, feasible $\boldsymbol{\theta}$ we have that $\omega(\mathbf{s}, \boldsymbol{\theta})$ is a real valued function of $\mathbf{s}$. Moreover, it follows from (7.1) that for every $\mathbf{s} = (k, \mathbf{x}, \mathbf{y}) \in \mathcal{S}$ we have that all the components x_{mn}, $m \in \mathcal{M}$, $n \in \mathcal{N}$, and y_t, $t \in \mathcal{T}$, are finite (non-negative) integers. Hence, it follows from (4.21) that for fixed $\boldsymbol{\theta}$, the function $\omega(\mathbf{s}, \boldsymbol{\theta})$ is bounded and $\omega(\mathbf{s}, \boldsymbol{\theta}) \in \mathcal{V}$. We then apply Lemma 4.2 to the inequality in (4.23) to conclude that $\omega(\mathbf{s}, \boldsymbol{\theta}) \leq v^*(\mathbf{s})$, for all $\mathbf{s} \in \mathcal{S}$.

Since $\omega(\mathbf{s}, \boldsymbol{\theta}) \leq v^*(\mathbf{s})$, for all $\mathbf{s} \in \mathcal{S}$, the objective function in (4.22a) is bounded from above for any feasible $\boldsymbol{\theta}$ because we get

$$\sum_{\mathbf{s} \in \mathcal{S}} \eta(\mathbf{s}) \left[\sum_{f \in \mathcal{F}} \theta_f \phi_f(\mathbf{s}) \right] = \sum_{\mathbf{s} \in \mathcal{S}} \eta(\mathbf{s}) \omega(\mathbf{s}, \boldsymbol{\theta}) \leq \sum_{\mathbf{s} \in \mathcal{S}} \eta(\mathbf{s}) v^*(\mathbf{s})$$

■

4.2.1 A further detailed presentation of the ALP and its dual

To solve the MRI appointment scheduling problem we chose the following affine approximation to $v(\mathbf{s}), \mathbf{s} \in \mathcal{S}$

$$v(\mathbf{s}) = v(k, \mathbf{x}, \mathbf{y}) = \begin{cases} V_0 + \sum_{m=1}^{M} \sum_{n=1}^{N} V_{mn} x_{mn} + \sum_{t=1}^{T} W_t y_t, & \text{if } k = 1, \\ \\ \widetilde{V}_0 + \sum_{m=1}^{M} \sum_{n=1}^{N} \widetilde{V}_{mn} x_{mn} + \sum_{t=1}^{T} \widetilde{W}_t y_t & \text{if } k = 2. \end{cases} \tag{4.24}$$

Here $V_0, \widetilde{V}_0 \in \mathbb{R}$, $\mathbf{V} = (V_{mn})_{m \in \mathcal{M}, n \in \mathcal{N}}$, $\widetilde{\mathbf{V}} = (\widetilde{V}_{mn})_{m \in \mathcal{M}, n \in \mathcal{N}} \in \mathbb{R}_+^{M \times N}$ and $\mathbf{W} = (W_t)_{t \in \mathcal{T}}$, $\widetilde{\mathbf{W}} = (\widetilde{W}_t)_{t \in \mathcal{T}} \in \mathbb{R}_+^T$. $\mathbb{R}_+^n$ is the space consisting of all n-tuples of non-negative real numbers, i.e., an element of $\mathbb{R}_+^n$ is a vector $\mathbf{z} = (z_1, z_2, \ldots, z_n)$, where each $z_i \in \mathbb{R}$ and satisfies $z_i \geq 0$.

The value of V_{mn} represents the marginal cost of having an additional appointment slot occupied on MRI scanner m on day n if we are at the first decision epoch of the day and the value of W_t represents the marginal cost of having one more appointment request of type t waiting to be booked at the first decision epoch of the day. Similarly, the value of $\widetilde{V}_{mn}$ represents the marginal cost of having an additional appointment slot occupied on MRI scanner m on day n if we are at the second decision epoch of the day and the value of $\widetilde{W}_t$ represents the marginal cost of having one more appointment request of type t waiting to be booked at the second decision epoch of the day.

Note that this is essentially the same approximation as using the architecture in (4.20) combined with the coefficients $\{\theta_f \in \mathbb{R} : f \in \mathcal{F}\}$ and the fitting functions given in (4.21). However, in (4.24), we use $V_0, \widetilde{V}_0, \mathbf{V}, \widetilde{\mathbf{V}}, \mathbf{W}, \widetilde{\mathbf{W}}$ as coefficients. A slight difference here is that elements of $\mathbf{V}, \widetilde{\mathbf{V}}, \mathbf{W}, \widetilde{\mathbf{W}}$ are restricted to be non-negative real numbers. From the interpretation of these vectors as explained above, it may be clear that it makes no sense if we allow the vector's elements to be negative.

Before we substitute the affine approximation (4.24) into the LP (4.19a)-(4.19c), we first rewrite the LP as this makes the substitution and further analysis easier. Since the state space, defined by (4.2), can be partitioned into states $\mathbf{s}_1 \in \mathcal{S}_1$ corresponding to the first decision epoch on a day and $\mathbf{s}_2 \in \mathcal{S}_2$ corresponding to the second decision epoch, the constraints (4.19b) can be rewritten accordingly. To see this, fix $\mathbf{s}_1 = (1, \mathbf{x}, \mathbf{y}) \in \mathcal{S}_1$ and recall that $\mathbf{a} = (a_{tmn})$, $t \in \mathcal{T}$, $m \in \mathcal{M}$, $n \in \mathcal{N}$. It then follows from the transition probabilities defined by (4.13) that for $\mathbf{s}_1 \in \mathcal{S}_1$ we can rewrite constraints (4.19b) as

$$c(\mathbf{s}_1, \mathbf{a}) + \gamma \sum_{\mathbf{y}' \in \mathcal{Y}_2} p_2(\mathbf{y}') v \left\{ \left(2, \left(x_{12} + \sum_{t \in \mathcal{T}(1)} d_t a_{t12}, \ldots, x_{MN} + \sum_{t \in \mathcal{T}(M)} d_t a_{tMN}, 0 \right), \mathbf{y}' \right) \right\} \geq v(\mathbf{s}_1),$$

for all $\mathbf{a} \in \mathcal{A}(\mathbf{s}_1)$.

Similarly, it follows from the transition probabilities defined by (4.15) that for $\mathbf{s}_2 = (2, \mathbf{x}, \mathbf{y}) \in \mathcal{S}_2$ constraints (4.19b) can be rewritten as

$$c(\mathbf{s}_2, \mathbf{a}) + \gamma \sum_{\mathbf{y}' \in \mathcal{Y}_1} p_1(\mathbf{y}') v \left\{ \left(1, \left(x_{13} + \sum_{t \in \mathcal{T}(1)} d_t a_{t13}, \ldots, x_{M,N-2} + \sum_{t \in \mathcal{T}(M)} d_t a_{tM,N-2}, 0, 0 \right), \mathbf{y}' \right) \right\} \geq v(\mathbf{s}_2),$$

for all $\mathbf{a} \in \mathcal{A}(\mathbf{s}_2)$.

Also based on the definition of the state space (given in Eq. (4.2)), we can rewrite the objective function (4.19a) as

$$\max_{\mathbf{v} \in \mathbb{R}^{|\mathcal{S}|}} \sum_{\mathbf{s}_1 \in \mathcal{S}_1} \alpha(\mathbf{s}_1) v(\mathbf{s}_1) + \sum_{\mathbf{s}_2 \in \mathcal{S}_2} \beta(\mathbf{s}_2) v(\mathbf{s}_2), \quad \alpha(\mathbf{s}_1) > 0, \text{ for all } \mathbf{s}_1 \in \mathcal{S}_1, \text{ and}$$

$\beta(\mathbf{s}_2) > 0$, for all $\mathbf{s}_2 \in \mathcal{S}_2$.

All together, by rearranging terms in the two preceding inequalities, we can rewrite the LP (4.19a)-(4.19c) as

$$\max_{\mathbf{v}} \quad \sum_{\mathbf{s}_1 \in \mathcal{S}_1} \alpha(\mathbf{s}_1) v(\mathbf{s}_1) + \sum_{\mathbf{s}_2 \in \mathcal{S}_2} \beta(\mathbf{s}_2) v(\mathbf{s}_2), \tag{4.25a}$$

s.t.

$$c(\mathbf{s}_1, \mathbf{a}_1) \geq v(\mathbf{s}_1) - \gamma \sum_{\mathbf{y}' \in \mathcal{Y}_2} p_2(\mathbf{y}') v \left\{ \left(2, \left(x_{12} + \sum_{t \in \mathcal{T}(1)} d_t a_{t12}, \ldots, x_{MN} + \sum_{t \in \mathcal{T}(M)} d_t a_{tMN}, 0 \right), \mathbf{y}' \right) \right\},$$

$$\text{for all } \mathbf{s}_1 = (1, \mathbf{x}, \mathbf{y}) \in \mathcal{S}_1, \mathbf{a}_1 \in \mathcal{A}(\mathbf{s}_1), \tag{4.25b}$$

$$c(\mathbf{s}_2, \mathbf{a}) \geq v(\mathbf{s}_2) - \gamma \sum_{\mathbf{y}' \in \mathcal{Y}_1} p_1(\mathbf{y}') v \left\{ \left(1, \left(x_{13} + \sum_{t \in \mathcal{T}(1)} d_t a_{t13}, \ldots, x_{M,N-2} + \sum_{t \in \mathcal{T}(M)} d_t a_{tM,N-2}, 0, 0 \right), \mathbf{y}' \right) \right\},$$

$$\text{for all } \mathbf{s}_2 = (2, \mathbf{x}, \mathbf{y}) \in \mathcal{S}_2, \mathbf{a}_2 \in \mathcal{A}(\mathbf{s}_2), \tag{4.25c}$$

$$\mathbf{v} \in \mathbb{R}^{|\mathcal{S}|} = \mathbb{R}^{|\mathcal{S}_1 \cup \mathcal{S}_2|}. \tag{4.25d}$$

We normalize $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ to $\sum_{\mathbf{s}_1 \in \mathcal{S}_1} \alpha(\mathbf{s}_1) = 1$ and $\sum_{\mathbf{s}_2 \in \mathcal{S}_2} \beta(\mathbf{s}_2) = 1$ and consider these as exogenous probability distributions over the feasible states at the first and second decision epoch on a day, respectively.

At this point, we substitute our affine approximation to $v(\mathbf{s})$, given by (4.24), into the LP (4.25a)-(4.25d). We first consider the objective function:

$$\sum_{\mathbf{s}_1 \in \mathcal{S}_1} \alpha(\mathbf{s}_1) v(\mathbf{s}_1) + \sum_{\mathbf{s}_2 \in \mathcal{S}_2} \beta(\mathbf{s}_2) v(\mathbf{s}_2) = \sum_{\mathbf{s}_1 \in \mathcal{S}_1} \alpha(\mathbf{s}_1) \left(V_0 + \sum_{m=1}^{M} \sum_{n=1}^{N} V_{mn} x_{mn}(\mathbf{s}_1) + \sum_{t=1}^{T} W_t y_t(\mathbf{s}_1) \right)$$
$$+ \sum_{\mathbf{s}_2 \in \mathcal{S}_2} \beta(\mathbf{s}_2) \left(\widetilde{V}_0 + \sum_{m=1}^{M} \sum_{n=1}^{N} \widetilde{V}_{mn} x_{mn}(\mathbf{s}_2) + \sum_{t=1}^{T} \widetilde{W}_t y_t(\mathbf{s}_2) \right).$$

Here $x_{mn}(\mathbf{s}_1)$ and $x_{mn}(\mathbf{s}_2)$ respectively denotes component x_{mn} of state vector $\mathbf{s}_1 \in \mathcal{S}_1$ and $\mathbf{s}_2 \in \mathcal{S}_2$. Since we defined $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ as probability distributions, we can simplify the right-hand side of expression above to

$$V_0 + \widetilde{V}_0 + \sum_{m=1}^{M} \sum_{n=1}^{N} \left\{ V_{mn} \mathbb{E}_\alpha[x_{mn}] + \widetilde{V}_{mn} \mathbb{E}_\beta[x_{mn}] \right\} + \sum_{t=1}^{T} \left\{ W_t \mathbb{E}_\alpha[y_t] + \widetilde{W}_t \mathbb{E}_\beta[y_t] \right\},$$

where

$$\mathbb{E}_\alpha[x_{mn}] = \sum_{\mathbf{s}\in\mathcal{S}} \alpha(\mathbf{s}) x_{mn}(\mathbf{s}), \qquad \mathbb{E}_\alpha[y_t] = \sum_{\mathbf{s}\in\mathcal{S}} \alpha(\mathbf{s}) y_t(\mathbf{s}),$$
$$\mathbb{E}_\beta[x_{mn}] = \sum_{\mathbf{s}\in\mathcal{S}} \beta(\mathbf{s}) x_{mn}(\mathbf{s}), \qquad \mathbb{E}_\beta[y_t] = \sum_{\mathbf{s}\in\mathcal{S}} \beta(\mathbf{s}) y_t(\mathbf{s}).$$

Next, fix $\mathbf{s}_1 = (1, \mathbf{x}, \mathbf{y}) \in \mathcal{S}_1$ and $\mathbf{a} \in \mathcal{A}(\mathbf{s}_1)$. Then, for this specific state-action pair $(\mathbf{s}_1, \mathbf{a})$, we substitute the affine approximation into Constraint (4.25b) to obtain

$$\begin{aligned}
c(\mathbf{s}_1, \mathbf{a}) \geq{}& V_0 + \sum_{m=1}^{M} \sum_{n=1}^{N} V_{mn} x_{mn} + \sum_{t=1}^{T} W_t y_t \\
& - \gamma \sum_{\mathbf{y}'\in\mathcal{Y}_2} p_2(\mathbf{y}') \left(\widetilde{V}_0 + \sum_{m=1}^{M} \sum_{n=1}^{N-1} \widetilde{V}_{mn} \left[x_{m,n+1} + \sum_{t\in\mathcal{T}(m)} d_t a_{tm,n+1} \right] + \sum_{t=1}^{T} \widetilde{W}_t y'_t \right) \\
={}& V_0 + \sum_{m=1}^{M} \sum_{n=1}^{N} V_{mn} x_{mn} + \sum_{t=1}^{T} W_t y_t \\
& - \gamma \left(\widetilde{V}_0 + \sum_{m=1}^{M} \sum_{n=1}^{N-1} \widetilde{V}_{mn} \left[x_{m,n+1} + \sum_{t\in\mathcal{T}(m)} d_t a_{tm,n+1} \right] \right) - \gamma \sum_{t=1}^{T} \widetilde{W}_t \mathbb{E}_{p_1}[y_t],
\end{aligned}$$

where we substituted the equality $\sum_{\mathbf{y}'\in\mathcal{Y}_2} p_2(\mathbf{y}') \sum_{t=1}^{T} \widetilde{W}_t y_t = \sum_{t=1}^{T} W_t \mathbb{E}_{p_2}[y_t]$ into the right-hand side of the last inequality. We show how to get this equality in the derivation that leads to (4.26). We now proceed with rearranging terms in the above to get

$$\begin{aligned}
c(\mathbf{s}_1, \mathbf{a}) \geq{}& V_0 - \gamma\widetilde{V}_0 + \sum_{m=1}^{M} \sum_{n=1}^{N} V_{mn} x_{mn} - \gamma \sum_{m=1}^{M} \sum_{n=1}^{N-1} \widetilde{V}_{mn} \left[x_{m,n+1} + \sum_{t\in\mathcal{T}(m)} d_t a_{tm,n+1} \right] \\
& + \sum_{t=1}^{T} W_t y_t - \gamma \sum_{t=1}^{T} \widetilde{W}_t \mathbb{E}_{p_2}[y_t] \\
={}& V_0 - \gamma\widetilde{V}_0 + \sum_{m=1}^{M} \sum_{n=1}^{N} \left\{ V_{mn} x_{mn} - \widetilde{V}_{mn} \mu_{mn}(\mathbf{s}_1, \mathbf{a}) \right\} + \sum_{t=1}^{T} \left\{ W_t y_t - \widetilde{W}_t \lambda_t \right\},
\end{aligned}$$

where we defined

$$\mu_{mn}(\mathbf{s}_1, \mathbf{a}) = \begin{cases} \gamma(x_{m,n+1}(\mathbf{s}_1) + \sum_{t \in \mathcal{T}(m)} d_t a_{tm,n+1}) & \text{for all } m = 1, , \dots, M, n = 1, \dots, N-1, \\ 0 & \text{for all } m = 1, \dots, M \text{ and } n = N, \end{cases}$$

and

$$\lambda_t = \gamma \mathbb{E}_{p_2}[y_t].$$

We proceed with the derivation of the equality $\sum_{\mathbf{y} \in \mathcal{Y}_2} p_2(\mathbf{y}) \sum_{t=1}^{T} \widetilde{W}_t y_t = \sum_{t=1}^{T} W_t \mathbb{E}_{p_2}[y_t]$:

$$\begin{aligned}
\sum_{\mathbf{y} \in \mathcal{Y}_2} p_2(\mathbf{y}) \sum_{t=1}^{T} \widetilde{W}_t y_t &= \sum_{y_1'=0}^{Q_1^1} \sum_{y_2=0}^{Q_1^1} \cdots \sum_{y_T=0}^{Q_T^1} p_2(y_1, y_2, \dots, y_T) \left(\widetilde{W}_1 y_1 + \widetilde{W}_2 y_2 + \dots \widetilde{W}_T y_T \right) \\
&= \sum_{y_1=0}^{Q_1^1} \sum_{y_2=0}^{Q_1^1} \cdots \sum_{y_{T-1}=0}^{Q_{T-1}^1} \left\{ \widetilde{W}_1 y_1 \left(\sum_{y_T=0}^{Q_T^1} p_2(y_1, y_2, \dots, y_T) \right) + \right. \\
&\quad + \widetilde{W}_2 y_2 \left(\sum_{y_T=0}^{Q_T^1} p_2(y_1, y_2, \dots, y_T) \right) + \cdots \\
&\quad \left. + \widetilde{W}_T \left(\sum_{y_T'=0}^{Q_T^1} y_T p_2(y_1, y_2, \dots, y_T) \right) \right\} \\
&= \widetilde{W}_1 \sum_{y_1=0}^{Q_1^1} y_1 \left(\sum_{y_2=0}^{Q^1} \cdots \sum_{y_T=0}^{Q_T^1} p_2(y_1, y_2, \dots, y_T) \right) \\
&\quad + \widetilde{W}_2 \sum_{y_1=0}^{Q_1^1} \sum_{y_2=0}^{Q_2^1} y_2 \left(\sum_{y_3=0}^{Q_3^1} \cdots \sum_{y_T=0}^{Q_T^1} p_2(y_1, y_2, \dots, y_T) \right) + \cdots \\
&\quad + \widetilde{W}_T \sum_{y_1=0}^{Q_1^1} \sum_{y_2=0}^{Q_1^1} \cdots \sum_{y_T=0}^{Q_T^1} y_T p_2(y_1, y_2, \dots, y_T).
\end{aligned}$$

Recognizing that $p_2(y_1) = \sum_{y_2=0}^{Q_1^1} \cdots \sum_{y_T=0}^{Q_T^1} p_2(y_1, y_2, \dots, y_T)$ and in general

$p(y_t) = \sum_{y_t' \neq y_t} \sum_{y_t'=0}^{Q_t^1} p_2(y_1, y_2, \ldots, y_T)$, we obtain

$$\sum_{\mathbf{y} \in \mathcal{Y}_2} p_2(\mathbf{y}) \sum_{t=1}^{T} \widetilde{W}_t y_t = \widetilde{W}_1 \sum_{y_1'=0}^{Q_1^1} y_1 p(y_1) + \widetilde{W}_2 \sum_{y_2=0}^{Q_2^1} y_2 p_2(y_2) + \cdots \widetilde{W}_T \sum_{y_T=0}^{Q_T^1} y_T p_2(y_T)$$

$$= \sum_{t=1}^{T} \widetilde{W}_t \mathbb{E}_{p_2}[y_t]. \tag{4.26}$$

Substituting the probability distribution p_1 for p_2 and $\mathcal{Y}_1$ for $\mathcal{Y}_2$ in the derivation that led to (4.26), it follows that $\sum_{\mathbf{y} \in \mathcal{Y}_1} p_1(\mathbf{y}) \sum_{t=1}^{T} W_t y_t = \sum_{t=1}^{T} W_t \mathbb{E}_{p_1}[y_t]$.

Next, fix $\mathbf{s}_2 = (2, \mathbf{x}, \mathbf{y}) \in \mathcal{S}_2$ and $\mathbf{a} \in \mathcal{A}(\mathbf{s}_2)$. Then, for this specific state-action pair $(\mathbf{s}_1, \mathbf{a})$, we substitute the affine approximation into Constraint (4.25c) to obtain

$$\begin{aligned}
c(\mathbf{s}_2, \mathbf{a}) \geq\ & \widetilde{V}_0 + \sum_{m=1}^{M} \sum_{n=1}^{N} \widetilde{V}_{mn} x_{mn} + \sum_{t=1}^{T} \widetilde{W}_t y_t \\
& - \gamma \sum_{\mathbf{y}' \in \mathcal{Y}_1} p_1(\mathbf{y}') \left(V_0 + \sum_{m=1}^{M} \sum_{n=1}^{N-2} V_{mn} \left[x_{m,n+2} + \sum_{t \in \mathcal{T}(m)} d_t a_{tm,n+2} \right] + \sum_{t=1}^{T} W_t y_t' \right) \\
=\ & \widetilde{V}_0 - \gamma V_0 + \sum_{m=1}^{M} \sum_{n=1}^{N} \widetilde{V}_{mn} x_{mn} - \gamma \sum_{m=1}^{M} \sum_{n=1}^{N-2} V_{mn} \left[x_{m,n+2} + \sum_{t \in \mathcal{T}(m)} d_t a_{tm,n+2} \right] \\
& + \sum_{t=1}^{T} \widetilde{W}_t y_t - \gamma \sum_{t=1}^{T} W_t \mathbb{E}_{p_1}[y_t] \\
=\ & \widetilde{V}_0 - \gamma V_0 + \sum_{m=1}^{M} \sum_{n=1}^{N} \left\{ \widetilde{V}_{mn} x_{mn} - V_{mn} \tilde{\mu}_{mn}(\mathbf{s}_2, \mathbf{a}) \right\} + \sum_{t=1}^{T} \left\{ \widetilde{W}_t y_t - W_t \tilde{\lambda}_t \right\},
\end{aligned}$$

where we defined

$$\tilde{\mu}_{mn}(\mathbf{s}_2, \mathbf{a}) = \begin{cases} \gamma(x_{m,n+2}(\mathbf{s}_2) + \sum_{t \in \mathcal{T}(m)} d_t a_{tm,n+2}) & \text{for all } m = 1, , \ldots, M \text{ and } n = 1, \ldots, N-2, \\ 0 & \text{for all } m = 1, \ldots, M,\ n = N-1, N, \end{cases}$$

and

$$\tilde{\lambda}_t = \gamma \mathbb{E}_{p_1}[y_t].$$

At this point, we substituted the affine approximation to $v(\mathbf{s})$ given by (4.24) into all parts of the LP (4.25a)-(4.25d). The result is the ALP in (4.27a)-(4.27d), whose solution leads to an approximate optimal policy for the MRI scheduling MDP.

$$\max_{\substack{V_0,\widetilde{V}_0,\mathbf{V},\\ \widetilde{\mathbf{V}},\mathbf{W},\widetilde{\mathbf{W}}}} \quad V_0+\widetilde{V}_0+\sum_{m=1}^{M}\sum_{n=1}^{N}\left\{V_{mn}\mathbb{E}_\alpha[x_{mn}]+\widetilde{V}_{mn}\mathbb{E}_\beta[x_{mn}]\right\}+\sum_{t=1}^{T}\left\{W_t\mathbb{E}_\alpha[y_t]+\widetilde{W}_t\mathbb{E}_\beta[y_t]\right\}, \tag{4.27a}$$

$$\text{s.t.} \quad V_0-\gamma\widetilde{V}_0+\sum_{m=1}^{M}\sum_{n=1}^{N}\left\{V_{mn}x_{mn}-\widetilde{V}_{mn}\mu_{mn}(\mathbf{s}_1,\mathbf{a})\right\}+\sum_{t=1}^{T}\left\{W_ty_t-\widetilde{W}_t\lambda_t\right\}\leq c(\mathbf{s}_1,\mathbf{a}),$$

$$\text{for all } \mathbf{s}_1=(1,\mathbf{x},\mathbf{y})\in\mathcal{S}_1, \mathbf{a}_1\in\mathcal{A}(\mathbf{s}_1), \tag{4.27b}$$

$$\widetilde{V}_0-\gamma V_0+\sum_{m=1}^{M}\sum_{n=1}^{N}\left\{\widetilde{V}_{mn}x_{mn}-V_{mn}\tilde{\mu}_{mn}(\mathbf{s}_2,\mathbf{a})\right\}+\sum_{t=1}^{T}\left\{\widetilde{W}_ty_t-W_t\tilde{\lambda}_t\right\}\leq c\left(\mathbf{s}_2,\mathbf{a}\right),$$

$$\text{for all } \mathbf{s}_2=(2,\mathbf{x},\mathbf{y})\in\mathcal{S}_2, \mathbf{a}_2\in\mathcal{A}(\mathbf{s}_2), \tag{4.27c}$$

$$V_0,\widetilde{V}_0\in\mathbb{R}, \quad \mathbf{V},\widetilde{\mathbf{V}}\in\mathbb{R}_+^{M\times N}, \quad \mathbf{W},\widetilde{\mathbf{W}}\in\mathbb{R}_+^{T}, \tag{4.27d}$$

where

$$\mathbb{E}_\alpha[x_{mn}]=\sum_{\mathbf{s}\in\mathcal{S}}\alpha(\mathbf{s})x_{mn}(\mathbf{s}), \qquad \mathbb{E}_\alpha[y_t]=\sum_{\mathbf{s}\in\mathcal{S}}\alpha(\mathbf{s})y_t(\mathbf{s}),$$

$$\mathbb{E}_\beta[x_{mn}]=\sum_{\mathbf{s}\in\mathcal{S}}\beta(\mathbf{s})x_{mn}(\mathbf{s}), \qquad \mathbb{E}_\beta[y_t]=\sum_{\mathbf{s}\in\mathcal{S}}\beta(\mathbf{s})y_t(\mathbf{s}),$$

$$\lambda_t=\gamma\mathbb{E}_{p_2}[y_t],$$

$$\tilde{\lambda}_t=\gamma\mathbb{E}_{p_1}[y_t],$$

$$\mu_{mn}(\mathbf{s}_1,\mathbf{a})=\begin{cases}\gamma(x_{m,n+1}(\mathbf{s}_1)+\sum_{t\in\mathcal{T}(m)}d_ta_{tm,n+1}), & \text{for all } m=1,,\ldots,M \text{ and} n=1,\ldots,N-1,\\ 0, & \text{for all } m=1,\ldots,M \text{ and } n=N,\end{cases}$$

$$\tilde{\mu}_{mn}(\mathbf{s}_2,\mathbf{a})=\begin{cases}\gamma(x_{m,n+2}(\mathbf{s}_2)+\sum_{t\in\mathcal{T}(m)}d_ta_{tm,n+2}), & \text{for all } m=1,\ldots,M \text{ and} n=1,\ldots,N-2,\\ 0, & \text{for all } m=1,\ldots,M \text{ and} n=N-1,N.\end{cases}$$

The (approximate) linear programming model in (4.25a)-(4.25d) has a tractable number of variables, $2(1+M\times N+T)$, but still an intractable number of constraints (one for every state-action pair). For this reason, we solve its dual (4.28a)-(4.28h) using column generation.

$$\min_{\mathbf{X}} \sum_{\substack{\mathbf{s}\in\mathcal{S},\\ \mathbf{a}\in\mathcal{A}(\mathbf{s})}} c(\mathbf{s},\mathbf{a})X(\mathbf{s},\mathbf{a}), \tag{4.28a}$$

$$\text{s.t.} \sum_{\substack{\mathbf{s}_1\in\mathcal{S}_1,\\ \mathbf{a}_1\in\mathcal{A}(\mathbf{s}_1)}} X(\mathbf{s}_1,\mathbf{a}_1) \;-\; \gamma \sum_{\substack{\mathbf{s}_2\in\mathcal{S}_2,\\ \mathbf{a}_2\in\mathcal{A}(\mathbf{s}_2)}} X(\mathbf{s}_2,\mathbf{a}_2) = 1 \tag{4.28b}$$

$$\sum_{\substack{\mathbf{s}_2\in\mathcal{S}_2,\\ \mathbf{a}_2\in\mathcal{A}(\mathbf{s}_2)}} X(\mathbf{s}_2,\mathbf{a}_2) \;-\; \gamma \sum_{\substack{\mathbf{s}_1\in\mathcal{S}_1,\\ \mathbf{a}_1\in\mathcal{A}(\mathbf{s}_1)}} X(\mathbf{s}_1,\mathbf{a}_1) = 1 \tag{4.28c}$$

$$\sum_{\substack{\mathbf{s}_1\in\mathcal{S}_1,\\ \mathbf{a}_1\in\mathcal{A}(\mathbf{s}_1)}} x_{mn}(\mathbf{s}_1)X(\mathbf{s}_1,\mathbf{a}_1) \;-\; \sum_{\substack{\mathbf{s}_2\in\mathcal{S}_2,\\ \mathbf{a}_2\in\mathcal{A}(\mathbf{s}_2)}} \tilde{\mu}_{mn}(\mathbf{s}_2,\mathbf{a}_2)X(\mathbf{s}_2,\mathbf{a}_2) \geq \mathbb{E}_\alpha[x_{mn}], \quad \forall m\in\mathcal{M}, \forall n\in\mathcal{N} \tag{4.28d}$$

$$\sum_{\substack{\mathbf{s}_2\in\mathcal{S}_2,\\ \mathbf{a}_2\in\mathcal{A}(\mathbf{s}_2)}} x_{mn}(\mathbf{s}_2)X(\mathbf{s}_2,\mathbf{a}_2) \;-\; \sum_{\substack{\mathbf{s}_1\in\mathcal{S}_1,\\ \mathbf{a}_1\in\mathcal{A}(\mathbf{s}_1)}} \mu_{mn}(\mathbf{s}_1,\mathbf{a}_1)X(\mathbf{s}_1,\mathbf{a}_1) \geq \mathbb{E}_\beta[x_{mn}], \quad \forall m\in\mathcal{M}, \forall n\in\mathcal{N} \tag{4.28e}$$

$$\sum_{\substack{\mathbf{s}_1\in\mathcal{S}_1,\\ \mathbf{a}_1\in\mathcal{A}(\mathbf{s}_1)}} y_t(\mathbf{s}_1)X(\mathbf{s}_1,\mathbf{a}_1) \;-\; \tilde{\lambda}_t \sum_{\substack{\mathbf{s}_2\in\mathcal{S}_2,\\ \mathbf{a}_2\in\mathcal{A}(\mathbf{s}_2)}} X(\mathbf{s}_2,\mathbf{a}_2) \geq \mathbb{E}_\alpha[y_t], \qquad \forall t\in\mathcal{T} \tag{4.28f}$$

$$\sum_{\substack{\mathbf{s}_2\in\mathcal{S}_2,\\ \mathbf{a}_2\in\mathcal{A}(\mathbf{s}_2)}} y_t(\mathbf{s}_2)X(\mathbf{s}_2,\mathbf{a}_2) \;-\; \lambda_t \sum_{\substack{\mathbf{s}_1\in\mathcal{S}_1,\\ \mathbf{a}_1\in\mathcal{A}(\mathbf{s}_1)}} X(\mathbf{s}_1,\mathbf{a}_1) \geq \mathbb{E}_\beta[y_t], \qquad \forall t\in\mathcal{T} \tag{4.28g}$$

$$X(\mathbf{s},\mathbf{a}) \in \mathbb{R}_+, \text{ for all } \mathbf{s}\in\mathcal{S}, \mathbf{a}\in\mathcal{A}(\mathbf{s}). \tag{4.28h}$$

The dual variable $X(\mathbf{s},\mathbf{a})$ can be interpreted as the frequency of taking action $\mathbf{a}$ when in state $\mathbf{s}$.

4.2.2 Column generation

Column generation (sometimes also called *delayed column generation*) is a method for dealing with LPs that have a huge number of variables compared to a moderate number of constraints. For an introduction to column generation, we refer to [5, Chapter 6]. Column generation finds the optimal solution to (4.28a)-(4.28h), the *master problem* (MP), starting with a small set of feasible state-action pairs (or actually with their corresponding variables $X(\mathbf{s},\mathbf{a})$) and then iteratively adding new state-action pairs that lead to a better solution. The algorithm iterates until no such state-action pairs can be found.

We introduce the subset of states $\mathcal{S}' \subset \mathcal{S}$ and subsets of actions $\mathcal{A}'(\mathbf{s}) \subseteq \mathcal{A}(\mathbf{s})$, for $\mathbf{s}\in\mathcal{S}'$ and observe that, based on the definition of $\mathcal{S}$ (see (4.2)), we can define $\mathcal{S}_1'$ and $\mathcal{S}_2'$ and let $\mathcal{S}' = \mathcal{S}_1' \cup \mathcal{S}_2'$. Analogous to this we have $\mathcal{A}'(\mathbf{s}_1) \subseteq \mathcal{A}(\mathbf{s}_1)$ and $\mathcal{A}'(\mathbf{s}_2) \subseteq \mathcal{A}(\mathbf{s}_2)$ as a subset of feasible actions for the states $\mathbf{s}_1 \in \mathcal{S}_1'$ and $\mathbf{s}_2 \in \mathcal{S}_2'$, respectively. We then define the *Restricted Master Problem* (RMP) as the LP (4.28a)-(4.28h) with the modification that in (4.28a)-(4.28h) we included the variables $X(\mathbf{s},\mathbf{a})$ for all states $\mathbf{s}\in\mathcal{S}$ and feasible actions $\mathbf{a}\in\mathcal{A}(\mathbf{s})$ and in the RMP we only include the variables $X(\mathbf{s},\mathbf{a})$ for the states $\mathbf{s}\in\mathcal{S}'$ and the actions $\mathbf{a}\in\mathcal{A}'(\mathbf{s})$. This RMP is solved in each iteration of the column generation algorithm. After the RMP is solved, we look for a state-action pair $(\mathbf{s}, \mathbf{a})$ that, when we add the corresponding variable $X(\mathbf{s}, \mathbf{a})$ to the RMP, leads to a better objective function value of the RMP. If such a state-action pair $(\mathbf{s}, \mathbf{a})$ exists, the state is added to $\mathcal{S}'$ and the action to $\mathcal{A}'(\mathbf{s})$ and the algorithm proceeds to the next iteration. In this next iteration, the RMP is solved again, only this time with the updated $\mathcal{S}'$ and $\mathcal{A}'(\mathbf{s})$. If such a state-action pair does not exists, we conclude that the objective function value of the RMP cannot be further improved and we found the optimal solution to the MP (4.28a)-(4.28h) (which is equal to the solution of the RMP of the current iteration).

The problem still faced is how to generate a state-action pair $(\mathbf{s}, \mathbf{a})$, that improve the solution of the RMP. From column generation theory we know that we should generate state-action pairs $(\mathbf{s}, \mathbf{a})$ for which the corresponding variables $X(\mathbf{s}, \mathbf{a})$ have negative reduced costs. To generate such state-action pairs, we should not try all possible pairs and compute the reduced costs of the corresponding variable $X(\mathbf{s}, \mathbf{a})$. First, because there are many of them and second, and this is crucial, if we try all possible pairs and compute the reduced costs of the corresponding variable, then we have not saved anything compared to solving the whole LP with all variables. The idea is to find a new variable by solving a second optimization problem, called the *pricing problem*. From the dual perspective (that is from the perspective of (4.28a)-(4.28h)) we want to minimize the reduced costs of the variables $X(\mathbf{s}, \mathbf{a})$. From a primal perspective (that is from the perspective of (4.27a)-(4.27d)), this goal is equal to finding the most violated (primal) constraint. (If we do not have generated all variables of the dual, we do not have all constraints of the primal. If a solution is, though optimal for the (dual) RMP of the current iteration, not an optimal solution for the MP, then there exist not-yet-generated primal constraints that are violated by the (primal) solution of RMP of the current iteration.)

Given the dual values associated with the solution of the current RMP, V_0, $\widetilde{V}_0$, $\mathbf{V}$, $\widetilde{\mathbf{V}}$ $\mathbf{W}$, $\widetilde{\mathbf{W}}$, the pricing problem used to identify the state-action pair associated with the most violated primal constraint is given by:

$$\operatorname*{arg\,min}_{\substack{\mathbf{s}_1 \in \mathcal{S}_1, \mathbf{a}_1 \in \mathcal{A}(\mathbf{s}_1), \\ \mathbf{s}_2 \in \mathcal{S}_2, \mathbf{a}_2 \in \mathcal{A}(\mathbf{s}_2)}} \{g_1(\mathbf{s}_1, \mathbf{a}_1), g_2(\mathbf{s}_2, \mathbf{a}_2)\}, \tag{4.29}$$

where

$$\begin{aligned} g_1(\mathbf{s}_1, \mathbf{a}_1) = c(\mathbf{s}_1, \mathbf{a}_1) - V_0 + \gamma \widetilde{V}_0 - \sum_{m=1}^{M} \sum_{n=1}^{N} \left\{ V_{mn} x_{mn}(\mathbf{s}_1) - \widetilde{V}_{mn} \mu_{mn}(\mathbf{s}_1, \mathbf{a}_1) \right\} \\ - \sum_{t=1}^{T} \left\{ W_t y_t(\mathbf{s}_1) - \widetilde{W}_t \lambda_t \right\}, \\ g_2(\mathbf{s}_2, \mathbf{a}_2) = c(\mathbf{s}_2, \mathbf{a}_2) - \widetilde{V}_0 + \gamma V_0 - \sum_{m=1}^{M} \sum_{n=1}^{N} \left\{ \widetilde{V}_{mn} x_{mn}(\mathbf{s}_2) - V_{mn} \tilde{\mu}_{mn}(\mathbf{s}_2, \mathbf{a}_2) \right\} \\ - \sum_{t=1}^{T} \left\{ \widetilde{W}_t y_t(\mathbf{s}_2) - W_t \tilde{\lambda}_t \right\}. \end{aligned}$$

The column generation algorithm iterates until either no primal constraint is violated ore one is "close enough" to optimality to quit. Here, the first option is that the algorithm iterates as long as we find solutions to the pricing problem with objective value strictly less than zero. The implementation of the pricing model involves V_0, $\widetilde{V}_0$, $\mathbf{V}$, $\widetilde{\mathbf{V}}$ $\mathbf{W}$, $\widetilde{\mathbf{W}}$, which are the shadow prices corresponding to the constraints of the dual given in (4.28a)-(4.28h). Since most LP solvers calculate the shadow prices during each solve iteration, they can be directly accessed. However, it is important to allow numerical inaccuracies in the computed shadow prices. For this reason, it is generally advisable to use a small tolerance parameter $\delta < 0$ when verifying whether a new state-action pair will lead to improvement. The mathematical condition to be verified for progress then becomes to check whether the objective value of the solution of the pricing problem is greater than or equal to the tolerance δ. If this is the case, one is "close enough" to optimality to stop iterating. The value of δ is typically in the order of -10^{-4}.

Reformulation of the pricing problem

The pricing problem in (4.29) suggests that we need to compute $g_1(\mathbf{s}_1, \mathbf{a}_1)$, for all $\mathbf{s}_1 \in \mathcal{S}_1$ and $\mathbf{a}_1 \in \mathcal{A}(\mathbf{s}_1)$ and $g_2(\mathbf{s}_2, \mathbf{a}_2)$, for all $\mathbf{s}_2 \in \mathcal{S}_2$ and $\mathbf{a}_2 \in \mathcal{A}(\mathbf{s}_2)$. However, this requires that much computation and storage capacity that we reformulate the pricing problem (4.29) as two "actual" optimization problems. For this, we first define

$$\begin{aligned} g_1^* &= \min_{\mathbf{s}_1 \in \mathcal{S}_1, \mathbf{a}_1 \in \mathcal{A}(\mathbf{s}_1)} g_1(\mathbf{s}_1, \mathbf{a}_1) \\ &= \min_{\substack{(1,\mathbf{x},\mathbf{y}) \in \mathcal{S}_1 \\ \mathbf{a}_1 \in \mathcal{A}(1,\mathbf{x},\mathbf{y})}} \left\{ c((1,\mathbf{x},\mathbf{y}),\mathbf{a}_1) - \sum_{m=1}^{M} \sum_{n=1}^{N} \left\{ V_{mn} x_{mn} - \widetilde{V}_{mn} \mu_{mn}((1,\mathbf{x},\mathbf{y}),\mathbf{a}_1) \right\} - \sum_{t=1}^{T} W_t y_t \right\} - U_1, \end{aligned} \tag{4.30}$$

where the constant $U_1 = V_0 - \gamma \widetilde{V}_0 - \sum_{t=1}^{T} \widetilde{W}_t \lambda_t$. Similarly, we define

$$g_2^* = \min_{\substack{(2,\mathbf{x},\mathbf{y}) \in \mathcal{S}_2 \\ \mathbf{a}_2 \in \mathcal{A}(2,\mathbf{x},\mathbf{y})}} \left\{ c((2,\mathbf{x},\mathbf{y}),\mathbf{a}_2) - \sum_{m=1}^{M} \sum_{n=1}^{N} \left\{ \widetilde{V}_{mn} x_{mn} - V_{mn} \tilde{\mu}_{mn}((2,\mathbf{x},\mathbf{y}),\mathbf{a}_2) \right\} - \sum_{t=1}^{T} \widetilde{W}_t y_t \right\} - U_2, \tag{4.31}$$

where the constant $U_2 = \widetilde{V}_0 - \gamma V_0 - \sum_{t=1}^{T} W_t \tilde{\lambda}_t$.

The optimization problem in (4.30) is almost an integer linear program (ILP) with decision variables $(\mathbf{x}, \mathbf{y}) \in \mathcal{S}_1$ (as defined in (7.1)) and $\mathbf{a}_1 \in \mathcal{A}(1, \mathbf{x}, \mathbf{y})$ (defined by the constraints (4.3)-(4.7)), except that the direct costs incorporates the non-linear part f^{AS}. However, we can linearize f^{AS}, and thereby the objective in (4.30), to create an ILP that we can solve as the pricing problem corresponding to state-action pairs feasible at the first decision epoch of a day. To do this, we first introduce the additional binary variables ϕ_{mn} and ξ_{mn}, for $m \in \mathcal{M}$ and $n \in \mathcal{N}$, and the constraints:

$$x_{mn} + \sum_{t \in \mathcal{T}(m)} d_t a_{tmn} \geq \phi_{mn} C_{mn}^{R,1} \tag{4.32}$$

$$x_{mn} + \sum_{t \in \mathcal{T}(m)} d_t a_{tmn} \leq (1 - \phi_{mn}) C_{mn}^{R,1} + \phi_{mn}(C_{mn}^{R,1} + C_{mn}^{OT,1}) \tag{4.33}$$

$$x_{mn} \geq \xi_{mn} C_{mn}^{R,1} \tag{4.34}$$

$$x_{mn} \leq (1 - \xi_{mn}) C_{mn}^{R,1} + \xi_{mn}(C_{mn}^{R,1} + C_{mn}^{OT,1}) \tag{4.35}$$

The introduced variables $\phi_m \in \{0, 1\}$ and $\xi_{mn} \in \{0, 1\}$ function as indicators whether or not $x_{mn} + \sum_{t \in \mathcal{T}(m)} d_t a_{tmn} \geq C_{mn}^{R,1}$ and $x_{mn} \geq C_{mn}^{R,1}$, respectively. Using these indicators, we can rewrite $f^{AS}((1, \mathbf{x}, \mathbf{y}), \mathbf{a})$ now as

$$\begin{aligned} f^{AS}((1, \mathbf{x}, \mathbf{y}), \mathbf{a}) &= \sum_{m \in \mathcal{M}} \sum_{n \in \mathcal{N}} h \left[x_{mn} + \sum_{t \in \mathcal{T}(m)} d_t a_{tmn} - C_{mn}^{R,1} \right]^+ - \sum_{m \in \mathcal{M}} \sum_{n \in \mathcal{N}} h \left[x_{mn} - C_{mn}^{R,1} \right]^+ \\ &= \sum_{m \in \mathcal{M}} \sum_{n \in \mathcal{N}} h \left[\left(x_{mn} + \sum_{t \in \mathcal{T}(m)} d_t a_{tmn} \right) \phi_{mn} - x_{mn} \xi_{mn} + (\xi_{mn} - \phi_{mn}) C_{mn}^{R,1} \right]. \end{aligned}$$

This is, however, still a non-linear function since it involves the multiplications of the integer variables $x_{mn} + \sum_{t \in \mathcal{T}(m)} d_t a_{tm1}$ and the binary variables ϕ_{mn} and x_{mn} and ξ_{mn}. Fortunately, such a multiplication of variables can be linearized by introducing additional integer variables θ_{mn} and κ_{mn}, for $m \in \mathcal{M}$ and $n \in \mathcal{N}$. We then can express the multiplication using the following conditions:

$$\theta_{mn} \leq \phi_{mn}(C_{mn}^{R,1} + C_{mn}^{OT,1}) \tag{4.36}$$

$$\theta_{mn} \leq x_{mn} + \sum_{t \in \mathcal{T}(m)} d_t a_{tmn}, \tag{4.37}$$

$$\theta_{mn} \geq x_{mn} + \sum_{t \in \mathcal{T}(m)} d_t a_{tmn} - (C_{mn}^{R,1} + C_{mn}^{OT,1})(1 - \phi_{mn}), \tag{4.38}$$

$$\kappa_{mn} \leq \xi_{mn}(C_{mn}^{R,1} + C_{mn}^{OT,1}), \tag{4.39}$$

$$\kappa_{mn} \leq x_{mn}, \tag{4.40}$$

$$\kappa_{mn} \geq x_{mn} - (C_{mn}^{R,1} + C_{mn}^{OT,1})(1 - \xi_{mn}), \tag{4.41}$$

$$\theta_{mn} \in \mathbb{N}_0, \tag{4.42}$$

$$\kappa_{mn} \in \mathbb{N}_0, \tag{4.43}$$

and write

$$f^{AS}(\mathbf{s}_1, \mathbf{a}) = \sum_{m \in \mathcal{M}} \sum_{n \in \mathcal{N}} h \left[\theta_{mn} - \kappa_{mn} + (\xi_{mn} - \phi_{mn}) C^{R,1}_{mn} \right]. \tag{4.44}$$

From the definition of ϕ_{mn} and ξ_{mn} it follows that if $\phi_{mn} = \xi_{mn} = 1$ we should incur $\sum_{m \in \mathcal{M}} \sum_{n \in \mathcal{N}} h \left[\sum_{t \in \mathcal{T}(m)} d_t a_{tmn} - x_{mn} \right]$ overtime costs. Since (4.37) and (4.38) imply that $\theta_{mn} = x_{mn} + \sum_{t \in \mathcal{T}(m)} d_t a_{tmn}$ and (4.40) and (4.41) imply that $\kappa_{mn} = x_{mn}$, it follows from (4.44) that

$$f^{AS}(\mathbf{s}_1, \mathbf{a}) = \sum_{m \in \mathcal{M}} \sum_{n \in \mathcal{N}} h \left[\sum_{t \in \mathcal{T}(m)} d_t a_{tmn} - x_{mn} \right].$$

Similarly, using other combinations of constraints (4.36)-(4.43), we can observe that (4.44) also gives the correct overtime costs for other possible combinations of ϕ_{mn} and ξ_{mn}.

All together, we get an ILP that is equivalent to the optimization problem in (4.30). In this ILP the constraints are given by the state space constraints, the action set constraints (4.3)-(4.7) and the constraints needed to linearize $f^{AS}(\mathbf{s}_1, \mathbf{a})$, (4.32)-(4.43). For the sake of completeness, we present this ILP in its completeness on the following page.
The optimization problem in (4.31) can be linearized in a similar way. We present the resulting ILP from this linearization in Appendix B.1.

Note that constraints (4.45c), (4.45g) and (4.45h) just fix action variables to be zero. Hence these variables do not contribute to the objective value of the ILP. So, to limit computation time, instead of implementing these constraints, we do not construct the decision variables that do appear in these constraints. Afterwards, to generate the state-action pair $(\mathbf{s}_1, \mathbf{a}_1)$ that corresponds to the ILP solution $(\mathbf{x}, \mathbf{y}, \mathbf{a})$, we fix the not-generated variables to be zero.

$$\min_{\substack{\mathbf{x},\mathbf{y},\mathbf{a},\\ \boldsymbol{\phi},\boldsymbol{\theta},\boldsymbol{\xi},\boldsymbol{\kappa}}} \quad \sum_{n=1}^{N}\sum_{t=1}^{T}\sum_{m\in\mathcal{M}(t)} f^{AT}(1,t,n)a_{tmn} + \sum_{n=1}^{N}\sum_{m=1}^{M} h\left[\theta_{mn} - \kappa_{mn} + (\xi_{mn} - \phi_{mn})C_{mn}^{R,1}\right]$$

$$- \sum_{m=1}^{M}\left\{\sum_{n=1}^{N} V_{mn}x_{mn} - \gamma\sum_{n=1}^{N-1}\widetilde{V}_{mn}(x_{m,n+1} + \sum_{t\in\mathcal{T}(m)} d_t a_{tm,n+1})\right\} - \sum_{t=1}^{T} W_t y_t, \quad (4.45a)$$

s.t.

$$x_{mn} \leq C_{mn}^{R,1} + C_{mn}^{OT,1}, \qquad \forall m\in\mathcal{M},\ \forall n = 1,2,\ldots,N-2, \quad (4.45b)$$

$$x_{mn} = 0, \qquad \forall m\in\mathcal{M},\ n = N-1, N, \quad (4.45c)$$

$$y_t \leq Q_t^1, \qquad \forall t\in\mathcal{T}, \quad (4.45d)$$

$$\sum_{n\in\mathcal{N}}\sum_{m\in\mathcal{M}(t)} a_{tmn} = y_t, \qquad \forall t\in\mathcal{T}, \quad (4.45e)$$

$$x_{mn} + \sum_{t\in\mathcal{T}(m)} d_t a_{tmn} \leq C_{mn}^{R,1} + C_{mn}^{OT,1}, \qquad \forall m\in\mathcal{M},\ n\in\mathcal{N}, \quad (4.45f)$$

$$a_{tmn} = 0, \qquad \forall t\in\mathcal{T},\ m\notin\mathcal{M}(t),\ n\in\mathcal{N}, \quad (4.45g)$$

$$a_{tm1} = a_{tm2} = 0, \qquad \forall t\in\mathcal{T}',\ m\in\mathcal{M}(t), \quad (4.45h)$$

$$x_{mn} + \sum_{t\in\mathcal{T}(m)} d_t a_{tmn} \geq \phi_{mn}C_{mn}^{R,1}, \qquad \forall m\in\mathcal{M},\ \forall n\in\mathcal{N} \quad (4.45i)$$

$$x_{mn} + \sum_{t\in\mathcal{T}(m)} d_t a_{tmn} \leq (1-\phi_{mn})C_{mn}^{R,1} + \phi_{mn}(C_{mn}^{R,1} + C_{mn}^{OT,1}), \quad \forall m\in\mathcal{M},\ n\in\mathcal{N}, \quad (4.45j)$$

$$x_{mn} \geq \xi_{mn}C_{mn}^{R,1}, \qquad \forall m\in\mathcal{M},\ n\in\mathcal{N}, \quad (4.45k)$$

$$x_{mn} \leq (1-\xi_{mn})C_{mn}^{R,1} + \xi_{mn}(C_{mn}^{R,1} + C_{mn}^{OT,1}), \qquad \forall m\in\mathcal{M},\ n\in\mathcal{N}, \quad (4.45l)$$

$$\theta_{mn} \leq \phi_{mn}(C_{m,1}^{R,1} + C_{m,1}^{OT,1}), \qquad \forall m\in\mathcal{M},\ n\in\mathcal{N}, \quad (4.45m)$$

$$\theta_{mn} \leq x_{mn} + \sum_{t\in\mathcal{T}(m)} d_t a_{tmn}, \qquad \forall m\in\mathcal{M},\ n\in\mathcal{N}, \quad (4.45n)$$

$$\theta_{mn} \geq x_{mn} + \sum_{t\in\mathcal{T}(m)} d_t a_{tmn} - (C_{mn}^{R,1} + C_{mn}^{OT,1})(1-\phi_{mn}), \qquad \forall m\in\mathcal{M},\ n\in\mathcal{N}, \quad (4.45o)$$

$$\kappa_{mn} \leq \xi_{mn}(C_{mn}^{R,1} + C_{mn}^{OT,1}), \qquad \forall m\in\mathcal{M},\ n\in\mathcal{N}, \quad (4.45p)$$

$$\kappa_{mn} \leq x_{mn}, \qquad \forall m\in\mathcal{M},\ n\in\mathcal{N}, \quad (4.45q)$$

$$\kappa_{mn} \geq x_{mn} - (C_{mn}^{R,1} + C_{mn}^{OT,1})(1-\xi_{mn}), \qquad \forall m\in\mathcal{M},\ n\in\mathcal{N}, \quad (4.45r)$$

$$\mathbf{x}\in\mathbb{N}_0^{MN},\ \mathbf{y}\in\mathbb{N}_0^{T},\ \mathbf{a}\in\mathbb{N}_+^{TMN}, \quad (4.45s)$$

$$\boldsymbol{\phi}\in\{0,1\}^{MN},\ \boldsymbol{\xi}\in\{0,1\}^{MN}, \boldsymbol{\theta}\in\mathbb{N}_0^{MN}, \boldsymbol{\kappa}\in\mathbb{N}_0^{MN}. \quad (4.45t)$$

At this point, we have split the pricing problem (4.29) into two ILPs. It is evident that g_1^* is the solution to the ILP (4.45a)-(4.45t) minus the constant U_1 and g_2^* is the solution to (B.1a)-(B.1t) minus the constant U_2. In the classical column generation we would now add the state-action pair associated with the most violated primal constraint. That is, if $g_1^* < g_2^*$ and $(\mathbf{s}_1^* = (1, \mathbf{x}^*, \mathbf{y}^*), \mathbf{a}^*)$ is the solution of (4.45a)-(4.45t), we would add the state-action pair $(\mathbf{s}_1^*, \mathbf{a}^*)$ to the RMP. If, instead, $g_2^* < g_1^*$, we would add $(\mathbf{s}_2^*, \mathbf{a}^*)$ to the RMP, where $(\mathbf{s}_2^* = (1, \mathbf{x}^*, \mathbf{y}^*), \mathbf{a}^*)$ is the solution of (B.1a)-(B.1t). However, the idea behind column generation is to add variables with negative reduced costs, since these variables may improve the solution. Thus, as long as $g_1^* < 0$, the solution $(\mathbf{s}_1^* = (1, \mathbf{x}^*, \mathbf{y}^*), \mathbf{a}^*)$ may improve the solution to the RMP and we should add this state-action pair. Similarly, if $g_2^* < 0$ the solution $(\mathbf{s}_2^* = (2, \mathbf{x}^*, \mathbf{y}^*), \mathbf{a}^*)$ may improve the solution to the RMP and we should add this state-action pair. In conclusion, there is no need to only add the state-action pair corresponding to $\min\{g_1^*, g_2^*\}$, but instead we could add both $(\mathbf{s}_1^* = (1, \mathbf{x}^*, \mathbf{y}^*), \mathbf{a}^*)$ and $(\mathbf{s}_2^* = (2, \mathbf{x}^*, \mathbf{y}^*), \mathbf{a}^*)$ simultaneously (in the same iteration) as long as both g_1^*, g_2^* are strictly less than zero. In fact, this is probably advantageous for the computation time of the column generation algorithm, since you need to solve both g_1^*, g_2^* in each iteration anyway, so by adding both $(\mathbf{s}_1^* = (1, \mathbf{x}^*, \mathbf{y}^*), \mathbf{a}^*)$ and $(\mathbf{s}_2^* = (2, \mathbf{x}^*, \mathbf{y}^*), \mathbf{a}^*)$ to the RMP the number of iterations is probably reduced. Algorithm 4.1 summarizes the column generation algorithm.

Algorithm 4.1: Column generation algorithm

Input : the MP(4.28a)-(4.28h) and
tolerance parameter δ

Output : Optimal solution to the MP (4.28a)-(4.28h)

Initialization: find initial $\mathcal{S}_1'$ $\mathcal{S}_2'$ (and set $\mathcal{S}' = \mathcal{S}_1' \cup \mathcal{S}_2'$), and $\mathcal{A}'(\mathbf{s}_1)$, for $\mathbf{s}_1 \in \mathcal{S}_1'$, and $\mathcal{A}'(\mathbf{s}_2)$, for $\mathbf{s}_2 \in \mathcal{S}_2'$, that satisfy the constraints (4.28b)-(4.28h) ;
while $\min\{g_1^*, g_2^*\} < \delta$ **do**
  Solve the RMP with $\mathcal{S}_1'$ $\mathcal{S}_2'$ and $\mathcal{A}'(\mathbf{s}_1)$, for $\mathbf{s}_1 \in \mathcal{S}_1'$, $\mathcal{A}'(\mathbf{s}_2)$, for $\mathbf{s}_2 \in \mathcal{S}_2'$;
  Solve the ILPs (4.45a)-(4.45t) and (B.1a)-(B.1t) and compute g_1^* and g_2^*;
  if $g_1^* < 0$ **then**
    Set $\mathcal{S}_1' := \mathcal{S}_1' \cup \{\mathbf{s}_1^*\}$ and $\mathcal{A}'(\mathbf{s}_1^*) := \mathcal{A}'(\mathbf{s}_1^*) \cup \{\mathbf{a}_1^*\}$, where $(\mathbf{s}_1^* = (1, \mathbf{x}^*, \mathbf{y}^*), \mathbf{a}_1^*)$ is the solution to the ILP (4.45a)-(4.45t) ;
  else if $g_2^* < 0$ **then**
    Set $\mathcal{S}_2' := \mathcal{S}_2' \cup \{\mathbf{s}_2^*\}$ and $\mathcal{A}'(\mathbf{s}_2^*) := \mathcal{A}'(\mathbf{s}_2^*) \cup \{\mathbf{a}_2^*\}$, where $(\mathbf{s}_2^* = (2, \mathbf{x}^*, \mathbf{y}^*), \mathbf{a}_2^*)$ is the solution to the ILP (B.1a)-(B.1t) ;
end

Finding initial state-action pairs

The last issue to tackle arises from the initialization of Algorithm 4.1: we need to find an initial set of state-action pairs. This set determines the RMP of the first iteration and needs to satisfy the constraints of the MP (4.28a)-(4.28h). Unfortunately, there is no easy way to identify an initial set of feasible state-action pairs. To find initial state-action pairs, the right-hand sides (RHSs) of the constraints (4.28b)-(4.28h) play an important role: the smaller these RHSs are, the easier it is to find feasible initial state-action pairs. At least, that is what we have experienced during experiments. In Appendix B.2 we present a convenient way to determine the RHSs. The presented manner is not only an easy way to determine them, but also offers the freedom to scale them to a certain extent.

Next, we first find one state-action pair per patient type for both decision epochs. That is, for all $t \in \mathcal{T}$, we find $\mathbf{s}_1(\mathbf{t}) \in \mathcal{S}_1$, $\mathbf{a}_1(\mathbf{t}) \in \mathcal{A}(\mathbf{s}_1(\mathbf{t}))$ and $\mathbf{s}_2(\mathbf{t}) \in \mathcal{S}_2$, $\mathbf{a}_2(\mathbf{t}) \in \mathcal{A}(\mathbf{s}_2(\mathbf{t}))$ from

$$\underset{\substack{\mathbf{s}_1 \in \mathcal{S}_1, \mathbf{a}_1 \in \mathcal{A}(\mathbf{s}_1) \\ \mathbf{s}_2 \in \mathcal{S}_2, \mathbf{a}_2 \in \mathcal{A}(\mathbf{s}_2)}}{\arg\max} \left\{ \min\left\{ q_1(\mathbf{s}_1, \mathbf{a}), q_2(\mathbf{s}_2, \mathbf{a}) \right\} \,\middle|\, y_t(\mathbf{s}_1) = Q_t^1, y_t(\mathbf{s}_2) = Q_t^2 \right\},$$

where

$$\begin{aligned} q_1(\mathbf{s}_1, \mathbf{a}_1) &= \sum_{m,n} x_{mn}(\mathbf{s}_1) - \tilde{\mu}_{mn}(\mathbf{s}_2, \mathbf{a}_2) \;\text{ and} \\ q_2(\mathbf{s}_2, \mathbf{a}_2) &= \sum_{m,n} x_{mn}(\mathbf{s}_2) - \mu_{mn}(\mathbf{s}_1, \mathbf{a}_1). \end{aligned}$$

Then we pass this initial state-action pairs, $\mathbf{s}_1(\mathbf{t})$, $\mathbf{a}_1(\mathbf{t})$, $\mathbf{s}_2(\mathbf{t})$, $\mathbf{a}_2(\mathbf{t})$ to the following LP, which we solve iteratively until $\psi = 0$. In each iteration we add state-action pairs based on the pricing problems (4.45a)-(4.45t) and (B.1a)-(B.1t), that are also used to solve the RMP.

$$\min_{\mathbf{X}, \psi} \quad \psi,$$

$$\begin{aligned}
\text{s.t.} \quad & \sum_{\substack{\mathbf{s}_1 \in \mathcal{S}_1, \\ \mathbf{a}_1 \in \mathcal{A}(\mathbf{s}_1)}} X(\mathbf{s}_1, \mathbf{a}_1) \;-\; \gamma \sum_{\substack{\mathbf{s}_2 \in \mathcal{S}_2, \\ \mathbf{a}_2 \in \mathcal{A}(\mathbf{s}_2)}} X(\mathbf{s}_2, \mathbf{a}_2) = 1 \\
& \sum_{\substack{\mathbf{s}_2 \in \mathcal{S}_2, \\ \mathbf{a}_2 \in \mathcal{A}(\mathbf{s}_2)}} X(\mathbf{s}_2, \mathbf{a}_2) \;-\; \gamma \sum_{\substack{\mathbf{s}_1 \in \mathcal{S}_1, \\ \mathbf{a}_1 \in \mathcal{A}(\mathbf{s}_1)}} X(\mathbf{s}_1, \mathbf{a}_1) = 1 \\
& \sum_{\substack{\mathbf{s}_1 \in \mathcal{S}_1, \\ \mathbf{a}_1 \in \mathcal{A}(\mathbf{s}_1)}} x_{mn}(\mathbf{s}_1) X(\mathbf{s}_1, \mathbf{a}_1) \;-\; \sum_{\substack{\mathbf{s}_2 \in \mathcal{S}_2, \\ \mathbf{a}_2 \in \mathcal{A}(\mathbf{s}_2)}} \tilde{\mu}_{mn}(\mathbf{s}_2, \mathbf{a}_2) X(\mathbf{s}_2, \mathbf{a}_2) \geq \mathbb{E}_\alpha[x_{mn}] - \psi, \quad && \forall m \in \mathcal{M}, \forall n \in \mathcal{N} \\
& \sum_{\substack{\mathbf{s}_2 \in \mathcal{S}_2, \\ \mathbf{a}_2 \in \mathcal{A}(\mathbf{s}_2)}} x_{mn}(\mathbf{s}_2) X(\mathbf{s}_2, \mathbf{a}_2) \;-\; \sum_{\substack{\mathbf{s}_1 \in \mathcal{S}_1, \\ \mathbf{a}_1 \in \mathcal{A}(\mathbf{s}_1)}} \mu_{mn}(\mathbf{s}_1, \mathbf{a}_1) X(\mathbf{s}_1, \mathbf{a}_1) \geq \mathbb{E}_\beta[x_{mn}] - \psi, \quad && \forall m \in \mathcal{M}, \forall n \in \mathcal{N} \\
& \sum_{\substack{\mathbf{s}_1 \in \mathcal{S}_1, \\ \mathbf{a}_1 \in \mathcal{A}(\mathbf{s}_1)}} y_t(\mathbf{s}_1) X(\mathbf{s}_1, \mathbf{a}_1) \;-\; \tilde{\lambda}_t \sum_{\substack{\mathbf{s}_2 \in \mathcal{S}_2, \\ \mathbf{a}_2 \in \mathcal{A}(\mathbf{s}_2)}} X(\mathbf{s}_2, \mathbf{a}_2) \geq \mathbb{E}_\alpha[y_t] - \psi, \quad && \forall t \in \mathcal{T} \\
& \sum_{\substack{\mathbf{s}_2 \in \mathcal{S}_2, \\ \mathbf{a}_2 \in \mathcal{A}(\mathbf{s}_2)}} y_t(\mathbf{s}_2) X(\mathbf{s}_2, \mathbf{a}_2) \;-\; \lambda_t \sum_{\substack{\mathbf{s}_1 \in \mathcal{S}_1, \\ \mathbf{a}_1 \in \mathcal{A}(\mathbf{s}_1)}} X(\mathbf{s}_1, \mathbf{a}_1) \geq \mathbb{E}_\beta[y_t] - \psi, \quad && \forall t \in \mathcal{T} \\
& X(\mathbf{s}, \mathbf{a}) \in \mathbb{R}_+, \text{ for all } \mathbf{s} \in \mathcal{S}, \mathbf{a} \in \mathcal{A}(\mathbf{s}), \quad \psi \geq 0
\end{aligned}$$

4.2.3 Approximate optimal policy

In this section, we discuss how to derive a policy from the solution to the(dual) ALP. We refer to this policy as the approximate optimal policy (AOP). In traditional LP method for directly solving the MDP, we would have solved the dual of (4.19a)-(4.19c), yielding an optimal value $X^*(\mathbf{s}, \mathbf{a})$, for all states $\mathbf{s} \in \mathcal{S}$ and for all actions $\mathbf{a} \in \mathcal{A}(\mathbf{s})$. From this we would be able to construct an optimal policy by setting the probability of using action $\mathbf{a}$ in state $\mathbf{s}$ equal to the value of the dual variable, $X^*(\mathbf{s}, \mathbf{a})$, divided by the sum of the dual variables over all possible actions in state $\mathbf{s}$. The viability of this method depends on the fact that a direct solution to the LP (4.19a)-(4.19c) will have at least one positive variable, $X^*(\mathbf{s}, \mathbf{a})$, for all states $\mathbf{s}$. 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 the optimal values V_0^*, $\widetilde{V}_0^*$, $\mathbf{V}^*$, $\widetilde{\mathbf{V}}^*$ $\mathbf{W}^*$, $\widetilde{\mathbf{W}}^*$ into the right-hand side of the optimality equations (4.18) and then solve them for $\mathbf{a} \in \mathcal{A}(\mathbf{s})$. For states $\mathbf{s}_1 = (1, \mathbf{x}, \mathbf{y}) \in \mathcal{S}_1$ this gives the following optimization problem:

$$\begin{aligned}
&\min_{\mathbf{a}\in\mathcal{A}(\mathbf{s}_1)} \left\{ c(\mathbf{s}_1, \mathbf{a}) + \gamma \sum_{\mathbf{s}'\in\mathcal{S}} p(\mathbf{s}'|\mathbf{s}, \mathbf{a}) v(\mathbf{s}') \right\}, \\
&= \min_{\mathbf{a}\in\mathcal{A}(\mathbf{s}_1)} \left\{ \begin{array}{l} \displaystyle\sum_{n=1}^{N}\sum_{t=1}^{T}\sum_{m\in\mathcal{M}(t)} f^{AT}(t,n)a_{tmn} + f^{AS}(\mathbf{s}_1,\mathbf{a}) \\ \displaystyle + \gamma \sum_{\mathbf{y}'\in\mathcal{Y}_1} p_1(\mathbf{y}') \left(\widetilde{V}_0^* + \sum_{m=1}^{M}\sum_{n=1}^{N-1} \widetilde{V}_{mn}^*(x_{m,n+1} + \sum_{t\in\mathcal{T}(m)} d_t a_{tm,n+1}) + \sum_{t=1}^{T} \widetilde{W}_t^* y'_t, \right) \end{array} \right\}, \\
&= \min_{\mathbf{a}\in\mathcal{A}(\mathbf{s}_1)} \left\{ \sum_{n=1}^{N}\sum_{t=1}^{T}\sum_{m\in\mathcal{M}(t)} f^{AT}(t,n)a_{tmn} + f^{AS}(\mathbf{s}_1,\mathbf{a}) + \gamma \sum_{m=1}^{M}\sum_{n=1}^{N-1} \widetilde{V}_{mn}^* \left(\sum_{t\in\mathcal{T}(m)} d_t a_{tm,n+1} \right) \right\} + Z_1,
\end{aligned}$$

where the constant

$$Z_1 = \gamma \left(\widetilde{V}_0^* + \sum_{m=1}^{M}\sum_{n=1}^{N-1} \widetilde{V}_{mn}^* x_{m,n+1} \right) + \sum_{t=1}^{T} \widetilde{W}_t^* \lambda_t.$$

Note that to determine the optimal action in state $\mathbf{s}_1$, the constant Z_1 is of no importance. To reformulate this optimization program into an equivalent ILP, we use the same linearization of f^{AS} as in (4.45a)-(4.45t):

$$\begin{aligned}
\min_{\substack{\mathbf{a},\boldsymbol{\phi},\boldsymbol{\theta}\\ \boldsymbol{\xi},\boldsymbol{\kappa}}} \quad & \sum_{n=1}^{N}\sum_{t=1}^{T}\sum_{m\in\mathcal{M}(t)} f^{AT}(1,t,n)a_{tmn} + \sum_{m=1}^{M}\sum_{n=1}^{N} h \left[\theta_{mn}\kappa_{mn} + (\xi_{mn} - \phi_{mn}) C_{mn}^{R,1} \right] \\
& + \gamma \sum_{m=1}^{M}\sum_{n=1}^{N-1} \widetilde{V}_{mn}^* \left(\sum_{t\in\mathcal{T}(m)} d_t a_{tm,n+1} \right), \\
\text{s.t.} \quad & \text{constraints (4.45b)-(4.45t).}
\end{aligned}$$

Note that, in constrast to (4.45a)-(4.45t), here the x_{mn}'s are no decision variables but an input parameter of the ILP.

Analogously, to determine the approximate optimal action $\mathbf{a} \in \mathcal{A}(\mathbf{s}_2)$, for states $\mathbf{s}_2 = (2, \mathbf{x}, \mathbf{y})$, we get

$$\min_{\mathbf{a}\in\mathcal{A}(\mathbf{s}_2)} \left\{ \sum_{n=1}^{N} \sum_{t=1}^{T} \sum_{m\in\mathcal{M}(t)} f^{AT}(t,n) a_{tmn} + f^{AS}(\mathbf{s}_2, \mathbf{a}) + \gamma \sum_{m=1}^{M} \sum_{n=1}^{N-2} V_{mn}^* \left(\sum_{t\in\mathcal{T}(m)} d_t a_{tm,n+2} \right) \right\} + Z_2,$$

where the constant

$$Z_2 = \gamma \left(V_0^* + \sum_{m=1}^{M} \sum_{n=1}^{N-2} V_{mn}^* x_{m,n+2} \right) + \sum_{t=1}^{T} W_t^* \tilde{\lambda}_t.$$

Or, equivalently,

$$\begin{aligned} \min_{\substack{\mathbf{a},\boldsymbol{\phi},\boldsymbol{\theta}\\ \boldsymbol{\xi},\boldsymbol{\kappa}}} \quad & \sum_{n=1}^{N} \sum_{t=1}^{T} \sum_{m\in\mathcal{M}(t)} f^{AT}(2,t,n) a_{tmn} + \sum_{m=1}^{M} \sum_{n=1}^{N} h \left[\theta_{mn}\kappa_{mn} + (\xi_{mn} - \phi_{mn}) C_{mn}^{R,2} \right] \\ & + \gamma \sum_{m=1}^{M} \sum_{n=1}^{N-2} V_{mn}^* \left(\sum_{t\in\mathcal{T}(m)} d_t a_{tm,n+2} \right), \\ \text{s.t.} \quad & \text{constraints (B.1b)-(B.1t).} \end{aligned}$$

In practice, rather than computing and storing the approximate optimal actions for each state, a resource-intensive task, we only compute them as needed.

The implementation of the column generation algorithm was performed in AIMMS 4.39 with CPLEX 12.7 as the solver.

5

Data study

This chapter reports on the results of a data analysis we performed on a data set provided by the Radiology department of Rijnstate. The purpose of this data analysis is twofold. Firstly, it provides a quantitative (approximate) summary of the performances of the department during the period that is included in the data set. Secondly, to evaluate the performance of the Markov decision process (MDP)-based MRI appointment system (AS) of Chapter 4, we need to derive statistical estimators for the required parameters of this model. Section 5.1 is about the analysis of the current performances and Section 5.2 about deriving the input for the MDP from the previous chapter. In this latter section, we assume the reader to be familiar with the MDP and its notation. The results in this chapter have been calculated using the statistical software program IBM SPSS Statistics 23, because this was available at Rijnstate's computers. For easy reference, in the remaining chapters of this thesis, we will refer to the 1.5 Tesla MRI scanner at location Arnhem as MRI scanner 1, the 3 Tesla MRI scanner at location Arnhem as MRI scanner 2 and the 1.5 Tesla MRI scanner at location Zevenaar as MRI scanner 3.

5.1 Data analysis of current performances of Rijnstate's Radiology department

The analyses in this section are based on a data set provided by the Radiology department of Rijnstate. This data set initially consists of data on all MRI scans performed on MRI scanner 2 and MRI scanner 3 from January 2, 2019 up to and including August 30, 2019. This data comes from the department's computer system HIX and initially consisted of 7135 data points. However, we had to remove 75; 6 points because these are maintenance visits to the MRI scanners (3 maintenance visits to each MRI scanner) and another 69 data points because they are incorrect or incomplete. By incomplete we mean that these data points could not be used to score any of the key performance indicators (KPIs). By incorrect data points we mean any irregularities such as negative service times. On some days, according to the data, three MRI scans take place simultaneously. Since there are only two MRI scanners in the data set, this is impossible. In these cases there is always one MRI scan that is entirely within the service period of one of the other two. We do not know how these irregularities ended up in the data, but we have removed these data points before doing any further analyses. At the end, we have a data set of 7060 data points.

Based on the data set we evaluate the performance of Rijnstate's Radiology department on the following key performance indicators (KPIs): (direct) waiting time, access time, average daily utilization and average daily overtime. Utilization is defined as the percentage of available regular appointment slots used for delivering MRI examinations. Note that utilization values different than 100% are associated with either overtime or idle time and thus utilization values close to 100% are preferred. We calculate the utilization and overtime on a daily basis, because this offers us the flexibility to exclude days in the calculation if we observe any abnormalities in the data of this day. Table 5.1 shows the scores on these KPIs. Apart from the response time and the (direct) waiting time, all KPIs are blueprint dependent. During the time span of the data, the department altered the blueprint calendars twice (and occasionally for a single day). We observed the scores on the KPIs for all three periods and found that the difference was very small. Hence we only report on the overall scores on the KPIs. In the simulation model in the next period, we used the blueprint calendars used by the department during the period from January 2, 2019 up to and including April 26, 2019. The absolute difference between the score on a KPI during this specific period and overall was at most 1. Besides the scores on the KPIs, Table 5.1 also shows the number of data points these scores are based on. These numbers are included because not all data points could be used to evaluate the score on a KPI. We will explain which data points needed to be excluded for each KPI later on. Note the available data set does not allow us to perform analyzes per individual MRI scanner.

Three important performance indicators for which the data set does not offer any information are the no-show percentage and the cancellation percentage. We have no information about the no-show percentage because we only have data available from MRI scans that actually took place.
If a cancellation occurs where no alternative appointment is planned (immediately), the same consequence as with a no-show applies and this appointment did not appear in our data set. If today a cancellation occurs where an alternative appointment is planned immediately, the HIX system works in such a way that the order date (the date on which the MRI request was first submitted) is not adjusted. However, the original booking date (the date on which the appointment was booked that is rescheduled today) is adjusted to today's date. If we include these MRI appointments in the data set in this way (original order date and adjusted booking date), this has negative influence on the response time and almost always negative influence on the access time (almost always the alternative appointment is planned on a day further in the future than the original appointment day). In addition, the access time for the alternative appointment is determined by what the agenda looks like on the day the cancellation (of the original appointment) takes place. Hence, because it gives a more realistic picture of the response time and access time, we asked Rijnstate's data specialist to adjust the original order date to the booking date of the alternative appointment if a cancellation has taken place. Due to the complexity of the HIX system, it was impossible for us to do this ourselves. The price we pay for this is that we do not gain insight into the cancellation percentage based on the given data set. A final comment on this data adjustment is that it may incorrectly reduce the mean response time slightly.

Table 5.1: Current performances (mean and standard deviation) of the Radiology department of Rijnstate. The mean and standard deviation calculations were made based on the number of data points from the data set from as displayed in the last column of the table. Access times are calculated in workdays instead of calendar days.

KPI	mean	std. dev.	n
(Direct) waiting time	4.95 min.	23.40 min.	6971
Access time	22.9 days	18.5 days	7060
Daily utilization	83.10%	9.98%	173 days
Daily overtime	8.57 min.	21.79 min.	173 days

Theoretically it can happen that an alternative MRI appointment is not immediately booked in the event of a cancellation, but this happens the following day. However, the data adjustment ensures that the response time when booking an alternative appointment is always equal to zero days.
The service levels for patients, computed for each priority class as the percentage of patients booked within the corresponding access time target, can also not be determined from the data set.

In the calculations for the average response time and its standard deviation, we have omitted data points concerning emergency patients. This type patients randomly arrive throughout the day (also after regular working hours and on weekend days) and for their MRI examinations no appointment has been booked in advance. Hence, these data points would reduce the response time in an unfair manner.
For the calculation of the mean (direct) waiting time we excluded the same MRI examinations for which no appointment has been booked in advance and beside 20 more data points which we considered as outliers. Since the patients for which no appointment has been booked in advance can be considered as emergency patients, these patients are *squeezed into* the regular program (or arrive after the regular program) and therefore have a different perception of waiting time than the other patients. The outliers passed the test(s) to be irregularities but have a direct waiting time longer than four hours. This seemed to be so unlikely to happen that we excluded them.
For the access time in Table 5.1 we included all patients. In Table 5.2 we recalculated the mean access time without the emergency patients. Of course, this number is slightly higher since we excluded data points concerning an access time of zero days. In Table 5.2 we also calculate the mean access time without emergency patients and (half-) yearly check-ups.

In Table 5.2 we also calculated the number of MRI scans done per day, taking into account different types of days. The time span of the data set is 240 days. On 195 of these 240 days at least one MRI scan is done, weekend days and holidays included. According to the blueprint calendars used during the data period, there were 181 workdays during the time span of the data set. This means that there were 14 days on which an extra shift was inserted, for example an extra Sunday shift, or on which at least one MRI scan was performed without it being an (extra inserted) working day. If we only consider weekdays, we end up with 173 days and on average 39.8 MRI scans per day. That is on average 19.6 scans per MRI scanner per day. We used only

Table 5.2: Other quantities calculated from the data set. The mean and standard deviation calculations were made based on the number of data points as displayed in the last column of the table.

Quantity	mean	Stand. dev.	n
Number of MRI scans/examinations done per day; all days	36.2	13.8	195 days
Number of MRI scans/examinations done per day; workdays according to blueprint(s)	38.7	10.8	181 days
Number of MRI scans/examinations done per day; only weekdays	39.8	10.4	173 days
Access time without emergency data points	23.2 days	17.8 days	6991
Access time without emergency & (half-)yearly check-up data points	16.1	11.0	6729

weekdays in the calculation for the utilization and overtime, because the other days deviated too much.

5.2 Obtaining model input parameters from the data

For the analyses in this section we used historical data from January 2, 2019 to May 1, 2019. For this period we know how many MRI appointment requests arrived every day, the time of arrival and the date time when the appointment is booked. We only look at weekdays because MRI appointments are booked on weekdays. This implies that the MRI requests that are submitted on Friday evening, Saturday or Sunday and do not have the priority to be scanned during the weekend occur in our data on the following Monday and that we have 108 days in our data set. The average response time in this data set is 0.19 days ($\approx$ 4.5 hrs.), with an standard deviation of 0.46 days ($\approx$ 11 hrs.).

In the next chapter, we evaluate the performance of the approximate optimal policy (AOP) through simulation. One goal of this simulation is that we would like to make statements about the access times obtained by the AOP if we would implement this policy in practice in Rijnstate. Hence, we would like a valid simulation model. To perform validation exercise(s), we need to code the current scheduling practice and need the current blueprint calendars. In Section 5.2.1 we determine the input for the validation of the simulation model. This model uses a subset of the currently defined patient types (which are dictated by the current blueprint calendars). However, some of these patient types are identical for the MDP (they have the same access time target, capacity requirement and patient type-MRI scanner compatibility restrictions) and hence we merge into one single patient type to keep the problem size small as possible. In Section 5.2.2 we present the MDP patient types and their attributes, based on the data set.

5.2.1 Input parameters for validation of the simulation model

Since we only know the aggregated current performance of MRI scanner 2 and MRI scanner 3, the validation of any simulation model must be done based on the results presented in Section 5.1 (if we want to do validation exercise(s) based on data and not, for instance, on expert opinion). Hence, we need a simulation model for the fictitious Rijnstate situation with two MRI scanners, patient types and their arrival distribution for this system and the blueprint calendars for the two MRI scanners.

If we consider the current real Rijnstate situation with all the three MRI scanners and their blueprint calendars (the blueprint calendars used from January 2, 2019 up to and including April 26, 2019), these blueprint calendars dictates patient types in the way we illustrated in Figure2.1 in Chapter 2. Based on the blueprint calendars, we can distinguish 51 different patient types, which are all listed in Table C.1 in Appendix C.1. For the validation of our simulation model, we take the subset of all patient types that can at least feasibly be served by either MRI scanner 2 or MRI scanner 3. The characteristics of these patient types are shown in Table 5.3. Here patient types 1, 3-5 have an access time target of 0 days (that is they need to be examined in the remaining part of today; access time target of 2 sessions observed from the first decision epoch of the day and an access time target of 1 session observed from the second decision epoch of the day); patient types 6-14 have an access time target of 3 (future) days (access time target of 11 sessions observed from the first decision epoch of the day and an access time target of 10 session observed from the second decision epoch of the day); patient types 15-21, 24-29 haven an access time target of 5 (future) days (access time target of 17 sessions observed from the first decision epoch of the day and an access time target of 16 session observed from the second decision epoch of the day); and finally patient types 30-39, 43-51 have an access time target of 10 (future) days (access time target of 32 sessions observed from the first decision epoch of the day and an access time target of 31 session observed from the second decision epoch of the day). For simplicity we incorporated (half-)yearly check-up appointment requests into one of the patient types 30-39, 43-51. All patient types with an access time target greater than 0 days are assumed to be outpatient and cannot be booked into the remaining part of today, i.e. all patient types other than types 1, 3-5 cannot be booked into session 1 and 2 if are at the first decision epoch of the day and not into session 1 if we are at the second decision epoch of the day.

Table 5.3: Characteristics of patient types in the simulation model validation instance. Here we have only MRI scanner 2 and MRI scanner 3. The duration describes the capacity requirement of each patient in number of 10-minute appointment slots. The third and fourth column show whether a MRI scanner is suitable for the type of MRI examination (✓) or not (✗).

Patient type	**Duration (# slots)**	**Compatibility MRI scanner ...**		**Occurrence in data at DE ... (in %)**		**Arrival rate (# reqs./day) at DE ...**	
		2	**3**	**1**	**2**	**1**	**2**
1	3	✓	✗	1.047	0.552	0.231	0.104
3	3	✓	✗	1.710	1.161	0.377	0.218

(Continued on next page)

Table 5.3 – *Continued*

Patient type	Duration (# slots)	Compatibility MRI scanner ...		Occurrence in data at DE ... (in %)		Arrival rate (# reqs./day) at DE ...	
		2	3	1	2	1	2
4	6	✓	✗	0.281	0.244	0.062	0.046
5	3	✗	✓	1.680	0.464	0.370	0.087
6	2	✓	✓	3.784	2.931	0.835	0.552
7	3	✓	✓	4.722	3.728	1.041	0.701
8	3	✓	✓	0.847	1.054	0.187	0.198
9	3	✓	✗	0.961	1.017	0.212	0.191
10	3	✓	✓	4.545	4.235	1.002	0.797
11	3	✓	✓	4.914	3.357	1.084	0.632
12	3	✓	✓	2.276	2.621	0.502	0.493
13	3	✓	✗	1.967	1.234	0.434	0.232
14	2	✗	✓	0.836	1.361	0.184	0.256
15	3	✓	✓	3.818	5.346	0.842	1.006
16	3	✓	✓	3.988	3.718	0.880	0.700
17	2	✓	✓	0.405	0.594	0.089	0.112
18	3	✓	✗	0.899	1.516	0.198	0.285
19	2	✓	✗	0.416	0.262	0.092	0.049
20	2	✓	✗	0.199	0.428	0.044	0.080
21	2	✗	✓	1.315	1.501	0.290	0.283
24	3	✓	✓	4.660	5.411	1.028	1.018
25	4	✓	✗	1.353	1.197	0.298	0.225
26	3	✗	✓	1.422	1.964	0.314	0.370
27	2	✗	✓	0.516	0.221	0.114	0.042
28	2	✗	✓	0.984	0.584	0.217	0.110
29	2	✗	✓	1.241	0.548	0.274	0.103
30	2	✓	✓	2.398	1.776	0.529	0.334
31	2	✓	✓	4.457	5.110	0.983	0.962
32	3	✓	✓	4.010	5.376	0.884	1.012
33	3	✓	✓	4.319	4.942	0.953	0.930
34	3	✓	✓	2.405	2.046	0.530	0.385
35	3	✓	✓	1.056	1.298	0.233	0.244
36	3	✓	✓	1.229	1.720	0.271	0.324
37	3	✓	✗	1.252	0.689	0.276	0.130
38	3	✓	✗	1.835	1.517	0.405	0.285
39	4	✓	✗	1.273	1.561	0.281	0.294
43	3	✓	✓	3.660	5.843	0.807	1.100

(Continued on next page)

Table 5.3 – *Continued*

Patient type	Duration (# slots)	Compatibility MRI scanner ...		Occurrence in data at DE ... (in %)		Arrival rate (# reqs./day) at DE ...	
		2	3	1	2	1	2
44	3	✓	✓	5.351	4.185	1.180	0.788
45	3	✓	✓	3.719	2.785	0.820	0.524
46	3	✓	✓	3.357	6.445	0.740	1.213
47	6	✓	✗	2.494	2.577	0.550	0.485
48	3	✓	✗	1.831	1.068	0.404	0.201
49	3	✓	✗	1.781	0.658	0.393	0.124
50	3	✓	✗	1.107	1.169	0.244	0.220
51	3	✓	✗	1.681	1.985	0.371	0.373
				100.000	**100.000**	**22.056**	**18.817**

The MDP has two decision epochs every day. In the simulation model we present in the next chapter, we also use these decision epochs. The first decision epoch on a day will be at 12.30 p.m. and the second decision epoch of the day will be at 5 p.m. At 5p.m. the other outpatient clinics of the hospital are closed and so all same-day priority appointment requests are known. However, the MRI scanners still operate for at least an hour and so we can try to book same-day priority patients in this last hour. At 12.30 p.m., the morning session of outpatient clinics is finished, so the same-day priority MRI requests emerging from these sessions is known. Also, the morning rounds on the wards are completed, so the majority of inpatient same-day priority MRI requests is known.

For all patient types in Table 5.3 we want the arrival rate for requests at both decision epochs. However, we only want the arrival rates for the appointment requests submitted to MRI scanner 2 or MRI scanner 3. Fortunately, for each data point in the data set we know on which MRI scanner it was eventually booked. For instance, patient type 1 can be booked on both MRI scanner 1 and MRI scanner 2, but in the simulation model of the fictitious situation with only MRI scanner 2 and MRI scanner 3, we only want those requests from patient type 1 that eventually have been booked on MRI scanner 2. In the data set we observe that appointment requests arrive at an average rate of 22.06 requests at the first decision epoch of a day, with a standard deviation of 5.45. At the second decision epoch of a day, we observe requests to arrive at an average rate of 18.82, with a standard deviation of 4.57. For simulation, it is necessary to model the appointment request arrivals at both decision epochs by a probability distribution. To this end, we first tried to fit Poisson distributions for all types separately. However, in none of the cases we could find a satisfying fit. Instead we model the total number of appointment requests at both decision epochs by a probability distribution and each arriving patient in the simulation is randomly assigned to be one of the 45 patient types, based on the empirically defined probability distribution in the sixth and seventh column of Table 5.3.

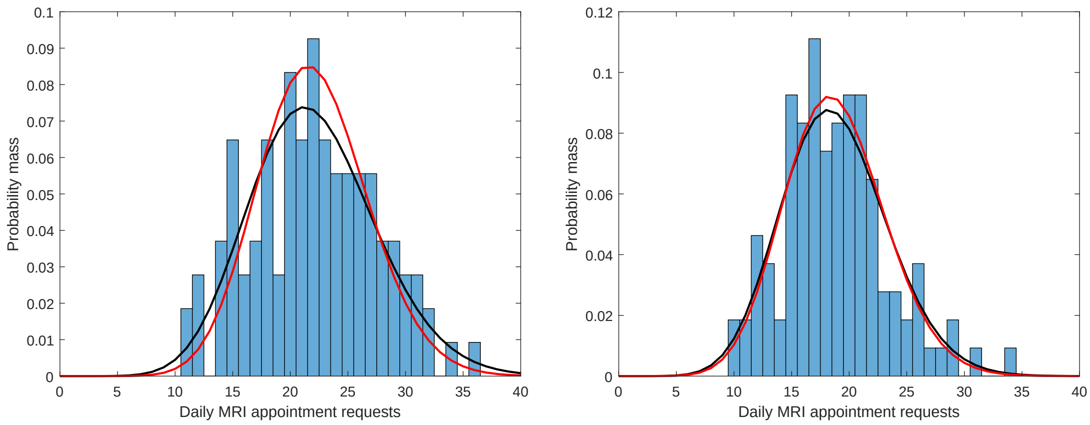

(a) Daily MRI appointment requests submitted to MRI scanner 2 or MRI scanner 3 at the first decision epoch plotted as a histogram along with the Poisson (red) and negative binomial (black) distribution curve fitted to this data.

(b) Daily MRI appointment requests submitted to MRI scanner 2 or MRI scanner 3 at the second decision epoch plotted as a histogram along with the Poisson (red) and negative binomial (black) distribution curve fitted to this data.

Figure 5.1: MRI appointment requests submitted to MRI scanner 2 or MRI scanner 3 at the first (a) and second (b) decision epoch of a day plotted as a histogram along with the Poisson (red) and negative binomial (black) distribution curve fitted to this data.

In Figure 5.1a the number of MRI appointment requests observed at the first decision epoch is plotted as a histogram along with the Poisson and negative binomial distribution curve fitted to it. In case of the negative binomial distribution, the random variable models the number of 'failures' before a specified number of successes, r^{th}, is reached in a series of independent, identical trials. The parameter p is the probability of success in a single trial. Both fitted distributions are found using the Maximum Likelihood Estimator (MLE) method. From this it follows immediately that the Poisson distribution fitted to the data has a mean $\lambda = 22.06$. The parameters for the negative binomial fitted to the data are $r = 64$, $p = 0.744$. The fit of both distributions with respect to the data is statistically tested using the Chi-squared goodness of fit test. For the Poisson distribution we found p-value $= 0.232$. For the negative binomial distribution we found p-value $= 0.943$. Hence we model the total number of appointment requests at the first decision epoch as a negative binomial distribution with parameters $r = 64$, $p = 0.744$. Figure 5.1b shows the equivalent of Figure 5.1a for the second decision epoch. The Poisson distribution fitted to the data has a mean $\lambda = 18.82$. The parameters for the negative binomial fitted to the data are $r = 173$, $p = 0.902$. For the Poisson distribution we found p-value $= 0.300$. For the negative binomial distribution we found p-value $= 0.535$. Hence we model the total number of appointment requests at the second decision epoch as a negative binomial distribution with parameters $r = 173$, $p = 0.902$.

5.2.2 Patient types and the arrival distribution of appointment requests for the MDP

In the next chapter, we eventually want to simulate various policies among which the AOP for the complete system with all three MRI scanners. Hence for this system, we also want to model the appointment request arrivals for all patient types at both decision epochs by a probability distribution. However, some of the patient types in Table C.1 are identical to the MDP and we will explain why. The patient types in this table are dictated by the blueprint calendars, i.e., every HIX code that has some appointment slots uniquely reserved for it, yields as a patient type. The MDP cannot distinguish different days in terms of total available capacity and it can also not reserve appointment slots for specific HIX codes. Hence, all patient types from Table C.1 that have the same access time target, plus the same capacity requirement, plus the same MRI scanner compatibility restrictions are identical to the MDP. Table C.2 shows the MDP patient types and how they follow from the patient types dictated by the blueprint calendars in Table C.1.

Table 5.5 shows the characteristics of the MDP patient types. Here patient types 1-5 have an access time target of 0 days (that is they need to be examined in the remaining part of today; access time target of 2 sessions observed from the first decision epoch of the day and an access time target of 1 session observed from the second decision epoch of the day); patient types 6-11 have an access time target of 3 (future) days (access time target of 11 sessions observed from the first decision epoch of the day and an access time target of 10 session observed from the second decision epoch of the day); patient types 12-19 haven an access time target of 5 (future) days (access time target of 17 sessions observed from the first decision epoch of the day and an access time target of 16 session observed from the second decision epoch of the day); and finally patient types 20-27 have an access time target of 10 (future) days (access time target of 32 sessions observed from the first decision epoch of the day and an access time target of 31 session observed from the second decision epoch of the day). For simplicity we incorporated (half-)yearly check-up appointment requests into one of the patient types 20-27. All patient types with an access time target greater than 0 days are assumed to be outpatient and cannot be booked into the remaining part of today, i.e. all patient types other than types 1-5 cannot be booked into session 1 and 2 if are at the first decision epoch of the day and not into session 1 if we are at the second decision epoch of the day.

Table 5.5: Characteristics of the MDP patient types. The duration describes the capacity requirement of each patient in number of 10-minute appointment slots. The third, fourth and fifth column show whether a MRI scanner is suitable for the type of MRI examination (✓) or not (✗).

Patient type	Duration (# slots)	Compatibility MRI scanner ...			Occurrence in data at DE ... (in %)		Arrival rate (# reqs./day) at DE ...	
		1	**2**	**3**	**1**	**2**	**1**	**2**
1	3	✓	✓	✗	1.612	0.109	0.495	0.030
2	3	✓	✗	✗	1.284	0.414	0.394	0.115
3	3	✗	✓	✗	1.462	0.520	0.449	0.144
4	6	✗	✓	✗	0.240	0.211	0.074	0.059
5	3	✗	✗	✓	1.436	0.401	0.441	0.111

(Continued on next page)

Table 5.5 – *Continued*

Patient type	Duration (# slots)	Compatibility MRI scanner ...			Occurrence in data at DE ... (in %)		Arrival rate (# reqs./day) at DE ...	
		1	**2**	**3**	**1**	**2**	**1**	**2**
6	2	✓	✓	✓	3.236	2.531	0.994	0.702
7	3	✓	✓	✓	4.763	4.129	1.463	1.145
8	3	✓	✓	✓	1.929	1.398	0.592	0.388
9	3	✓	✓	✓	10.036	8.818	3.082	2.445
10	3	✗	✓	✗	1.682	1.066	0.516	0.295
11	2	✗	✗	✓	0.715	1.175	0.220	0.326
12	3	✓	✓	✓	6.676	7.826	2.050	2.170
13	2	✓	✓	✗	2.475	3.037	0.760	0.842
14	3	✓	✓	✗	1.273	1.427	0.391	0.396
15	3	✓	✗	✗	2.983	2.963	0.916	0.822
16	3	✗	✓	✓	3.986	4.672	1.224	1.295
17	4	✗	✓	✗	1.157	1.033	0.355	0.287
18	3	✗	✗	✓	1.216	1.696	0.373	0.470
19	2	✗	✗	✓	2.344	1.168	0.720	0.324
20	2	✓	✓	✓	5.862	5.945	1.800	1.648
21	3	✓	✓	✓	11.134	15.281	3.419	4.237
22	3	✓	✓	✗	4.716	4.098	1.448	1.136
23	4	✓	✓	✓	2.152	2.239	0.661	0.621
24	3	✓	✗	✗	4.265	3.931	1.310	1.090
25	3	✗	✓	✓	13.757	16.628	4.225	4.611
26	6	✗	✓	✗	2.133	2.225	0.655	0.617
27	3	✓	✓	✗	5.474	5.060	1.681	1.403
					100.000	**100.000**	**30.709**	**27.728**

In the data set we observe that appointment requests arrive at an average rate of 30.71 requests at the first decision epoch of a day, with a standard deviation of 6.92. At the second decision epoch of a day, we observe requests to arrive at an average rate of 27.73, with a standard deviation of 6.14. Again, it turned out to be impossible to fit probability distributions for all types separately. Instead we model the total number of appointment requests at both decision epochs by a probability distribution and each arriving patient in the simulation is randomly assigned to be one of the 27 patient types, based on the empirically defined probability distribution in the sixth and seventh column of Table 5.5.
In Figure 5.2a the number of MRI appointment requests observed at the first decision

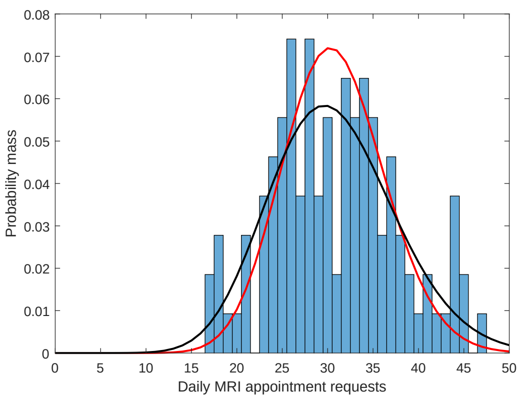

(a) Daily MRI appointment requests at the first decision epoch plotted as a histogram along with the Poisson (red) and negative binomial (black) distribution curve fitted to this data.

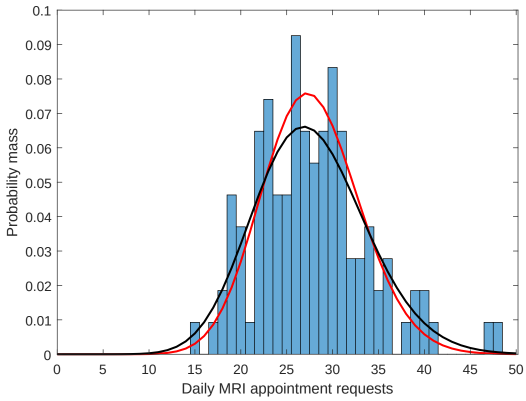

(b) Daily MRI appointment requests at the second decision epoch plotted as a histogram along with the Poisson (red) and negative binomial (black) distribution curve fitted to this data.

Figure 5.2: MRI appointment requests at the first (a) and second (b) decision epoch of a day plotted as a histogram along with the Poisson (red) and negative binomial (black) distribution curve fitted to this data.

epoch is plotted as a histogram along with the Poisson and negative binomial distribution curve fitted to it. The Poisson distribution fitted to the data has a mean $\lambda = 30.71$. The parameters for the negative binomial fitted to the data are $r = 56$, $p = 0.646$. The fit of both distributions with respect to the data is again statistically tested using the Chi-squared goodness of fit test. For the Poisson distribution we found p-value $= 0.134$. For the negative binomial distribution we found p-value $= 0.792$. Hence we model the total number of appointment requests at the first decision epoch as a negative binomial distribution with parameters $r = 56$, $p = 0.646$.

Figure 5.2b shows the equivalent of Figure 5.1a for the second decision epoch. The Poisson distribution fitted to the data has a mean $\lambda = 27.73$. The parameters for the negative binomial fitted to the data are $r = 84$, $p = 0.752$. For the Poisson distribution we found p-value $= 0.753$. For the negative binomial distribution we found p-value $= 0.794$. Hence we model the total number of appointment requests at the second decision epoch as a negative binomial distribution with parameters $r = 84$, $p = 0.752$.

6

Results

In this chapter we present the results of the approximate optimal policy (AOP) found for Rijnstate. Although the Markov Decision Process (MDP) model described in Chapter 4 is formulated in terms of the expected discounted costs over the infinite horizon, the performance of the approximate optimal scheduling policy is evaluated in terms of service levels, mean daily average capacity utilization, mean daily overtime slots used, mean average access times and mean service levels. The service levels are computed for each priority class as the percentage of patients booked within the corresponding access time target. Utilization is defined as the percentage of available regular appointment slots used for delivering MRI examinations and is calculated on a daily basis. Note that utilization values different than 100% are associated with either overtime or idle time and thus utilization values close to 100% are preferred. To measure the AOP in terms of these performance measures, we use a simulation model. In Section 6.1 we shortly explain the used simulation model.

We compare the AOP to the two other policies described below, which we both use with and without blueprint calendars.

- First Available Slot (FAS): Patients are booked as soon as possible. The order in which patients are booked is determined by their access time target. This policy resorts to overtime only when there is no available regular capacity within the booking horizon. Overtime is then booked starting with session 1 and working up to session N of the booking horizon . If this policy is used with blueprint calendars, we refer to the policy as FAS-wBC, and the blueprint calendar determines which appointment slots are suitable for each patient type. Patients can only be booked into slots feasible for their type. If we use this policy without a blueprint calendar, we refer to the policy as FAS-nBC, and patients can be booked into every empty slot.
- Myopic (M): Patients are booked as soon as possible. The order in which patients are booked is determined by their access time target. Unlike the previous policy, this policy resorts to overtime for patients of type t only when there is no available regular capacity within the first n_t sessions of the booking horizon, where $n_t = \max\{n : f^{AT}(k, t, n) < h\}$. Overtime is then booked starting with session 1 and working up to session N. If this policy is used with blueprint calendars, we refer to the policy as M-wBC, and the blueprint calendar determines which

appointment slots are suitable for each patient type. Patients can only be booked into slots feasible for their type. If we use this policy without a blueprint calendar, we refer to the policy as M-nBC, and patients can be booked into every empty slot.

The goal of simulating the FAS-wBC and M-wBC policies is to replicate the MRI appointment scheduling process at the Radiology department in Rijnstate. The simulation must be validated to ensure that the coded policies replicate the actual performance of the department. We report in the results of our validation run in Section 6.2. We present the actual results of all the five policies for Rijnstate in Section 6.3.

6.1 Simulation model

In the simulation model we make booking decisions on the decision epochs as defined in the MDP. That means that we make booking decisions at 12.30 p.m. and at 5 p.m. and therefore deviate from the current practice in Rijnstate, since there patients are booked at times decided by booking agents. It also implies that we book into the MDPs sessions and not on days. As said, we can simulate the AOP policy and, besides, the FAS and Myopic policy both with and without an underlying blueprint calendar. If we choose to simulate a policy with an underlying blueprint calendar, we come across a problem with the currently defined blueprint calendars in Rijnstate and the fact that we use the decision epochs from the MDP.

As we illustrated Chapter 2, in the blueprint calendar time slots are reserved for patients with a small access time target. In practice, these slots are accessible for patients with an access time target of either zero days (same-day patients) or three days. For same-day patients, MRI requests often come in by telephone and an attempt is made to book them today. However, the combination when the request is submitted and the moment on the day where the urgent patient slots are located determine whether this can be achieved or not. In the original blueprint calendar of Rijnstate, the urgent patient slots are located in the morning and afternoon. For a system with the first decision moment at 12.30 p.m., urgent patient slots before this moment are useless. Also, if there are no urgent patient slots after the second decision moment, patients with an access time target of zero days that arrive between the first and second decision epoch on the day will never booked within their access time target. Hence, we slightly adjust the original blueprint calendars of Rijnstate. The total capacity dedicated to a patient type on a day has not changed, but the location of the urgent patient slots has. All urgent patient slots that were originally in the morning were moved to the afternoon on the same day. In addition, we have the urgent patient slots distributed so that there are at least three before 5 p.m. and three after 5 p.m. (which means that there are at least three in the afternoon sessions of the MDP as well as in the evening session of the MDP).

As we explained in Section 5.2, the currently used blueprint calendars determine the patient types that the Radiology department currently distinguishes. However, some of these patient types are identical for the MDP and hence have been been merged into one single patient type to keep the problem size small as possible. This implies that we also need to adjust the blueprint calendars for the merged patient types. We did this as follows. Suppose that patient type 1 and 2 are merged into the new (MDP) patient type 1. Then, if an appointment slot in the original blueprint calendars was either

suitable for patient type 1 or 2 (or both), then this appointment slot is now suitable for the new (MDP) patient type 1.
This implies that we can never compare the service levels for patient types achieved by the AOP to those achieved by the current way of scheduling MRI appointments, simply because both scheduling policies distinguish other patient types. Hence, we compare the performance of the AOP with the FAS-wB and M-wB policies where we use the MDP patient types and the adjusted blueprint calendars. These blueprint calendars can be found in Figure C.3-C.5 in Appendix C.4.
Nevertheless, we would like to make a statement about the access times obtained by the AOP if we would implement this policy in practice in Rijnstate. Hence, we would like a valid simulation model. We illustrate the process of validation and the condition to be able to say something about the performance of the AOP in practice in Figure 6.1.

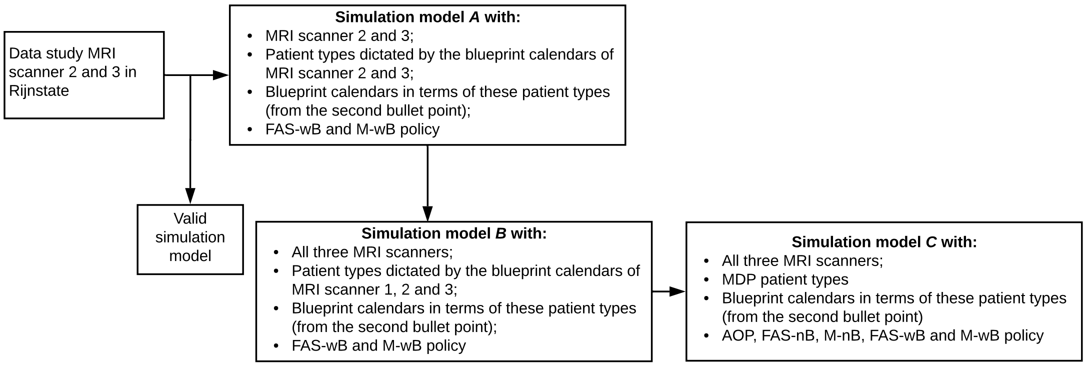

Figure 6.1: Simulation validation process to be able to say something about the performance of the AOP if we would implement this policy in practice in Rijnstate.

We labeled the simulation model used for validation as simulation model *A*. Because we have data available from MRI scanner 2 and MRI scanner 3, we will have to do validation based on a simulation model in which only these two MRI scanners occur. We can then compare the results of this simulation run with the results of our data study in Section 5.1. As Figure 6.1 shows, Simulation model *B* extends model *A* by including MRI scanner 1. Now, if model *A* turned out to be valid for the fictitious Rijnstate situation with two MRI scanners, we may assume model *B* to be valid for the real Rijnstate situation with three MRI scanners. Simulation model *C* uses the MDP patient types and the corresponding adjusted blueprint calendars. Since these blueprint calendars offer more flexibility than those used in simulation model *B*, it may be assumed that the system performance in model *C* is better than in model *B*. Hence, model *C* illustrates the advantages with respect to the system performance if more flexibility is allowed in the blueprint calendars. In summary, if simulation model *A* is considered to be valid, the access times obtained by the AOP in simulation model *C* and other system performances may also be expected in practice if we would implement this policy in practice in Rijnstate.

Since the simulation models includes an integer linear program to determine the AOP at every decision epoch, we implemented the simulation model using AIMMS 4.39 and CPLEX 12.7. To run our simulations, we use a computer with a 2.60 gigahertz Quad Core CPU with 16 gigabyte of RAM.

6.2 Results of the validation run of the simulation model

For the validation of the simulation model (model A in Figure 6.1), we use the patient types and their characteristics of Table 5.3 and the blueprint calendars are given by Figure C.1-C.2 (we have 10 different *blueprint days*). We use a booking horizon of 25 future days, which is equal to a booking horizon of 77 sessions. Overtime is only available in the evening sessions and is for both MRI scanners limited to 9 10-minute slots at Day 1-6 and 10 of the blueprint cycle and 3 slots at Day 7-9.
MRI appointment requests arrive at an average of 22.06 requests at the first decision epoch and another 18.82 requests at the second decision epoch. As said in the previous chapter, the arrival process of MRI appointment requests follow a negative binomial distribution with parameters $r = 64$ and $p = 0.744$ at the first decision epoch and $r = 173$ and $p = 0.902$. Each arriving patient in the simulation is randomly assigned to be on of the 45 patient types, based on the empirically defined probability distribution Table 5.3. The upper bounds on the number of newly arriving appointment request of each type is set at the ceil of three times the mean number of arriving requests at both decision epochs. The access time penalties are defined in Table 6.1 and the overtime cost is 400. The system load (average demand divided by average supply) is 0.98. Table 6.2 displays the results of this simulation run.

Table 6.1: Access time penalties (in sessions) for the fictitious setting with only MRI scanner 2 and MRI scanner 3 present for validation of the simulation model.

Penalty per session within session interval at the first decision epoch of the day										
Patient types	$[0, 2]$	$[3, 5]$	$[6, 11]$	$[12, 14]$	$[15, 17]$	$[18, 20]$	$[21, 32]$	$[33, 35]$	$[36, 38]$	$[39, 77]$
1, 3-5	0	100	250	250	250	250	250	250	250	250
6-14	0	0	0	80	100	150	150	150	150	150
15-21, 24-29	0	0	0	0	0	50	80	125	125	125
30-39, 43-51	0	0	0	0	0	0	0	40	90	100
Penalty per session within session interval at the second decision epoch of the day										
Patient types	$[0, 1]$	$[2, 4]$	$[5, 10]$	$[11, 13]$	$[14, 16]$	$[17, 19]$	$[20, 31]$	$[32, 34]$	$[35, 37]$	$[38, 77]$
1, 3-5	0	100	250	250	250	250	250	250	250	250
6-14	0	0	0	80	100	150	150	150	150	150
15-21, 24-29	0	0	0	0	0	50	80	125	125	125
30-39, 43-51	0	0	0	0	0	0	0	40	90	100

Table 6.2: Results from the validation of the coded FAS-wBC and M-wBC policies (Simulation model A from Figure 6.1). Simulation average is abbreviated to Sim. av. and absolute different to Abs. dif.

Performance metric	**Data analysis**	**FAS-wB**		**M-wB**	
		Sim. av.	**Abs. dif.**	**Sim.av.**	**Abs. dif.**
Mean number of scans per day	39.80	40.87	1.07	41.42	1.62
Mean access time (in days)	16.1	3.53	12.57	2.76	13.34
Mean daily capacity utilization (in %)	83.10	93.10	10.00	93.17	10.17
Mean daily overtime (in min.)	8.57	0.00	8.57	20.49	11.92

In Table 6.2 we only included the performance indicators for which we were also able to perform a data study. The absolute difference in the number of scans per day is small enough so that it may pass a validation exercise. On the other hand, the absolute differences we observe for the other performance metrics are so large that with most probably no known validation exercises will pass. Hence, to our regret, we need to conclude the simulation model is not a valid representation of the current practice. Consequently, we cannot draw conclusions about the access times obtained by the AOP if we would implement this policy in practice in Rijnstate. Or at least we have to be very careful about this. Nevertheless, we can compare the AOP to the other policies for the case study obtained from the data provided by Rijnstate.

If we take a closer look to the results obtained in this simulation, we observe the mean access time in the simulation to be much smaller than the mean access time from the data study. Since the system load is less than one, requests received at either the first or the second decision epoch of a specific day in the simulation, can be booked in the near future because the system does not *fill up*, i.e., the daily demand is not structurally higher than the available capacity which would have forced the system to book patients either further into the future or in overtime. We could not find a satisfactory explanation why the average access time in the data is so high cannot be explained. We have carefully examined the data and looked at whether there might have been a backlog, so that the first available time slot is just very far in the future. If this were the case, we would observe for a large part of the data points that access time would be greater than a certain value, i.e., the 10^{th} or 25^{th} percentile would be large.
An explanation for the higher utilization in simulation may be partly due to the adjustment of the blueprint calendars. The urgent patient slots can always be used in the simulation because they are located after the decision moments. In reality, an urgent patient slot may remain unused because the first urgent MRI request is only submitted after the first urgent patient slot on the day. However, it is very doubtful whether this can cause the big absolute difference that we observe.
The average overtime observed in the data is not that bad in itself. For the FAS-wB it applies that it never books in overtime, because the regular capacity of the entire booking horizon is never fully booked. In addition, the simulation is based on deterministic service times, which implies that overtime is only induced (in the M-wB policy) by patients who are booked in overtime due to their access time target restriction. In reality, of course, there are many more factors that determine overtime. It is therefore to be expected that the simulation model will not validate for this performance metric because the factors that determine the overtime differ strongly.

6.3 Results for Rijnstate

We now evaluate the AOP by simulating its performance for a large scale case study at the Radiology department of Rijnstate. The patient types, their capacity requirements and the patient type-MRI scanner compatibility restrictions are displayed in Table 5.5 in Chapter 5. Patients can receive treatment on three MRI scanners with on average 1.70 suitable MRI scanners per patient type.
The arrival process of MRI appointment requests follow a negative binomial distribution with parameters $r = 56$ and $p = 0.646$ at the first decision epoch and $r = 84$ and $p = 0.752$. Each arriving patient in the simulation is randomly assigned to be on of the 27 patient types, based on the empirically defined probability distribution Table 5.5. The upper bounds on the number of newly arriving appointment request of each type

is set at the ceil of three times the mean number of arriving requests at both decision epochs. In total, the expected daily numbers of MRI requests consist of 3.0 patients with an access time penalty from day 2 (1.9 at the first decision epoch and 1.1 at the second), 12.2 patients with a penalty from day 4 (6.9 at the first decision epoch and 5.3 at the second), 13.4 with a penalty from day 6 (6.8 at the first decision epoch and 6.6 at the second) and 30.0 with a penalty from day 11 (15.2 at the first decision epoch and 14.8 at the second). The access time penalties are defined in Table 6.3 and the overtime cost is 400.

Table 6.3: Access time penalties (in sessions).

Penalty per session within session interval at the first decision epoch of the day										
Patient types	$[0,2]$	$[3,5]$	$[6,11]$	$[12,14]$	$[15,17]$	$[18,20]$	$[21,32]$	$[33,35]$	$[36,38]$	$[39,77]$
1-5	0	100	250	250	250	250	250	250	250	250
6-11	0	0	0	80	100	150	150	150	150	150
12-19	0	0	0	0	0	50	80	125	125	125
20-27	0	0	0	0	0	0	0	40	90	100
Penalty per session within session interval at the second decision epoch of the day										
Patient types	$[0,1]$	$[2,4]$	$[5,10]$	$[11,13]$	$[14,16]$	$[17,19]$	$[20,31]$	$[32,34]$	$[35,37]$	$[38,77]$
1-5	0	100	250	250	250	250	250	250	250	250
6-11	0	0	0	80	100	150	150	150	150	150
12-19	0	0	0	0	0	50	80	125	125	125
20-27	0	0	0	0	0	0	0	40	90	100

Table 6.4 displays the system configurations used to simulate the various policies. The discount factor is 0.99. Note that the system load is different for the AOP, FAS-nB and M-nB policies compared to the FAS-wB and M-wB policies because the average regular capacity following from the blueprint calendars is slightly higher than in the MDP. The number of states in the MDP model (and thus the number of variables in the linear program) would be in the order of 10^{629} and the number state-action combinations (which equals the number of constraints in the original linear program) would even be more. All five policies were simulated for 1000 days with statistics collected for each of 10 runs after a warm-up period of 500 days. The execution time of the (column generation) algorithm used to obtain the AOP was around 10 hours. Additionally, the simulation took around 24 hours. These execution times are satisfactory, considering that the coefficients defining the approximate optimal policies only need to be redetermined when there is a significant change in the problem parameters. It is also important to note that solving the integer linear programming model used to identify the approximate optimal actions, which is done on twice a day, takes less than two seconds.

Table 6.4: System configurations for simulating the performance of the AOP, FAS-nB, M-nB, FAS-wB and M-wB policies for a large scale case study at the Radiology department of Rijnstate.

Configurations for the AOP, FAS-nB, M-nB, FAS-wB and M-wB policy:	
Number of MRI scanners	3
Length of appointment slots	10 minutes
Length of booking horizon	77 sessions (5 weeks with 5 working days)
Overtime cost	400
Configurations for the AOP, FAS-nB and M-nB policy:	
$C_{mn}^{R,1} = C_{m,n-1}^{R,2}$	27, for all $m \in \mathcal{M}, n = 3, 6, 9, \ldots, 75$ (duration of morning sessions in num. of slots.)
$C_{mn}^{R,1} = C_{m,n-1}^{R,2}$	27, for all $m \in \mathcal{M}, n = 1, 4, 7, \ldots, 76$ (duration of afternoon sessions in num. of slots. Note that we always need to have $n \geq 1$.)
$C_{mn}^{R,1} = C_{m,n-1}^{R,2}$	6, for all $m \in \mathcal{M}, n = 2, 5, 8, \ldots, 77$ (duration of evening sessions in num. of slots.)
$C_{mn}^{OT,1} = C_{m,n-1}^{OT,2}$	1, for all $m \in \mathcal{M}, n = 3, 6, 9, \ldots, 75$ (number of overtime slots in morning sessions.)
$C_{mn}^{OT,1} = C_{m,n-1}^{OT,2}$	1, for all $m \in \mathcal{M}, n = 1, 4, 7, \ldots, 76$ (number of overtime slots in morning sessions. Note that we always need to have $n \geq 1$.)
$C_{mn}^{OT,1} = C_{m,n-1}^{OT,2}$	9, for all $m \in \mathcal{M}, n = 2, 5, 8, \ldots, 77$ (number of overtime slots in evening sessions.)
System load (average demand divided by average regular supply)	0.97
Configurations for the FAS-wB and M-wB policy:	
Blueprints are given in Figure C.3-C.3 in Appendix C.4	
Blueprint cycle length	10
Overtime slots Day 1-6 & 10	9
Overtime slots Day 7, 8, 9	3
System load	0.94

Table 6.5: Meand and 95%-confidence intervals of the service levels for each patient type, defined as the percentages of patients examined within a certain number of days (their access time targets in bold) and patient types' average access time. Note that today is defined as day 0, tomorrow day 1 and so on.

Patient type	AOP		FAS-nB		M-nB		FAS-wB		M-wB	
	SL	AT	SL	AT	SL	AT	SL	AT	SL	AT
1	80.25 ± 1.36	0.23 ± 0.02	17.52 ± 2.55	0.84 ± 0.03	17.68 ± 0.76	0.83 ± 0.02	66.59 ± 1.96	0.27 ± 0.03	68.75 ± 0.79	0.26 ± 0.03
2	83.81 ± 2.53	0.18 ± 0.03	6.33 ± 0.84	0.99 ± 0.02	6.79 ± 0.18	0.97 ± 0.01	46.25 ± 3.70	0.36 ± 0.03	48.77 ± 2.91	0.36 ± 0.03
3	73.50 ± 1.83	0.33 ± 0.03	10.88 ± 1.63	0.94 ± 0.03	11.02 ± 0.51	0.93 ± 0.02	53.03 ± 1.00	0.57 ± 0.02	56.17 ± 1.10	0.56 ± 0.02
4	47.61 ± 3.96	1.25 ± 0.13	4.31 ± 2.55	1.07 ± 0.09	5.83 ± 0.59	1.03 ± 0.06	15.43 ± 1.76	0.64 ± 0.03	15.80 ± 1.02	0.64 ± 0.03
5	79.69 ± 1.44	0.25 ± 0.02	42.66 ± 5.76	0.59 ± 0.07	42.82 ± 1.78	0.58 ± 0.06	61.29 ± 2.42	0.17 ± 0.02	68.27 ± 0.88	0.15 ± 0.02
6	100.00 ± 0.00	2.35 ± 0.05	100.00 ± 0.00	1.01 ± 0.01	100.00 ± 0.00	1.01 ± 0.01	100.00 ± 0.00	1.79 ± 0.26	100.00 ± 0.00	1.73 ± 0.26
7	98.58 ± 0.32	2.53 ± 0.03	100.00 ± 0.00	1.03 ± 0.02	100.00 ± 0.00	1.03 ± 0.02	100.00 ± 0.00	1.11 ± 0.05	100.00 ± 0.00	1.11 ± 0.05
8	98.68 ± 0.56	2.55 ± 0.05	100.00 ± 0.00	1.03 ± 0.02	100.00 ± 0.00	1.03 ± 0.02	100.00 ± 0.00	1.73 ± 0.35	100.00 ± 0.00	1.73 ± 0.35
9	98.38 ± 0.26	2.55 ± 0.04	100.00 ± 0.00	1.04 ± 0.02	100.00 ± 0.00	1.04 ± 0.02	100.00 ± 0.00	1.76 ± 0.18	100.00 ± 0.00	1.74 ± 0.18
10	86.02 ± 1.53	2.84 ± 0.01	100.00 ± 0.00	1.07 ± 0.02	100.00 ± 0.00	1.07 ± 0.02	100.00 ± 0.00	1.58 ± 0.18	100.00 ± 0.00	1.57 ± 0.18
11	100.00 ± 0.00	2.49 ± 0.02	100.00 ± 0.00	1.05 ± 0.03	100.00 ± 0.00	1.04 ± 0.02	100.00 ± 0.00	1.69 ± 0.16	100.00 ± 0.00	1.69 ± 0.16
12	100.00 ± 0.00	4.87 ± 0.01	100.00 ± 0.00	1.06 ± 0.03	100.00 ± 0.00	1.06 ± 0.03	100.00 ± 0.00	3.54 ± 0.52	100.00 ± 0.00	3.54 ± 0.52
13	100.00 ± 0.00	4.68 ± 0.06	100.00 ± 0.00	1.09 ± 0.03	100.00 ± 0.00	1.09 ± 0.03	100.00 ± 0.00	3.30 ± 0.59	100.00 ± 0.00	3.30 ± 0.59
14	100.00 ± 0.00	4.89 ± 0.04	100.00 ± 0.00	1.11 ± 0.02	100.00 ± 0.00	1.11 ± 0.02	100.00 ± 0.00	3.41 ± 0.51	100.00 ± 0.00	3.41 ± 0.51
15	99.68 ± 0.32	4.86 ± 0.03	100.00 ± 0.00	1.21 ± 0.04	100.00 ± 0.00	1.21 ± 0.04	100.00 ± 0.00	3.48 ± 0.96	100.00 ± 0.00	3.48 ± 0.96
16	100.00 ± 0.00	4.93 ± 0.01	100.00 ± 0.00	1.08 ± 0.03	100.00 ± 0.00	1.07 ± 0.03	100.00 ± 0.00	4.11 ± 0.07	100.00 ± 0.00	4.11 ± 0.07
17	99.06 ± 0.25	4.91 ± 0.02	100.00 ± 0.00	1.15 ± 0.05	100.00 ± 0.00	1.14 ± 0.04	100.00 ± 0.00	4.15 ± 0.73	100.00 ± 0.00	4.15 ± 0.73
18	96.71 ± 0.71	4.76 ± 0.02	100.00 ± 0.00	1.09 ± 0.03	100.00 ± 0.00	1.09 ± 0.02	100.00 ± 0.00	4.83 ± 0.71	100.00 ± 0.00	4.83 ± 0.71
19	99.72 ± 0.16	4.79 ± 0.02	100.00 ± 0.00	1.04 ± 0.02	100.00 ± 0.00	1.04 ± 0.02	100.00 ± 0.00	4.71 ± 0.38	100.00 ± 0.00	4.71 ± 0.38
20 - 27	100.00 ± 0.00	9.94 ± 0.04	100.00 ± 0.00	1.43 ± 0.04	100.00 ± 0.00	1.43 ± 0.04	100.00 ± 0.00	4.93 ± 0.91	100.00 ± 0.00	4.93 ± 0.91

The results for this simulation are summarized in Table 6.5, Table 6.6 and Figure 6.2. Table 6.5 shows the average service levels and the average access time for each patient type and the corresponding 95%-confidence intervals. Table 6.6 shows the average daily utilization of all MRI scanners for the five policies. Figure 6.2 shows the patient booking distribution across MRI scanners. Additionally results can be found in Appendix D.1. Here, Table D.1 shows the percentages of patients examined within a certain number of days for patient types 1-5. Table D.3 shows the patient booking distribution, in percentages, across MRI scanners that are graphically presented in Figure 6.2.

We first focus on the the service levels and access times. With respect to the service levels for patient types with an access time target of 0 days (needs to be examined today), the AOP obtains the best results and the FAS-nB policy performs worst. We observe the following for the FAS-nB policy. Since the system load is less than one, requests received at either the first or the second decision epoch of day k in the simulation, can often be booked on day $k+1$.Then, if the simulation proceeds to day $k+1$, there is not enough capacity available to serve patients whose request arrives at day $k+1$ and needs to be examined on the same day. Hence, the patients are all booked on the next day (which would be day $k+2$ in the simulation) and are one day late. For the majority of patients with access time targets greater than 0 days, it applies that they are also booked on day $k+2$ of the simulation and if this is not possible anymore, they are booked on day $k+3$. However, for all patient types this is still within their access time target and hence the FAS-nB policy always manage to schedule patient types with an access time target greater than 0 days within their access time target.
The same applies to the M-nB policy, since the M-nB copies the action of the FAS-nB policy for patient types with an access time target greater than 0 days. However, the M-nB policy performs slightly better for patient types with an access time target of 0 days, since the policy resorts to overtime for these patient types. This results in more overtime slots used for this policy, as Table 6.6 also displays.

Based on the values of V^*_{mn} and $\widetilde{V}^*_{mn}$ we would expect that the AOP obtains the best results with respect to the service times for patient with an access time target of 0 days. The values of V^*_{mn} and $\tilde{V}^*_{mn}$ corresponding to sessions n today and tomorrow are much larger than values of V^*_{mn} and $\tilde{V}^*_{mn}$ corresponding to sessions n further into the future, i.e., the AOP protects capacity in the sessions of today and tomorrow for same-day patients. Furthermore, we observe that for fixed m, V^*_{mn} and $\widetilde{V}^*_{mn}$ decreases if n increases, but only if n corresponds to a different day of the booking horizon. That is, if $n_1 < n_2$ are sessions both belonging to the same day of the booking horizon, then $V^*_{mn_1} = V^*_{mn_2}$, but if n_1 and n_2 belong to consecutive days of the booking horizon, then $V^*_{mn_1} > V^*_{mn_2}$, where still $n_1 < n_2$. Also, for fixed n there are small differences between the different MRI scanners. The system prefers to first book patients on MRI scanner 3, followed by MRI scanner 2 and MRI scanner 1.

For the FAS-wB and M-wB policy, the system performs excellently with respect to the service levels for patient types 6-11. These are the types with an access time target of 3 days and have the most feasible time slots in the blueprint calendars. Regarding the patient types with an access time target of 0 days, thanks to the blueprint calendars there are slots reserved for this group of patients and the FAS-wB and M-wB policies outperform the FAS-nB and M-nB, respectively, but the AOP still performs best. Finally, after patient types 1-11 have been booked, there are always still enough suitable slots free within 5 days to book patient types 12-19 within their 5-day access time target. Logically, for patient types 20-27, there are always enough suitable time slots within the next 10 days to book these patient types within their access time target. The AOP books

Table 6.6: Average daily utilization of each MRI scanner in percentages and the average daily overtime slots used, for all of the five policies. Note that utilization values greater than 100% are associated with overtime.

MRI scanner	Utilization				
	AOP	**FAS-nB**	**M-nB**	**FAS-wB**	**M-wB**
1	86.51 ± 0.37	98.18 ± 0.08	103.19 ± 0.08	92.82 ± 0.35	92.67 ± 0.34
2	99.37 ± 0.30	97.57 ± 0.15	103.57 ± 0.16	93.68 ± 0.22	94.00 ± 0.21
3	97.96 ± 0.20	88.16 ± 0.66	88.13 ± 0.65	90.15 ± 0.27	90.00 ± 0.27
Overall	94.61 ± 0.26	94.64 ± 0.29	97.63 ± 0.29	92.20 ± 0.23	92.20 ± 0.23
	Overtime slots used				
	AOP	**FAS-nB**	**M-nB**	**FAS-wB**	**M-wB**
1	0.25 ± 0.01	0.00 ± 0.00	5.59 ± 0.20	0.00 ± 0.00	2.86 ± 0.20
2	1.30 ± 0.06	0.00 ± 0.00	6.51 ± 0.18	0.00 ± 0.00	3.18 ± 0.17
3	0.85 ± 0.08	0.00 ± 0.00	0.01 ± 0.00	0.00 ± 0.00	0.01 ± 0.00
Overall	2.40 ± 0.14	0.00 ± 0.00	12.10 ± 0.36	0.00 ± 0.00	6.05 ± 0.36

preferably on MRI scanner 3, followed by MRI scanner 2. This can also been observed from Figure 6.2a. Apparently. the blueprint calendars divide the suitable slots of patient types in such a way over all three MRI scanners that the booking distribution across MRI scanners of the FAS-wB and M-wB policies is quite similar to that of the AOP. Since the FAS-nB and M-wB policies always try to book on MRI scanner 1 first, followed by MIR scanner 2, we observe the opposite of Figure 6.2a in Figure 6.2b.

Table 6.6 shows the average daily utilization per MRI scanner for all policies. As expected, for MRI scanner 1 ,the utilization of the FAS-nB and M-nB policies are higher than for the AOP and the opposite applies to MRI scanner 3. Furthermore, we observe that the blueprint calendars reduce the utilization rate for both MRI scanner 1 and MRI scanner 2. Most probably, in the blueprint calendars for these two MRI scanners, the blueprint calendars dedicate more slots to one of the patient type than the demand of this type. This causes the MRI scanners to be idle for some periods. Finally, we also observe a slight difference in the overall system utilization. However, here we do have to remind that the daily average capacity is slightly higher for the FAS-wB and M-wB policies.
It is reasonable that the FAS-nB and FAS-wB policies do not generate overtime at all. As we explained in Section 6.2, if the system's load is strictly less than one, there is no overflow of appointment requests. So, the system does not *fill up*, i.e., the daily demand is not structurally higher than the available capacity which would have forced the system to book patients either further into the future or in overtime. As the FAS-nB and FAS-wB policies only book patients in overtime after all regular slots of the complete booking horizon are filled, this will never happen. We observe the AOP to

perform better than the M-nB and the M-wB policy as the AOP reserves capacity for patient type1-5 tomorrow, whilst the M-nB tries to book tomorrow as full as possible tomorrow, with the result that the system tomorrow needs to resort to overtime to book patient types 1-5 within their access time target. The blueprint calendars do also reserve some capacity for the urgent patient types. However, they cannot do this as well as the AOP.

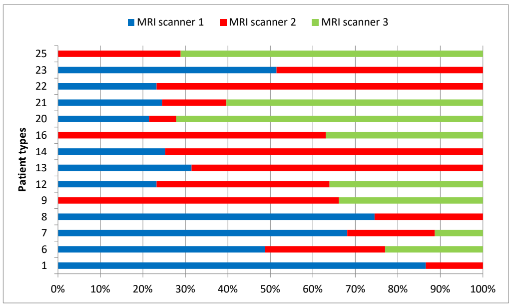

(a) Booking distribution across MRI scanners per patient type for the AOP.

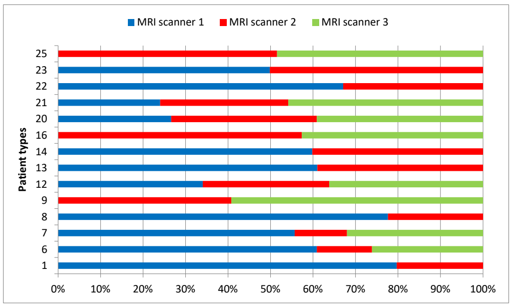

(b) Booking distribution across MRI scanners per patient type for the FAS-nB policy.

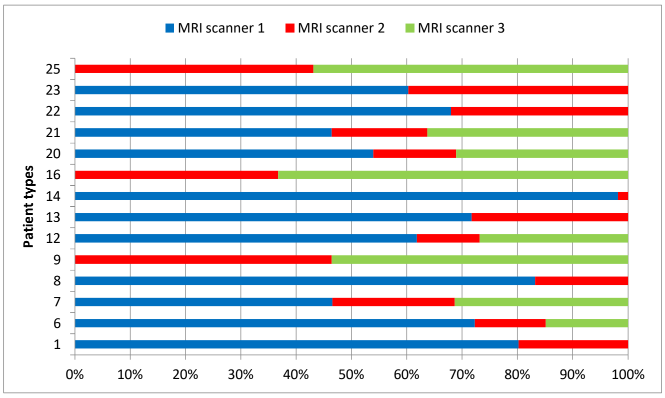

(c) Booking distribution across MRI scanners per patient type for the FAS-wB policy.

Figure 6.2: Booking distribution across MRI scanners per patient type. We only visualized the distribution for the AOP, FAS-nB and Fas-wB policy, since the distributions for the M-nB and M-wB policies are very similar to the distributions for the FAS-nB and Fas-wB policies, respectively.

6.4 Results for a heavier loaded system

The results from the previous section show well what influence the AOP has on the service levels of patient types 1-5. However, we are also curious about what service levels the AOP can obtain for patient types with a larger access time target if the capacity is a binding factor (i.e. we have a system load greater than 1). Hence, in this section, we study the performance of the AOP for the instance where we have the same patient types and arrival processes as in Section 6.3, but instead we set the daily available capacity of each MRI scanner equal to 9 hours (that is one hour less). Since the booking horizon consists, again, of 77 sessions, this means that we have

$$\begin{aligned}
C^{R,1}_{mn} &= C^{R,2}_{m,n-1} = 24, \text{ for all } m \in \mathcal{M}, n = 3,6,9,\ldots,75,\\
C^{R,1}_{mn} &= C^{R,2}_{m,n-1} = 24, \text{ for all } m \in \mathcal{M}, n = 1,4,7,\ldots,76,\\
C^{R,1}_{mn} &= C^{R,2}_{m,n-1} = 6, \text{ for all } m \in \mathcal{M}, n = 2,5,8,\ldots,77,\\
C^{OT,1}_{mn} &= C^{OT,2}_{m,n-1} = 1, \text{ for all } m \in \mathcal{M}, n = 3,6,9,\ldots,75,\\
C^{OT,1}_{mn} &= C^{OT,2}_{m,n-1} = 1, \text{ for all } m \in \mathcal{M}, n = 1,4,7,\ldots,76,\\
C^{OT,1}_{mn} &= C^{OT,2}_{m,n-1} = 9, \text{ for all } m \in \mathcal{M}, n = 2,5,8,\ldots,77,
\end{aligned}$$

The upper bounds on the number of newly arriving appointment request of each type is again set at the ceil of three times the mean number of arriving requests at both decision epochs. The access time penalties are already given in Table 6.3 and the overtime cost is again 400. The discount factor is again 0.99. The system load for this instance is 1.07. For this instance we compare the AOP only to the FAS-nB and M-nB policies and not to the FAS-wB and M-wB policies, because the system load is so different (1.07 against 0.94) that it would be an unfair comparison.
The execution time of the (column generation) algorithm used to obtain the AOP was again around 10 hours. Additionally, the simulation took around 21 hours.

The results for this simulation are summarized in Table 6.7, Table 6.8 and Figure 6.2. Additionally results can be found in Appendix D.2. Here, Table D.4 shows the percentages of patients examined within a certain number of days for patient types 1-5. Table D.2 shows the patient booking distribution, in percentages, across MRI scanners that are graphically presented in Figure 6.3.

With respect to the service levels for patient types with an access time target of 0 days (needs to be examined today), the M-nB policy obtains the best results and the FAS-nB policy performs worst. As follows from Table 6.8, the price the M-nB policy pays for these service levels, is a lot of overtime. So, what probably happens is that on the majority of simulation days the overtime limit of 9 appointment slots per MRI scanner is large enough to book all demand from patient type 1-5 in. On the other side, on approximately a quarter of the days, the AOP does not reserve enough capacity for patient types 1-5. For patient types other than 1-5, the AOP and M-nB achieve similar service levels. The FAS-nB policy achieves by far the worst performances.

As expected, the average daily utilization is more than 100%, since there is overflow of demand. Opposite to the results from the previous section, now the FAS-nB policy does generate overtime. Because of the overflow, all the regular appointment slots in the booking horizon become occupied and hence the FAS-nB needs to resort overtime. We observe the AOP to perform better than the FAS-nB and M-nB policy with respect to the average daily overtime slots used. The patient booking distribution is similar to the one in the previous section (compare Figure 6.3 to Figure 6.2).

Table 6.7: Mean and 95%-confidence intervals of the service levels for each patient type, defined as the percentages of patients examined within a certain number of days. Note that today is defined as day 0, tomorrow day 1 and so on.

Patient type	AOP	FAS-nB	M-nB
1	77.594 ± 1.976	0.000 ± 0.000	87.806 ± 2.105
2	76.283 ± 2.896	0.000 ± 0.000	74.909 ± 3.541
3	75.738 ± 1.597	0.000 ± 0.000	85.766 ± 1.710
4	44.808 ± 5.072	0.000 ± 0.000	69.169 ± 7.039
5	80.723 ± 1.433	0.000 ± 0.000	95.822 ± 1.648
6	99.878 ± 0.086	0.000 ± 0.000	75.913 ± 2.171
7	69.427 ± 1.809	0.000 ± 0.000	69.186 ± 0.513
8	72.206 ± 2.493	0.000 ± 0.000	66.413 ± 1.400
9	69.067 ± 1.118	0.000 ± 0.000	64.536 ± 1.805
10	73.042 ± 1.428	0.000 ± 0.000	63.055 ± 3.152
11	99.467 ± 0.268	0.000 ± 0.000	62.162 ± 3.313
12	99.668 ± 0.163	0.000 ± 0.000	100.000 ± 0.000
13	99.922 ± 0.082	0.000 ± 0.000	100.000 ± 0.000
14	99.151 ± 0.345	0.000 ± 0.000	100.000 ± 0.000
15	99.148 ± 0.260	3.539 ± 1.469	100.000 ± 0.000
16	97.507 ± 0.468	0.000 ± 0.000	100.000 ± 0.000
17	90.200 ± 0.566	0.271 ± 0.109	100.000 ± 0.000
18	93.950 ± 0.265	0.000 ± 0.000	100.000 ± 0.000
19	99.430 ± 0.163	0.000 ± 0.000	100.000 ± 0.000
20	100.000 ± 0.000	0.000 ± 0.000	100.000 ± 0.000
21	100.000 ± 0.000	0.162 ± 0.000	100.000 ± 0.000
22	100.000 ± 0.000	5.555 ± 0.000	100.000 ± 0.000
23	100.000 ± 0.000	0.599 ± 0.000	100.000 ± 0.000
24	100.000 ± 0.000	37.723 ± 2.402	100.000 ± 0.000
25	100.000 ± 0.000	2.5503 ± 0.257	100.000 ± 0.000
26	100.000 ± 0.000	39.832 ± 2.609	100.000 ± 0.000
27	100.000 ± 0.000	37.164 ± 1.512	100.000 ± 0.000

Table 6.8: Average daily utilization of each MRI scanner in percentages and the average daily overtime slots used, for all of the five policies. Note that utilization values greater than 100% are associated with overtime.

MRI scanner	Utilization		
	AOP	**FAS-nB**	**M-nB**
1	103.14 ± 0.68	108.39 ± 0.69	107.02 ± 0.66
2	107.48 ± 0.49	107.76 ± 0.75	106.22 ± 0.61
3	104.77 ± 0.37	99.237 ± 0.20	102.23 ± 0.41
Overall	105.13 ± 0.48	105.13 ± 0.53	105.16 ± 0.52
	Overtime slots used		
	AOP	**FAS-nB**	**M-nB**
1	2.60 ± 0.24	4.77 ± 0.36	8.01 ± 0.33
2	3.31 ± 0.27	4.37 ± 0.38	8.44 ± 0.30
3	1.01 ± 0.16	0.38 ± 0.08	1.71 ± 0.18
Overall	6.93 ± 0.65	9.52 ± 0.79	18.16 ± 0.76

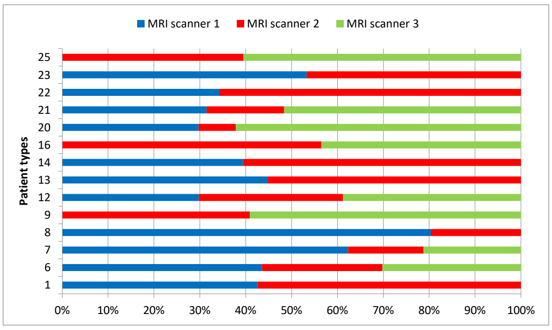

(a) Booking distribution across MRI scanners per patient type for the AOP.

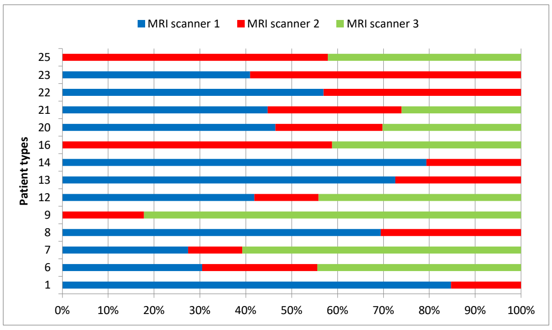

(b) Booking distribution across MRI scanners per patient type for the FAS-nB policy.

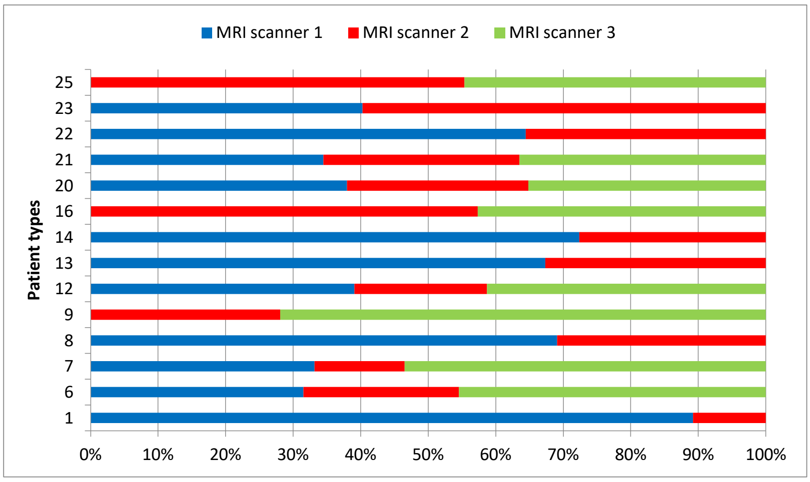

(c) Booking distribution across MRI scanners per patient type for the FAS-wB policy.

Figure 6.3: Booking distribution across MRI scanners per patient type for the AOP, FAS-nB and M-nB policy.

7

An extension to the near-online multipriority patient scheduling model

In this chapter, we present an extension to the near-online multipriority patient scheduling model from Chapter 4. We assume the reader to be familiar with this Markov Decision Process (MDP) and for a comprehensive overview of the notation used, we refer to Table 4.1. This extension is based on observations on the values and structure of the V^*_{mn}'s and $\widetilde{V}^*_{mn}$'s observed for the instances in Chapter 6.

As we mentioned in Section 6.2, the values of V^*_{mn} and $\widetilde{V}^*_{mn}$ corresponding to sessions n today and tomorrow are much larger than values of V^*_{mn} and $\tilde{V}^*_{mn}$ corresponding to sessions n further into the future, i.e., the AOP protects capacity in the sessions of today and tomorrow for same-day patients. Furthermore, we observe that for fixed m, V^*_{mn} and $\widetilde{V}^*_{mn}$ decreases if n increases, but only if n corresponds to a different day of the booking horizon. That is, if $n_1 < n_2$ are sessions both belonging to the same day of the booking horizon, then $V^*_{mn_1} = V^*_{mn_2}$, but if n_1 and n_2 belong to consecutive days of the booking horizon, then $V^*_{mn_1} > V^*_{mn_2}$, where still $n_1 < n_2$. Also, for fixed n there are small differences between the different MRI scanners. The system prefers to first book patients on MRI scanner 3, followed by MRI scanner 2 and MRI scanner 1. Hence, the AOP incorporates the following hierarchy. We assume the process to be at the first decision epoch of a day and need to schedule a patient with an access time target equal to 11 sessions. Then, we would first try to book this patient in the regular slots of session 11 of MRI scanner 3, followed by the regular slots of session 11 of MRI scanner 2 and MRI scanner 1. If this is not possible, we try to book the patient, in the same order of MRI scanner, into either session 10 or session 9. If this is not possible, we need to determine the *cheaper* of the following three options: book the patient either in the overtime of session 9, 10, or 11, or in a session prior to 9, or to violate the access time target and book the patient later than session 11. Now, this hierarchy is independent of the patient type, as $\tilde{V}^*_{mn}$ (we still assume to be at the first decision epoch of the day) does not depend on the patient type. Hence, in this chapter we present an extension in which the value associated with using capacity on a given day and MRI scanner does depend on the treatment type. This model has a larger state space and larger action sets and so it is interesting to study what this does to the computation time of the column generation algorithm. On the other side, it allows to define the direct costs in a more elegant way.

7.1 The MDP formulation

7.1.1 The state space

In this extended model we again assume each day to be divided into three sessions: a morning session, an afternoon session and an evening session. We also have a rolling booking horizon.

At both decision epochs of a day, the number of time slots already booked on each MRI scanner during all sessions of the booking horizon is known, as well as the number of appointment requests from each patient type to be scheduled at the current decision epoch. Thus, a typical state of the system, denoted by $\mathbf{s}$, takes the form

$$\mathbf{s} = (k, \mathbf{x}, \mathbf{y}) = (k, x_{tmn}, y_t), \quad t \in \mathcal{T}, m \in \mathcal{M}, n \in \mathcal{N},$$

where x_{tmn} is the number of appointment slots already reserved for patient type t on MRI scanner m during session n of the booking horizon and y_t is the number of type t patients waiting to be booked. In this state vector, k can take the value 1 or 2, corresponding to the decision epoch during the day.

Observed from the first decision epoch of the day, we have capacities $C_{mn}^{R,1}$ and limited overtime $C_{mn}^{OT,1}$. Observed from the second decision epoch of the day, we have capacities $C_{mn}^{R,2}$ and $C_{mn}^{OT,2}$, for all $m \in \mathcal{M}, n \in \mathcal{N}$.

For $k \in \{1, 2\}$, we define the sets

$$\mathcal{S}_k = \left\{ (k, \mathbf{x}, \mathbf{y}) \;\middle|\; \begin{array}{ll} \sum\limits_{t \in \mathcal{T}(m)} x_{tmn} \leq C_{mn}^{R,k} + C_{mn}^{OT,k}, & \text{for all } m \in \mathcal{M} \text{ and } n \in \mathcal{N}; \\ y_t \leq Q_t^k, & \text{for all } t \in \mathcal{T}; \\ (\mathbf{x}, \mathbf{y}) \in \mathbb{N}_0^{MN} \times \mathbb{N}_0^T & \end{array} \right\}. \tag{7.1}$$

Here Q_t^k is the maximum number of MRI appointment requests of type t present in the system (and thus waiting to be booked) at decision epoch k of a day. Truncating the maximum number of MRI appointment requests of each patient type waiting to be booked is necessary to keep the state space finite, but the maximum number can be set sufficiently high as to be of little practical significance. Note here that we here no longer assume that the number of appointment requests observed at a decision epoch only concerns newly arrived requests. It is therefore now possible to postpone booking decisions. We will come back to this in Section 7.1.2.

Because we experience a rolling booking horizon, in state $\mathbf{s}_1 = (1, \mathbf{x}, \mathbf{y})$ there are no appointments scheduled during sessions $N-1$ and N, that is, $x_{m,N-1} = x_{mN} = 0$, for all $m \in \mathcal{M}$. Similarly, in state $\mathbf{s}_2 = (2, \mathbf{x}, \mathbf{y})$ there are no bookings in session N, that is $x_{mN} = 0$, for all $m \in \mathcal{M}$. Furthermore, if $m \notin \mathcal{M}(t)$, it is required that $x_{tmn} = 0$, for all $n \in \mathcal{N}$.

The state space of the MDP is given by

$$\mathcal{S} = \mathcal{S}_1 \cup \mathcal{S}_2. \tag{7.2}$$

7.1.2 The action sets

At each decision epoch, the booking agent's task is to decide on which MRI scanner and during which session to schedule each of the patients waiting to be booked. Thus, a vector of possible actions can be written as $\mathbf{a} = (a_{tmn})_{t\in\mathcal{T},m\in\mathcal{M},n\in\mathcal{N}}$, where a_{tmn} is the number of type t patients to book on MRI scanner m during session n. Note that, once a patient is assigned to a session, a second level of scheduling is needed which assigns patients to specific appointment times. That goes beyond the scope of this MDP model.

Input parameters of the decision are y_t, the number of type t patients to be booked, and $\sum_{t\in\mathcal{T}(m)} x_{tmn}$, the number of time slots already occupied on MRI scanner m during session n of the booking horizon. We experience a rolling horizon and have two decision epochs in each day and hence we face different constraints for action $\mathbf{a}$ to be valid in states $\mathbf{s}_1 = (1, \mathbf{x}, \mathbf{y})$ and $\mathbf{s}_2 = (2, \mathbf{x}, \mathbf{y})$. To be valid, any action in $\mathbf{s}_1$ must satisfy the following constraints:

$$\sum_{n\in\mathcal{N}} \sum_{m\in\mathcal{M}(t)} a_{tmn} \leq y_t, \qquad \text{for all } t \in \mathcal{T}, \tag{7.3}$$

$$\sum_{t\in\mathcal{T}(m)} x_{tmn} + d_t a_{tmn} \leq C_{mn}^{R,1} + C_{mn}^{OT,1}, \qquad \text{for all } m \in \mathcal{M} \text{ and } n \in \mathcal{N}, \tag{7.4}$$

$$a_{tmn} = 0 \qquad \text{for all } t \in \mathcal{T}, m \notin \mathcal{M}(t) \text{ and } n \in \mathcal{N}, \tag{7.5}$$

$$a_{tm1} = a_{tm2} = 0 \qquad \text{for all } t \in \mathcal{T}' \text{ and for all } m \in \mathcal{M}(t), \tag{7.6}$$

$$a_{tmn} \in \mathbb{N}_0^{TMN} \tag{7.7}$$

Constraint (7.3) requires the number of bookings for each patient type to be less than or equal to the number of requests waiting to be booked. As we discussed earlier, in this model postponement of booking decision is allowed, but only between the first and second decision epoch of the same day. In Constraint (7.4), d_t describes the duration of a type t scan in number of time slots. Constraint (7.4) therefore restricts the total number of slots booked on MRI scanner m during session n of the booking horizon to be less than or equal to the number of available regular slots plus the overtime slots. Constraint (7.5) ensures patients are only booked on an MRI scanner suitable for their type of MRI scan. Recall that we defined the set $\mathcal{T}'$ to contain all outpatient patient types. We do not allow outpatient demand that does not have the priority to be scanned today to be booked into one of today's remaining sessions, because these patients are not necessarily already present in the hospital. Hence we have Constraint (7.6). Finally, all action variables are integer and non-negative (captured by Constraint (4.7)).

Similarly, to be valid, any action in $\mathbf{s}_2$ must satisfy the following constraints:

$$\sum_{n\in\mathcal{N}} \sum_{m\in\mathcal{M}(t)} a_{tmn} = y_t, \qquad \text{for all } t \in \mathcal{T}, \tag{7.8}$$

$$\sum_{t\in\mathcal{T}(m)} x_{tmn} + d_t a_{tmn} \leq C^{R,2}_{mn} + C^{OT,2}_{mn}, \qquad \text{for all } m \in \mathcal{M} \text{ and } n \in \mathcal{N}, \tag{7.9}$$

$$a_{tmn} = 0 \qquad \text{for all } t \in \mathcal{T}, m \notin \mathcal{M}(t) \text{ and } n \in \mathcal{N}, \tag{7.10}$$

$$a_{tm1} = 0 \qquad \text{for all } t \in \mathcal{T}' \text{ and for all } m \in \mathcal{M}(t), \tag{7.11}$$

$$a_{tmn} \in \mathbb{N}_0^{TMN} \tag{7.12}$$

We define the action sets $\mathcal{A}(\mathbf{s}_1)$, for any given state $\mathbf{s}_1 = (1, \mathbf{x}, \mathbf{y}) \in \mathcal{S}_1$, as the set of actions $\mathbf{a}$ satisfying Equations (7.3) to (7.7). We define the action sets $\mathcal{A}(\mathbf{s}_2)$, for any given state $\mathbf{s}_2 = (2, \mathbf{x}, \mathbf{y}) \in \mathcal{S}_2$, as the set of actions $\mathbf{a}$ satisfying Equations (7.8) to (7.12). In general we let $\mathcal{A}(s)$ denote the set of all feasible actions in state $\mathbf{s} \in \mathcal{S}$.

7.1.3 The transition probabilities

First consider the situation in which the system is at the first decision epoch of a given day. That is, the system is in some state $\mathbf{s}_1 = (1, \mathbf{x}, \mathbf{y}) \in \mathcal{S}_1$. Once a decision is made in this state, the system moves to a state corresponding to the second decision epoch of the given day, i.e., the system moves to some state $\mathbf{s}_2 = (2, \mathbf{x}', \mathbf{y}') \in \mathcal{S}_2$. The only source of uncertainty in this transition from $\mathbf{s}_1$ to $\mathbf{s}_2$ is the number of new appointment requests of each patient type. The other parameters are updated based on the capacity allocation to booked patients. Thus, as a result of choosing booking action $\mathbf{a}$ in state $\mathbf{s}_1 = (1, \mathbf{x}, \mathbf{y})$, and having q_t new requests of type t, the state of the system the next decision epoch, denoted by $\mathbf{s}_2 = (2, \mathbf{x}', \mathbf{y}')$, will be determined by the following probability distribution:

$$p(\mathbf{s}_2|\mathbf{s}_1, \mathbf{a}) = \begin{cases} p_2(\mathbf{q}) = \prod_{t\in\mathcal{T}} p_2(q_t), & \text{if } \mathbf{s}_2 = (2, \mathbf{x}', \mathbf{y}') \text{ satisfies Eqs. (7.14) and (7.15)}, \\ 0 & \text{otherwise.} \end{cases} \tag{7.13}$$

$$x'_{tmn} = \begin{cases} x_{tm,n+1} + \sum_{t\in\mathcal{T}(m)} d_t a_{tm,n+1}, & \text{for all } m \in \mathcal{M} \text{ and } n = 1, 2, \ldots, N-1, \\ 0, & \text{for all } m \in \mathcal{M} \text{ and } n = N, \end{cases} \tag{7.14}$$

$$y'_t = q_t + y_t - \sum_{n\in\mathcal{N}} \sum_{m\in\mathcal{M}} a_{tmn}, \quad \text{for all } t \in \mathcal{T}. \tag{7.15}$$

Note that these specific system dynamics differ from the model in Chapter 4, described by (4.13) and 4.14.
Similarly, if the system is at the second decision epoch on a given day, the system's next state is at the first decision epoch in the successive day. Thus, as a result of choosing booking action $\mathbf{a}$ in state $\mathbf{s}_2 = (2, \mathbf{x}, \mathbf{y})$, and having q_t new requests of type t, the state of the system the next decision epoch, denoted by $\mathbf{s}_1 = (1, \mathbf{x}', \mathbf{y}')$, will be determined

by the following probability distribution:

$$p(\mathbf{s}_1|\mathbf{s}_2,\mathbf{a}) = \begin{cases} p_1(\mathbf{q}) = \prod_{t\in\mathcal{T}} p_1(q_t), & \text{if } \mathbf{s}_1 = (1,\mathbf{x}',\mathbf{y}') \text{ satisfies Eqs. (7.17) and (7.18),} \\ 0 & \text{otherwise.} \end{cases} \tag{7.16}$$

$$x'_{tmn} = \begin{cases} x_{tm,n+2} + \sum_{t\in\mathcal{T}(m)} d_t a_{tm,n+2}, & \text{for all } m \in \mathcal{M} \text{ and } n = 1,2,\ldots,N-2, \\ 0, & \text{for all } m \in \mathcal{M} \text{ and } n = N-1, N, \end{cases} \tag{7.17}$$

$$y'_t = q_t, \quad \text{for all } t \in \mathcal{T}. \tag{7.18}$$

7.1.4 The direct costs

The direct costs associated with state-action pair $\mathbf{s} \in \mathcal{S}$, $\mathbf{a} \in \mathcal{A}(\mathbf{s})$ derives from three sources: (i) penalties associated with booking patients beyond their priority-specific access time targets; (ii) a cost associated with the day of service, expressed by overtime cost; and (iii) penalties associ- ated with delaying the booking decisions for some of the waiting demand. . We write the costs as

$$c(\mathbf{s},\mathbf{a}) = \sum_{n\in\mathcal{N}}\sum_{t\in\mathcal{T}} f^{AT}(k,t,n)\left[\sum_{m\in\mathcal{M}(t)} a_{tmn}\right] + f^{AS}(\mathbf{s},\mathbf{a}) \tag{7.19}$$

$$+ \sum_t f^D(t)\left(y_t - \sum_{m,n} a_{tmn}\right). \tag{7.20}$$

In Equation (7.20), $f^{AT}(k,t,n)$ is the cost associated with booking a type t patient into session n of the booking horizon, at decision epoch k of a day. The cost associated with the day of service is $f^{AS}(\mathbf{s},\mathbf{a})$ and $f^D(t)$ is the penalty associated with delaying the booking of a type t patient from the first to the second decision epoch of the day 1.

It is clearly reasonable to assume that $f^{AT}(t,n)$ should be zero if $n \leq AT(t)$. Furthermore, it would seem advisable to ensure that the penalty associated with delaying the booking of a patient, if possible, to the next decision epoch and then booking him/her within the corresponding wait time target should be equal to the penalty associated with booking the patient 1 session late initially. Thus, a natural form for the access time penalty is

$$f^{AT}(k,t,n) = \begin{cases} \sum_{j=1}^{n-AT(k,t)} \gamma^{j-1} f^D(t) & \text{for all } n > AT(t) \\ 0 & \text{otherwise.} \end{cases}$$

The (direct) within-day costs are defined as a cost for overtime and because we *jump* over one afternoon session if we *jump* from $\mathbf{s}_1 = (1,\mathbf{x},\mathbf{y})$ to $\mathbf{s}_2 = (2,\mathbf{x}',\mathbf{y}')$, the direct within-day costs for this session can be computed as

$$f^{AS}(\mathbf{s}_1, \mathbf{a}) = \sum_{m\in\mathcal{M}} h \left[x_{m1} + \sum_{t\in\mathcal{T}(m)} d_t a_{tm1} - C^{R,1}_{m1} \right]^+ ,$$

where h is the overtime cost per time slot and $[z]^+ = \max\{0, z\}$. Note that the overtime and cost are identical for all the MRI scanners $m \in \mathcal{M}$.

When we *jump* from $\mathbf{s}_2 = (1, \mathbf{x}, \mathbf{y})$ to $\mathbf{s}_1 = (2, \mathbf{x}', \mathbf{y}')$, we *jump* over two sessions: the evening session and the morning session of the upcoming day. Hence the direct within-day costs for this *jump* can be computed as

$$f^{AS}(\mathbf{s}_2, \mathbf{a}) = \sum_{m\in\mathcal{M}} h \left[x_{m1} + \sum_{t\in\mathcal{T}(m)} d_t a_{tm1} - C^{R,2}_{m1} \right]^+ + h \left[x_{m2} + \sum_{t\in\mathcal{T}(m)} d_t a_{tm2} - C^{R,2}_{m2} \right]^+ .$$

7.2 Solution Approach

We can use the same *approximate linear programming approach to ADP* as in Chapter 4. That means that we use an affine architecture to approximate the value function $v(\mathbf{s})$ of the MDP. We suggest to use the following affine approximation

$$v(\mathbf{s}) = v(k, \mathbf{x}, \mathbf{y}) = \begin{cases} V_0 + \sum_{t=1}^{T} \sum_{m=1}^{M} \sum_{n=1}^{N} V_{tmn} x_{tmn}, & \text{if } k = 1, \\ \\ \widetilde{V}_0 + \sum_{t=1}^{T} \sum_{m=1}^{M} \sum_{n=1}^{N} \widetilde{V}_{tmn} x_{tmn} + \sum_{t=1}^{T} \widetilde{W}_t y_t & \text{if } k = 2. \end{cases} \tag{7.21}$$

Note that this affine approximation uses more variables than the one in Chapter 4, given by (4.24). However, it enables us eventually to distinguish preferences between booking a type t at MRI scanner m in session n and booking a type t' at MRI scanner m in session n. Since we can not postpone booking decisions from the second decision epoch today to the first decision epoch tomorrow, we would like to study the possibility to leave out the term $\sum_{t\in\mathcal{T}} W_t y_t$ in the approximation. After doing the mathematical operations, we will end up again with a linear program with a reasonable number of variables and constraints for every state-action combination, which dual we hopefully (depending on the actual problem size) can solve with column generation.

-

Bibliography

[1] A. Ahmadi-Javid, Z. Jalali and K. J. Klassen. Outpatient appointment systems in healthcare: A review of optimization studies. *European Journal of Operational Research*, 258(1): 3–34, 2017. DOI: h10.1016/j.ejor.2016.06.064.

[2] N. T. J. Bailey. A Study of Queues and Appointment Systems in Hospital Out-Patient Departments, with Special Reference to Waiting-Times. *Journal of the Royal Statistical Society: Series B (Methodological)*, 14(2): 185–199, 1952. DOI: 10.1111/j.2517-6161.1952.tb00112.x.

[3] M. A. Begen and M. Queyranne. Appointment Scheduling with Discrete Random Durations. *Mathematics of Operations Research*, 36(2): 240–257, 2011. DOI: 10.1287/moor.1110.0489.

[4] D. Bertsekas and J. Tsitsiklis. Neuro-Dynamic Programming. Belmont, MA: Athena Scientific, 1996.

[5] D. Bertsimas and J. N. Tsitsiklis. Introduction to linear optimization. vol. 6. Athena Scientific Belmont, MA, 1997.

[6] I. A. Bikker. Organizing timely treatment in multi-disciplinary care. PhD thesis. University of Twente, 2018. DOI: 10.3990/1.9789402811896.

[7] A. Braaksma. Timely and Efficient planning of Treatments through Intelligent Scheduling. PhD thesis. University of Twente, 2015. Available at: http://www.utwente.nl/en/choir/research/Publications/Publications_Thesis. Last visited on September 24, 2019.

[8] M. W. Carter. Diagnosis: mismanagement of resources. *OR/MS Today*, 29(2): 26–32, 2002.

[9] T. Cayirli and E. Veral. Outpatient scheduling in health care: a review of literature. *Production and Operations Management*, 12(4): 519–549, 2003. DOI: 10.1111/j.1937-5956.2003.tb00218.x.

[10] S. Creemers, J. Beliën and M. Lambrecht. The optimal allocation of server time slots over different classes of patients. *European Journal of Operational Research*, 219(3): 508–521, 2012. DOI: 10.1016/j.ejor.2011.10.045.

[11] F. d'Epenoux. A Probabilistic Production and Inventory Problem. *Management Science*, 10(1): 98–108, 1963. DOI: 10.1287/mnsc.10.1.98.

[12] J. W. Daniel. Splines and efficiency in dynamic programming. *Journal of Mathematical Analysis and Applications*, 54(2): 402–407, 1976. DOI: 10.1016/0022-247X(76)90209-2.

[13] C. Derman. Finite State Markovian Decision Processes. Orlando, FL, USA: Academic Press, Inc., 1970. ISBN: 0122092503.

[14] A. Erdelyi and H. Topaloglu. Computing protection level policies for dynamic capacity allocation problems by using stochastic approximation methods. *IIE Transactions*, 41(6): 498–510, 2009. DOI: 10.1080/07408170802706543.

[15] A. Erdelyi and H. Topaloglu. Approximate dynamic programming for dynamic capacity allocation with multiple priority levels. *IIE Transactions*, 43(2): 129–142, 2010. DOI: 10.1080/0740817X.2010.504690.

[16] N. Geng and X. Xie. Optimizing contracted resource capacity with two advance cancelation modes. *European Journal of Operational Research*, 221(3): 501–512, 2012. DOI: 10.1016/j.ejor.2012.04.017.

[17] Y. Gocgun and M. Puterman. Dynamic scheduling with due dates and time windows: an application to chemotherapy patient appointment booking. *Health Care Management Science*, 17: 60–76, 2014. DOI: 10.1007/s10729-013-9253-z.

[18] D. Gupta and B. Denton. Appointment scheduling in health care: Challenges and opportunities. *IIE Transactions*, 40(9): 800–819, 2008. DOI: 10.1080/07408170802165880.

[19] S. Hahn-Goldberg, M. W. Carter, J. C. Beck, M. Trudeau, P. Sousa and K. Beattie. Dynamic optimization of chemotherapy outpatient scheduling with uncertainty. *Health Care Management Science*, 17: 379–392, 2014. DOI: 10.1007/s10729-014-9268-0.

[20] E. W. Hans, M. Van Houdenhoven and P. J. Hulshof. A framework for healthcare planning and control. in: *Handbook of healthcare system scheduling (p. 303-320)*. Springer, 2012.

[21] P. J. H. Hulshof, N. Kortbeek, R. J. Boucherie, E. W. Hans and P. J. M. Bakker. Taxonomic classification of planning decisions in health care: a structured review of the state of the art in OR/MS. *Health Systems*, 1(2): 129–175, 2012. DOI: 10.1057/hs.2012.18.

[22] P. J. Hulshof, P. T. Vanberkel, R. J. Boucherie, E. W. Hans, M. van Houdenhoven and J.-K. C. W. van Ommeren. Analytical models to determine room requirements in outpatient clinics. *OR Spectrum*, 34(2"): 391–405, 2012. DOI: 10.1007/s00291-012-0287-2.

[23] L. Kallenberg. Linear Programming and Finite Markovian Control Problems. Amsterdam.: Mathematisch Centrum, 1983.

[24] J. Marynissen and E. Demeulemeester. Literature review on multi-appointment scheduling problems in hospitals. *European Journal of Operational Research*, 272(2): 407–419, 2019. DOI: 10.1016/j.ejor.2018.03.001.

[25] M. Murray and D. M. Berwick. Advanced Access: Reducing Waiting and Delays in Primary Care. *JAMA*, 289(8): 1035–1040, 2003. DOI: 10.1001/jama.289.8.1035.

[26] M. S. Parizi and A. Ghate. Multi-class, multi-resource advance scheduling with no-shows, cancellations and overbooking. *Computers & Operations Research*, 67: 90–101, 2016. DOI: 10.1016/j.cor.2015.09.004.

[27] J. Patrick and M. Puterman. Reducing wait times through operations research: optimizing the use of surge capacity. *Healthcare quarterly (Toronto, Ont.)* 11(3): 77–83, 2008. DOI: 10.12927/hcq.2013.19844.

[28] J. Patrick, M. L. Puterman and M. Queyranne. Dynamic Multipriority Patient Scheduling for a Diagnostic Resource. *Operations Research*, 56(6): 1507–1525, 2008. DOI: 10.1287/opre.1080.0590.

[29] W. Powell. Approximate Dynamic Programming: Solving the Curses of Dimensionality. Hoboken, NJ: Wiley-Interscience, 2007.

[30] M. L. Puterman. Markov Decision Processes.: Discrete Stochastic Dynamic Programming. John Wiley & Sons, 2014.

[31] Rijnstate. Rijnstate jaarmagzine 2017-2018 [Patient information brochure]. Sept. 2018. Available at: http://www.rijnstate.nl/media/628384/rijnstate-jaarmagazine-2018.pdf. Brochure is in Dutch. Last visited on September 24, 2019.

[32] C. Sanmartin, F. Gendron, J.-M. Berthelot and K. Murphy. Access to health care services in Canada, 2003. *Ottawa: Statistics Canada, Catalogue 82-575-XIE*, 2004.

[33] A. Sauré, M. A. Begen and J. Patrick. Dynamic multi-priority, multi-class patient scheduling with stochastic service times. *European Journal of Operational Research*, 280(1): 254–265, 2020. DOI: 10.1016/j.ejor.2019.06.040.

[34] A. Sauré, J. Patrick, S. Tyldesley and M. L. Puterman. Dynamic multi-appointment patient scheduling for radiation therapy. *European Journal of Operational Research*, 223(2): 573–584, 2012. DOI: 10.1016/j.ejor.2012.06.046.

[35] H.-J. Schütz and R. Kolisch. Approximate dynamic programming for capacity allocation in the service industry. *European Journal of Operational Research*, 218(1): 239–250, 2012. DOI: 10.1016/j.ejor.2011.09.007.

[36] P. J. Schweitzer and A. Seidmann. Generalized polynomial approximations in Markovian decision processes. *Journal of Mathematical Analysis and Applications*, 110(2): 568–582, 1985. DOI: 10.1016/0022-247X(85)90317-8.

[37] I. B. Vermeulen, S. M. Bohte, S. G. Elkhuizen, H. Lameris, P. J. Bakker and H. L. Poutré. Adaptive resource allocation for efficient patient scheduling. *Artificial Intelligence in Medicine*, 46(1): 67–80, 2009. DOI: 10.1016/j.artmed.2008.07.019.

A

Configurations of MRI examinations carried out in Rijnstate

Table A.1: HIX codes of the MRI examinations carried out in Rijnstate. The third column shows the service time, in minutes, that is reserved for each MRI according to the currently used blueprint schedule. The last columns show whether a MRI scanner is suitable for the type of MRI examination (✓) or not (✗).

HIX code	Description	Duration	Compatibility MRI scanner ...		
			1	2	3
1390	MRI of the brain	20	✓	✓	✓
1390A	Carotid artery MRI	30	✓	✓	✓
1390C	MRI of the cholesteatoma	30	✓	✗	✓
1390E	MRI to see if there is an (obvious) reason for epileptic seizures	30	✓	✓	✓
1390H	MRI of the cranial nerves	30	✓	✗	✗
1390K	MRI of the brain under sedation	30	✓	✗	✗
2090	MRI of the face and/or neck	30	✓	✓	✓
2090A	MRI of the neck	30	✓	✓	✗
2290	MRI of the orbita	30	✓	✓	✗
2990	MRI of the thyroid	30	✓	✓	✗
3090T	MRI of the spinal column	30	✗	✓	✓
3190	MRI of the cervical spine (cervical vertebrae)	30	✗	✓	✓
3190PB	MRI of the brachial plexus (node of the nerves above the collarbone)	30	✓	✓	✓

(Continued on next page)

Table A.1 – *Continued*

HIX code	Description	Duration	Compatibility MRI scanner ... 1	2	3
3290	MRI of the thoracic spine (thoracic vertebrae)	30	✗	✓	✓
3390	MRI of the lumbar spine (lumbar vertebrae)	30	✗	✓	✓
3390P	MRI of the plexus lumbosacralis (spinal nerves in the lower part of the back)	30	✗	✓	✓
3490	MRI of the sacroiliac joints/the sacrum	30	✗	✓	✓
3615	MRI with injection in joint(s) (usually only done in combi with another MRI type)	10	✓	✓	✓
3690	MRI of myelum (spinal cord)	30	✗	✓	✗
4090 L/R	MRI of the upper extremities (shoulder, elbow, hand, wrist); left or right	30	✓	✓	✗
4090 AR/AL	MRI of hand/wrist combined with an arthrogram; left or right	30	✓	✗	✗
4292 L/R	MRI of shoulder combined with an arthrogram; left or right	30	✓	✗	✗
5090	MRI of the thorax (chest)	30	✗	✓	✓
5190/5192	MRI of the heart (5192 is with intake of Dobutamine)	60	✗	✓	✗
5191/5191R	MRI of heart with intake of Androsine/MRI of heart with intake of Regadenoson	60	✗	✓	✗
6990	MRI of the breasts (MRI Mammae)	30	✓	✗	✗
6990P	MRI of the breasts to take a biopsy (MRI vacuum biopsy Mammae)	60	✗	✓	✗
7090	MRI of the abdomen	30	✗	✓	✓
7090A	MRA (Magnetic Resonance Angiography): examination of blood vessel(s)/vascular system	30	✗	✓	✗
7091	CINE MRI of the abdomen; dynamic MRI scan to detect adhesions	20	✗	✗	✓
7490	MRI of the small intestine	30	✗	✓	✓
7690	MRI of the liver/pancreas	30	✗	✓	✓
7690P	MRI of the liver with intake of contrast medium Primovist	40	✗	✓	✗
7790	MRI of the bile duct/pancreas	30	✗	✓	✓

(Continued on next page)

Table A.1 – *Continued*

HIX code	Description	Duration	Compatibility MRI scanner ... 1	2	3
8490	MRI of the prostate	30	✓	✓	✓
9090 L/R	MRI of lower extremities (ankle, foot); left or right	30	✓	✓	✓
9090 AL/AR	MRI of hip(s) combined with an arthrogram; left or right	30	✓	✗	✗
9090A	MRI scan of the basin/pelvis and/or the upper legs (femur)	30	✗	✓	✗
9190	MRI of the lower abdomen	30	✓	✓	✓
9190E	MRI for determining or excluding endometriosis	30	✗	✗	✓
9290	MRI of the pelvis and/or hips	30	✓	✓	✓
9491 L/R	MRI of the knee; left or right	20	✓	✓	✓
9491 EL/ER	MRI of the knee to study a (nerve) entrapment	30	✗	✓	✗
ARTI	MRI scan of the brain (*Prostate RTG*)	30	✓	✓	✗
ARTIP	MRI scan of the prostate (*Prostate RTG*)	30	✗	✓	✗
VITE	MRI scan football players of Vitesse	30	✓	✓	✓
MRV	MRV (Magnetic Resonance Venography): examination of blood vessel(s)/vascular system	30	✗	✓	✗
MWEKE	MRI of soft tissue	30	✓	✓	✓
The HIX codes below correspond to MRI examinations that are booked decentrally.					
1390D	MRI of the brain; booked decentrally, i.e., directly by another outpatient clinic	20	✓	✓	✓
3190D	MRI of the cervical spine; booked decentrally, i.e., directly by another outpatient clinic	30	✗	✓	✓
3390D	MRI of the lumbar spine; booked decentrally, i.e., directly by another outpatient clinic	30	✗	✓	✓
The HIX codes below correspond to MRI examinations for one-stop-shop patients.					
AKP L/R	MRI of the knee for one-stop-shop patients; left or right	20	✓	✓	✓

(Continued on next page)

Table A.1 – *Continued*

HIX code	Description	Duration	Compatibility MRI scanner ... 1	2	3
DAG	MRI for one-stop-shop patients of the geriatrics department	20	✓	✓	✗
TIA/TIAS	MRI for one-stop-patients with suspicion/symptoms of a TIA; location Arnhem	20	✓	✓	✗
TIAZ	MRI for one-stop-patients with suspicion/symptoms of a TIA; location Zevenaar	20	✗	✗	✓
LWK	MRI of the lumbar spine (lumbar vertebrae) for one-stop-shop patients; location Zevenaar	20	✗	✗	✓
CWK	MRI of the cervical spine (cervical vertebrae) for one-stop-shop patients; location Zevenaar	20	✗	✗	✓
NEURO	MRI of the brain for one-stop-shop patients	20	✓	✓	✓
6990M	MRI of the breasts (MRI Mammae) for one-stop-shop patients	30	✓	✗	✗

B

The second ILP and constraint-related computations in the column generation algorithm

B.1 Linearization of the optimization problem in (4.31)

$$
\begin{aligned}
\min_{\substack{\mathbf{x},\mathbf{y},\mathbf{a},\\ \boldsymbol{\phi},\boldsymbol{\theta},\boldsymbol{\xi},\boldsymbol{\kappa}}} \quad & \sum_{n=1}^{N}\sum_{t=1}^{T}\sum_{m\in\mathcal{M}(t)} f^{AT}(2,t,n)a_{tmn} + \sum_{n=1}^{N}\sum_{m=1}^{M} h\left[\theta_{mn} - \kappa_{mn} + (\xi_{mn} - \phi_{mn})C_{mn}^{R,2}\right] \\
& - \sum_{m=1}^{M}\left\{\sum_{n=1}^{N}\widetilde{V}_{mn}x_{mn} - \gamma\sum_{n=1}^{N-2}V_{mn}(x_{m,n+2} + \sum_{t\in\mathcal{T}(m)} d_t a_{tm,n+2})\right\} - \sum_{t=1}^{T}\widetilde{W}_t y_t, && \text{(B.1a)} \\
\text{s.t.} \quad & x_{mn} \le C_{mn}^{R,2} + C_{mn}^{OT,2}, \qquad \forall m\in\mathcal{M},\ \forall n = 1,2,\ldots,N-1, && \text{(B.1b)} \\
& x_{mn} = 0, \qquad \forall m\in\mathcal{M},\ n = N, && \text{(B.1c)} \\
& y_t \le Q_t^2, \qquad \forall t\in\mathcal{T}, && \text{(B.1d)} \\
& \sum_{n\in\mathcal{N}}\sum_{m\in\mathcal{M}(t)} a_{tmn} = y_t, \qquad \forall t\in\mathcal{T}, && \text{(B.1e)} \\
& x_{mn} + \sum_{t\in\mathcal{T}(m)} d_t a_{tmn} \le C_{mn}^{R,2} + C_{mn}^{OT,2}, \qquad \forall m\in\mathcal{M},\ n\in\mathcal{N}, && \text{(B.1f)} \\
& a_{tmn} = 0, \qquad \forall t\in\mathcal{T},\ m\notin\mathcal{M}(t),\ n\in\mathcal{N}, && \text{(B.1g)} \\
& a_{tm1} == 0, \qquad \forall t\in\mathcal{T}',\ m\in\mathcal{M}(t), && \text{(B.1h)} \\
& x_{mn} + \sum_{t\in\mathcal{T}(m)} d_t a_{tmn} \ge \phi_{mn}C_{mn}^{R,2}, \qquad \forall m\in\mathcal{M},\ \forall n\in\mathcal{N} && \text{(B.1i)} \\
& x_{mn} + \sum_{t\in\mathcal{T}(m)} d_t a_{tmn} \le (1-\phi_{mn})C_{mn}^{R,2} + \phi_{mn}(C_{mn}^{R,2} + C_{mn}^{OT,2}), \quad \forall m\in\mathcal{M},\ n\in\mathcal{N}, && \text{(B.1j)}
\end{aligned}
$$

$$x_{mn} \geq \xi_{mn} C_{mn}^{R,2}, \qquad \forall m \in \mathcal{M},\ n \in \mathcal{N}, \tag{B.1k}$$

$$x_{mn} \leq (1-\xi_{mn}) C_{mn}^{R,2} + \xi_{mn}(C_{mn}^{R,2} + C_{mn}^{OT,2}), \quad \forall m \in \mathcal{M},\ n \in \mathcal{N}, \tag{B.1l}$$

$$\theta_{mn} \leq \phi_{mn}(C_{m,1}^{R,2} + C_{m,1}^{OT,2}), \qquad \forall m \in \mathcal{M},\ n \in \mathcal{N}, \tag{B.1m}$$

$$\theta_{mn} \leq x_{mn} + \sum_{t \in \mathcal{T}(m)} d_t a_{tmn}, \qquad \forall m \in \mathcal{M},\ n \in \mathcal{N}, \tag{B.1n}$$

$$\theta_{mn} \geq x_{mn} + \sum_{t \in \mathcal{T}(m)} d_t a_{tmn} - (C_{mn}^{R,2} + C_{mn}^{OT,2})(1-\phi_{mn}), \qquad \forall m \in \mathcal{M},\ n \in \mathcal{N}, \tag{B.1o}$$

$$\kappa_{mn} \leq \xi_{mn}(C_{mn}^{R,2} + C_{mn}^{OT,2}), \qquad \forall m \in \mathcal{M},\ n \in \mathcal{N}, \tag{B.1p}$$

$$\kappa_{mn} \leq x_{mn}, \qquad \forall m \in \mathcal{M},\ n \in \mathcal{N}, \tag{B.1q}$$

$$\kappa_{mn} \geq x_{mn} - (C_{mn}^{R,2} + C_{mn}^{OT,2})(1-\xi_{mn}), \qquad \forall m \in \mathcal{M},\ n \in \mathcal{N}, \tag{B.1r}$$

$$\mathbf{x} \in \mathbb{N}_0^{MN},\ \mathbf{y} \in \mathbb{N}_0^{T},\ \mathbf{a} \in \mathbb{N}_+^{TMN}, \tag{B.1s}$$

$$\boldsymbol{\phi} \in \{0,1\}^{MN},\ \boldsymbol{\xi} \in \{0,1\}^{MN}, \boldsymbol{\theta} \in \mathbb{N}_0^{MN}, \boldsymbol{\kappa} \in \mathbb{N}_0^{MN}. \tag{B.1t}$$

B.2 Calculation of the right-hand side of the constraints in the (restricted) master problem

For the constraints in the (restricted) master problem in (4.28a)-(4.28h) we need the quantities

$$\mathbb{E}_\alpha[x_{mn}] = \sum_{\mathbf{s_1} \in \mathcal{S}_1} \alpha(\mathbf{s}_1) x_{mn}(\mathbf{s}_1), \qquad \mathbb{E}_\alpha[y_t] = \sum_{\mathbf{s}_1 \in \mathcal{S}_1} \alpha(\mathbf{s}_1) y_t(\mathbf{s}_1),$$
$$\mathbb{E}_\beta[x_{mn}] = \sum_{\mathbf{s}_2 \in \mathcal{S}_2} \beta(\mathbf{s}_2) x_{mn}(\mathbf{s}_2), \qquad \mathbb{E}_\beta[y_t] = \sum_{\mathbf{s}_2 \in \mathcal{S}_2} \beta(\mathbf{s}_2) y_t(\mathbf{s}_2).$$

Experiments suggests that the larger these quantities are, the shorter the calculation time of the column generation becomes after feasible initial state-action pairs are found. Contradictory to this is that the smaller these quantities are, the less iterations are needed in the described procedure to find feasible initial state-action pairs (see Section 4.2.2). All together, we try to find values for $\mathbb{E}_\alpha[x_{mn}]$, $\mathbb{E}_\beta[x_{mn}]$, $\mathbb{E}_\alpha[y_t]$, $\mathbb{E}_\beta[y_t]$ as large as possible that still makes sure that finding feasible initial state-action pairs takes not too long.

We will show that for any $\epsilon > 0$, we can define the right-hand sides of the (restricted) master problem in (4.28a)-(4.28h) as

$$\mathbb{E}_\alpha[x_{mn}] = \epsilon \cdot \frac{1 + C_{mn}^{R,1} + C_{mn}^{OT,1}}{2}, \forall n = 1, \ldots, N-2, \qquad \mathbb{E}_\alpha[x_{mn}] = 0, \text{ for } n = N-1, N,$$

$$\mathbb{E}_\alpha[y_t] = \epsilon \cdot \frac{1 + Q_t^1}{2},$$

$$\mathbb{E}_\beta[x_{mn}] = \epsilon \cdot \frac{1 + C_{mn}^{R,2} + C_{mn}^{OT,2}}{2}, \forall n = 1, \ldots, N-1, \qquad \mathbb{E}_\beta[x_{mn}] = 0, \text{ for } n = N,$$

$$\mathbb{E}_{\beta}[y_t] = \epsilon \cdot \frac{1+Q_t^2}{2}.$$

To show this, we define support set $\Omega = \{0, 1, 2, \dots, K\}$ for the random variable X, defined as $X(\omega) = \omega$. Furthermore, we let Y be a (discrete) uniformly distributed random variable over the support set $\Omega \backslash \{0\}$. Finally, X has probability mass function p_X defined as

$$p_X(X=0) = 1-\epsilon, \quad p_X(X=x) = \frac{\epsilon}{|\Omega|-1} = \frac{\epsilon}{K}, \text{ for all } x \neq 0.$$

Then the following equality holds

$$\mathbb{E}[X] = 0 \cdot (1-\epsilon) + \sum_{k=1}^{K} k \cdot \frac{\epsilon}{K} = \frac{\epsilon}{K}\frac{K(K+1)}{2} = \epsilon \frac{K+1}{2} = \epsilon \mathbb{E}[Y].$$

Next, we define the random probability distribution $\boldsymbol{\alpha}$ over the states $\mathbf{s}_1 \in \mathcal{S}_1$ by $\alpha(\mathbf{0}) = 1-\epsilon$ and $\alpha(\mathbf{s}_1) = \epsilon/(|\mathcal{S}_1|-1)$, for all $\mathbf{s}_1 \neq \mathbf{0}$. It then follows from the preceding that, for all $m \in \mathcal{M}, n = 1, \dots, N-2$, $\mathbb{E}_{\alpha}[x_{mn}] = \epsilon \mathbb{E}_{\nu}[x_{mn}(\mathbf{s}_1)]$, where ν is the uniform distribution over the support set $\mathcal{S}_1 \backslash \{\mathbf{0}\}$. It now suffices to observe that

$$\mathbb{E}_{\nu}[x_{mn}(\mathbf{s}_1)] = \frac{1 + C_{mn}^{R,1} + C_{mn}^{OT,1}}{2},$$

because ν is the (discrete) uniform distribution over the support set $\mathcal{S}_1 \backslash \{\mathbf{0}\}$. Hence the expectation $\mathbb{E}_{\nu}[x_{mn}(\mathbf{s}_1)]$ reduces to the expectation of a uniformly distributed random variable over the support set $\{1, \dots, C_{mn}^{R,1} + C_{mn}^{OT,1}\}$. Hence, as this derivations shows, if take $\boldsymbol{\alpha}$ as the uniform distribution, it turns out that $\mathbb{E}_{\alpha}[x_{mn}]$ is the uniform distribution over its support set $\{0, 1, 2, \dots, C_{mn}^{R,1} + C_{mn}^{OT,1}\}$. Similar results hold for the other expectation appearing in the right-hand sides of the (4.28a)-(4.28h). Hence, we can calculate these as above.
We used $\epsilon = 0.5$.

C

Rijnstate's patient types and blueprint calendars used in the MDP and simulation

C.1 Patient types definition based on the blueprint calendars for all three MRI scanners used in Rijnstate

Table C.1: Patient types definition and some attributes. The second column shows the HIX codes belonging the patient type. Note that a HIX code may belong to multiple patient types. In that case the difference is in the access time target (abbreviated as AT in the header). Today is day 0 and tomorrow is counted as day 1. The third column shows the service time, in minutes, that is reserved for each patient type. The last columns shows whether a MRI scanner is suitable for the type of MRI examination (✓) or not (✗).

Patient type	HIX code(s)	AT (in days)	Duration (in min.)	Compatibility MRI scanner ...		
				1	2	3
1	emergency; machine 1 & 2: 1390, 1390D, 1390A, 1390E, 2090, 2090A, 2990, 3190PB, 8490, 9090 L/R, 9190, 9290, 9491 L/R, ARTI, MWEKE	0	30	✓	✓	✗
2	emergency; only on machine 1: 1390C, 1390H, 1390K, 2290, 4090 L/R, 4090 AL/AR, 4292 L/R, 6990	0	30	✓	✗	✗
3	emergency; only on machine 2: 3090T, 3190, 3190D, 3290, 3390, 3390D, 3390P, 3490, 3690, 5090, 6990P, 7090, 7090A, 7490, 7690, 7790, 9090A, 9491 EL/ER, MRV	0	30	✗	✓	✗
4	emergency; only on machine 2: 5190/5192, 5191/5191R, 7690P	0	60	✗	✓	✗

(Continued on next page)

Table C.1 – *Continued*

Patient type	HIX code(s)	AT (in days)	Duration (in min.)	Compatibility MRI scanner ...		
				1	2	3
5	emergency; only machine 3: 3090T, 3190, 3290, 3390, 3390P, 3490, 5090, 7090, 7091, 7490, 7690, 7790, 8490, 9090 L/R, 9190, 9190E, 9290, 9491, MWEKE	0	30	✗	✗	✓
6	1390, 1390D	3	20	✓	✓	✓
7	1390A, 2090, 3190PB; 3390P	3	30	✓	✓	✓
8	8490, 9190	3	30	✓	✓	✓
9	2290, 2990	3	30	✓	✓	✗
10	3190, 3190D	3	30	✗	✓	✓
11	3390, 3390D	3	30	✗	✓	✓
12	5090, 7090, 7490, 7690, 7790	3	30	✗	✓	✓
13	3690	3	30	✗	✓	✗
14	7091	3	20	✗	✗	✓
15	4090 L/R; 9090 L/R; 9290; MWEKE	5	30	✓	✓	✓
16	9491 L/R	5	30	✓	✓	✓
17	NEURO	5	20	✓	✓	✓
18	2090A	5	30	✓	✓	✗
19	DAG	5	20	✓	✓	✗
20	TIA/TIAS	5	20	✓	✓	✗
21	1390C	5	20	✓	✗	✓
22	1390H	5	30	✓	✗	✗
23	6990M	5	30	✓	✗	✗
24	3090T, 3290, 3490	5	30	✗	✓	✓
25	7690P	5	40	✗	✓	✗
26	9190E	5	30	✗	✗	✓
27	TIAZ	5	20	✗	✗	✓
28	LWK	5	20	✗	✗	✓
29	CWK	5	20	✗	✗	✓
30	AKP L/R	10	20	✓	✓	✓
31	1390, 1390D	10	20	✓	✓	✓
32	4090 L/R; 9090 L/R; 9290; MWEKE	10	30	✓	✓	✓
33	9491 L/R	10	30	✓	✓	✓
34	1390A, 2090, 3190PB; 3390P	10	30	✓	✓	✓
35	1390E	10	30	✓	✓	✓
36	8490, 9190	10	30	✓	✓	✓
37	2290, 2990	10	30	✓	✓	✗
38	ARTI	10	30	✓	✓	✗

(Continued on next page)

Table C.1 – *Continued*

Patient type	HIX code(s)	AT (in days)	Duration (in min.)	Compatibility MRI scanner ...		
				1	2	3
39	1390-3190 combi; 1390D-3190D combi; 1390-3090T combi; 1390-3690 combi; 3190-3190 combi; 3190-3390 combi; 3190D-3390D combi; 3190-3190PB combi; 3290-3390 combi; 3390-3390P combi; 3090T-3490 combi; 7690-7790 combi	10	40	✓	✓	✗
40	1390K	10	30	✓	✗	✗
41	6990	10	30	✓	✗	✗
42	4090 AR/AL, 9090 AL/AR, 4292 L/R	10	30	✓	✗	✗
43	3190, 3190D	10	30	✗	✓	✓
44	3390, 3390D	10	30	✗	✓	✓
45	3090T, 3290, 3490	10	30	✗	✓	✓
46	5090, 7090, 7490, 7690, 7790	10	30	✗	✓	✓
47	5190/5192, 5191/5191R	10	60	✗	✓	✗
48	6990P	10	30	✗	✓	✗
49	7090A, 9090A, MRV	10	30	✗	✓	✗
50	3690	10	30	✗	✓	✗
51	ARTIP	10	30	✗	✓	✗

C.2 Blueprint calendars used for validation of the simulation model

Monday (Day 1 and Day 6 in the biweekly cycle)

Time			#Slots	Feasible patient types	Remarks
08:00	-	08:30	3	10, 43	
08:30	-	09:30	6	6, 7, 9, 10, 11, 13, 15, 16, 24, 31, 32, 33, 34, 37, 39, 43, 44, 45, 50	
09:30	-	10:10	4	17 (6, 31 / 1)	*one-stop-shop* patient type
10:10	-	10:30	2	20 (6, 31 / 1)	*one-stop-shop* patient type
10:30	-	11:00	3	7, 10, 11, 24, 34, 43, 44, 45	
11:00	-	12:00	6	4, 47	
12:00	-	12:30	3	6, 7, 8, 9, 10, 11, 12, 13, 15, 16, 24, 25, 31, 32, 33, 34, 35, 36, 37, 39, 43, 44, 45, 46, 47, 49, 50, 51	Block for all types feasible on this MRI scanner
12:30	-	13:00	3	1, 3, 4, 6, 7, 8, 9, 10, 11, 12, 13	Emergency block: patient types with acces time targets less than or equal to 3 days can be booked here
13:00	-	13:30	3	7, 10, 11, 24, 34, 43, 44, 45	
13:30	-	16:30	18	4, 47	
16:30	-	17:00	3	6, 7, 8, 9, 10, 11, 12, 13, 15, 16, 24, 25, 31, 32, 33, 34, 35, 36, 37, 39, 43, 44, 45, 46, 47, 49, 50, 51	Block for all types feasible on this MRI scanner
17:00	-	18:00	6	1, 3, 4, 6, 7, 8, 9, 10, 11, 12, 13	Emergency block: patient types with acces time targets less than or equal to 3 days can be booked here

(a) Blueprint calendar for MRI scanner 2 for Day 1 and 6 in the biweekly cycle. This blueprint calendar is based on Rijnstate's blueprint calendar for Mondays in weeks with odd and even week numbers, respectively.

Tuesday (Day 2 (till 18:00) and Day 7 (till 21:00) in the biweekly cycle)

Time			#Slots	Feasible patient types	Remarks
08:00	-	08:30	3	10, 43	
08:30	-	09:00	3	6, 7, 9, 10, 11, 13, 15, 16, 24, 31, 32, 33, 34, 37, 39, 43, 44, 45, 50	
09:00	-	09:40	4	20 (6, 31 / 1)	*one-stop-shop* patient type
09:40	-	12:30	17	4, 47	
12:30	-	15:30	18	4, 47	
15:30	-	16:30	6	1, 3, 4, 6, 7, 8, 9, 10, 11, 12, 13	Emergency block: patient types with acces time targets less than or equal to 3 days can be booked here
16:30	-	17:00	3	8, 9, 12, 13, 15, 32, 37, 46, 50	
17:00	-	17:30	3	1, 3, 4, 6, 7, 8, 9, 10, 11, 12, 13	Emergency block: patient types with acces time targets less than or equal to 3 days can be booked here
17:30	-	18:00	3	51	
18:00	-	18:30	3	11, 24, 44, 45	
18:30	-	19:30	6	8, 9, 12, 13, 15, 32, 37, 46, 50	
19:30	-	20:30	6	7, 10, 11, 24, 34, 43, 44, 45	
20:30	-	21:00	3	6, 10, 11, 24, 31, 39, 43, 44, 45	

(b) Blueprint calendar for MRI scanner 2 for Day 2 and 7 in the biweekly cycle. This blueprint calendar is based on Rijnstate's blueprint calendar for Tuesdays in weeks with odd and even week numbers, respectively. In the weeks with even week numbers, the MRI scanner is operated till 21.00 p.m. In our biweekly cycle that is on Day 7.

Wednesday (Day 3 (till 18:00) and Day 8 (till 21:00) in the biweekly cycle)

Time			#Slots	Feasible patient types	Remarks
08:00	-	08:30	3	15, 32	
08:30	-	09:10	4	30 (16, 33 / 1)	*one-stop-shop* patient type
09:10	-	10:00	5	6, 7, 8, 9, 10, 11, 12, 13, 15, 16, 18, 24, 25, 31, 32, 33, 34, 35, 36, 37, 38, 39, 43, 44, 45, 46, 47, 48, 49, 50, 51	Block for all types feasible on this MRI scanner
10:00	-	10:30	3	48	
10:30	-	11:30	6	18	
11:30	-	12:00	3	38	
12:00	-	12:30	3	51	
12:30	-	13:00	3	6, 7, 9, 31, 34, 35, 37	
13:00	-	14:00	6	1, 3, 4, 6, 7, 8, 9, 10, 11, 12, 13	Emergency block: patient types with acces time targets less than or equal to 3 days can be booked here
14:00	-	14:30	3	38	
14:30	-	15:30	6	15, 32	
15:30	-	16:00	3	6, 7, 8, 9, 10, 11, 12, 13, 15, 16, 18, 24, 25, 31, 32, 33, 34, 35, 36, 37, 38, 39, 43, 44, 45, 46, 47, 48, 49, 50, 51	Block for all types feasible on this MRI scanner
16:00	-	16:30	3	16, 33	
16:30	-	17:00	3	6, 31	
17:00	-	17:30	3	1, 3, 4, 6, 7, 8, 9, 10, 11, 12, 13	Emergency block: patient types with acces time targets less than or equal to 3 days can be booked here
17:30	-	18:00	3	51	
18:00	-	18:30	3	11, 24, 44, 45	
18:30	-	19:30	6	8, 9, 12, 13, 15, 32, 37, 46, 50	
19:30	-	20:30	6	7, 10, 11, 24, 34, 43, 44, 45	
20:30	-	21:00	3	6, 10, 11, 24, 31, 39, 43, 44, 45	

(c) Blueprint calendars for MRI scanner 2 for Day 3 and Day 8 in the biweekly cycle. Blueprints are based on Rijnstate's blueprint calendars for Wednesdays in weeks with odd and even week numbers, respectively. In the weeks with even week numbers, the MRI scanner is operated till 21.00 p.m. In our biweekly cycle that is on Day 8.

Thursday (Day 4 (till 18:00) and Day 9 (till 21:00) in the biweekly cycle)

Time			#Slots	Feasible patient types	Remarks
08:00	-	08:30	3	10, 43	
08:30	-	09:00	3	6, 7, 9, 10, 11, 13, 15, 16, 24, 31, 32, 33, 34, 37, 39, 43, 44, 45, 50	
09:00	-	10:00	6	20 (6, 31 / 1)	*one-stop-shop* patient type
10:00	-	10:30	3	49	
10:30	-	11:30	6	7, 10, 11, 24, 34, 43, 44, 45	
11:30	-	12:30	6	8, 12, 25, 37, 46, 49, 51	
12:30	-	13:10	4	19 (6, 31 / 1)	
13:10	-	14:40	9	1, 3, 4, 6, 7, 8, 9, 10, 11, 12, 13	Emergency block: patient types with acces time targets less than or equal to 3 days can be booked here
14:40	-	15:10	3	6, 7, 8, 9, 10, 11, 12, 13, 15, 16, 18, 24, 25, 31, 32, 33, 34, 35, 36, 37, 38, 39, 43, 44, 45, 46, 47, 48, 49, 50, 51	Block for all types feasible on this MRI scanner
15:10	-	16:10	6	8, 12, 25, 37, 46, 49, 51	
16:10	-	16:40	3	6, 7, 11, 31, 34, 44	
16:40	-	17:00	2	6, 31	
17:00	-	17:30	3	1, 3, 4, 6, 7, 8, 9, 10, 11, 12, 13	Emergency block: patient types with acces time targets less than or equal to 3 days can be booked here
17:30	-	19:30	12	11, 24, 44, 45	
19:30	-	20:30	6	7, 10, 11, 24, 34, 43, 44, 45	
20:30	-	21:00	3	6, 10, 11, 24, 31, 39, 43, 44, 45	

(d) Blueprint calendar for MRI scanner 2 for Day 4 and 9 in the biweekly cycle. This blueprint calendar is based on Rijnstate's blueprint calendar for Thursdays in weeks with odd and even week numbers, respectively. In the weeks with even week numbers, the MRI scanner is operated till 21.00 p.m. In our biweekly cycle that is on Day 9.

Friday (Day 5 and Day 10 in the biweekly cycle)

Time	#Slots	Feasible patient types	Remarks
08:00 - 08:30	3	10, 43	
08:30 - 09:30	6	6, 7, 9, 10, 11, 13, 15, 16, 24, 31, 32, 33, 34, 37, 39, 43, 44, 45, 50	
09:30 - 10:10	4	20 (6, 31 / 1)	*one-stop-shop* patient type
10:10 - 11:10	6	8, 12, 25, 37, 46, 49	
11:10 - 12:30	8	6, 7, 8, 9, 10, 11, 12, 13, 15, 16, 18, 24, 25, 31, 32, 33, 34, 35, 36, 37, 38, 39, 43, 44, 45, 46, 47, 48, 49, 50, 51	Block for all types feasible on this MRI scanner
12:30 - 14:00	9		
14:00 - 15:00	6	8, 12, 25, 37, 46, 49	
15:00 - 15:30	3	1, 3, 4, 6, 7, 8, 9, 10, 11, 12, 13	Emergency block: patient types with acces time targets less than or equal to 3 days can be booked here
15:30 - 17:00	9	6, 7, 8, 9, 10, 11, 12, 13, 15, 16, 18, 24, 25, 31, 32, 33, 34, 35, 36, 37, 38, 39, 43, 44, 45, 46, 47, 48, 49, 50, 51	Block for all types feasible on this MRI scanner
17:00 - 17:30	3	1, 3, 4, 6, 7, 8, 9, 10, 11, 12, 13	Emergency block: patient types with acces time targets less than or equal to 3 days can be booked here
17:30 - 18:00	3	6, 7, 8, 9, 10, 11, 12, 13, 15, 16, 18, 24, 25, 31, 32, 33, 34, 35, 36, 37, 38, 39, 43, 44, 45, 46, 47, 48, 49, 50, 51	Block for all types feasible on this MRI scanner

(e) Blueprint calendar for MRI scanner 2 for Day 5 and 10 in the biweekly cycle. This blueprint calendar is based on Rijnstate's blueprint calendar for Fridays in weeks with odd and even week numbers, respectively.

Figure C.1: Blueprint calendars for MRI scanner 2 used for the validation of the simulation model.

Monday (Day 1 and Day 6 in the biweekly cycle)

Time	#Slots	Feasible patient types	Remarks
08:00 - 10:00	12	6, 7, 10, 11, 15, 16, 21, 24, 31, 32, 33, 34, 43, 44, 45	
10:00 - 11:30	9	16, 33	
11:30 - 12:30	6	6, 7, 8, 10, 11, 12, 14, 15, 16, 21, 24, 26, 31, 32, 33, 34, 35, 36, 43, 44, 45, 46	Block for all types feasible on this MRI scanner
12:30 - 13:00	3	6, 7, 10, 11, 15, 16, 21, 24, 31, 32, 33, 34, 43, 44, 45	
13:00 - 14:00	6	5, 6, 7, 8, 10, 11, 12, 14	Emergency block: patient types with acces time targets less than or equal to 3 days can be booked here
14:00 - 17:00	18	6, 7, 8, 10, 11, 12, 14, 15, 16, 21, 24, 26, 31, 32, 33, 34, 35, 36, 43, 44, 45, 46	Block for all types feasible on this MRI scanner
17:00 - 18:00	6	5, 6, 7, 8, 10, 11, 12, 14	Emergency block: patient types with acces time targets less than or equal to 3 days can be booked here

(a) Blueprint calendar for MRI scanner 3 for Day 1 and 6 in the biweekly cycle. This blueprint calendar is based on Rijnstate's blueprint calendar for Mondays in weeks with odd and even week numbers, respectively.

Tuesday (Day 2 and Day 7 in the biweekly cycle)

Time			#Slots	Feasible patient types	Remarks
08:00	-	09:30	9	6, 7, 10, 11, 15, 16, 21, 24, 31, 32, 33, 34, 43, 44, 45	
09:30	-	10:50	8	28, 29 (10, 11, 43, 44 / 1)	*one-stop-shop* patient types
10:50	-	11:10	2	27 (6, 31 / 1)	*one-stop-shop* patient type
11:10	-	12:30	8	6, 7, 10, 11, 12, 14, 15, 16, 21, 24, 26, 31, 32, 33, 34, 35, 36, 43, 44, 45, 46	Block for all types feasible on this MRI scanner
12:30	-	13:30	6	5, 6, 7, 8, 10, 11, 12, 14	Emergency block: patient types with acces time targets less than or equal to 3 days can be booked here
13:30	-	15:00	9	6, 7, 8, 10, 11, 12, 14, 15, 16, 21, 24, 26, 31, 32, 33, 34, 35, 36, 43, 44, 45, 46	Block for all types feasible on this MRI scanner
15:00	-	16:30	9	6, 7, 10, 11, 15, 16, 21, 24, 31, 32, 33, 34, 43, 44, 45	
16:30	-	17:00	3	16, 33	
17:00	-	17:30	3	12, 35, 46	
17:30	-	18:00	3	5, 6, 7, 8, 10, 11, 12, 14	Emergency block: patient types with acces time targets less than or equal to 3 days can be booked here

(b) Blueprint calendar for MRI scanner 3 for Day 2 and 7 in the biweekly cycle. This blueprint calendar is based on Rijnstate's blueprint calendar for Tuesdays in weeks with odd and even week numbers, respectively. In the weeks with even week numbers, the MRI scanner is operated till 21.00 p.m. In our biweekly cycle that is on Day 7.

Wednesday (Day 3 and Day 8 in the biweekly cycle)

Time			#Slots	Feasible patient types	Remarks
08:00	-	08:30	3	6, 7, 10, 11, 15, 16, 21, 24, 31, 32, 33, 34, 43, 44, 45	
08:30	-	10:30	12	6, 7, 8, 10, 11, 12, 14, 15, 16, 21, 24, 26, 31, 32, 33, 34, 35, 36, 43, 44, 45, 46	Block for all types feasible on this MRI scanner
10:30	-	11:30	6	5, 6, 7, 8, 10, 11, 12, 14	Emergency block: patient types with acces time targets less than or equal to 3 days can be booked here
11:30	-	11:50	2	27 (6, 31 / 1)	*one-stop-shop* patient type
11:50	-	12:30	4	17 (6, 31 / 1)	*one-stop-shop* patient type
12:30	-	13:30	6	5, 6, 7, 8, 10, 11, 12, 14	Emergency block: patient types with acces time targets less than or equal to 3 days can be booked here
13:00	-	15:00	12	6, 7, 8, 10, 11, 12, 14, 15, 16, 21, 24, 26, 31, 32, 33, 34, 35, 36, 43, 44, 45, 46	Block for all types feasible on this MRI scanner
15:30	-	17:00	9	16, 33	
17:00	-	18:00	6	5, 6, 7, 8, 10, 11, 12, 14	Emergency block: patient types with acces time targets less than or equal to 3 days can be booked here

(c) Blueprint calendars for MRI scanner 2 for Day 3 and Day 8 in the biweekly cycle. Blueprints are based on Rijnstate's blueprint calendars for Wednesdays in weeks with odd and even week numbers, respectively. In the weeks with even week numbers, the MRI scanner is operated till 21.00 p.m. In our biweekly cycle that is on Day 8.

Thursday (Day 4 and Day 9 in the biweekly cycle)

Time	#Slots	Feasible patient types	Remarks
08:00 - 09:00	6	6, 7, 10, 11, 15, 16, 21, 24, 31, 32, 33, 34, 43, 44, 45	
09:00 - 09:40	4	27 (6, 31 / 1)	*one-stop-shop* patient type
09:40 - 10:40	6	28, 29 (10, 11, 43, 44 / 1)	*one-stop-shop* patient types
10:40 - 11:10	3	16, 33	
11:10 - 11:50	4	17 (6, 31 / 1)	*one-stop-shop* patient type
11:50 - 12:30	4	6, 7, 8, 10, 11, 12, 14, 15, 16, 21, 24, 26, 31, 32, 33, 34, 35, 36, 43, 44, 45, 46	Block for all types feasible on this MRI scanner
12:30 - 13:30	6	5, 6, 7, 8, 10, 11, 12, 14	Emergency block: patient types with acces time targets less than or equal to 3 days can be booked here
13:30 - 14:30	6	30 (16, 33 / 1)	*one-stop-shop* patient type
14:30 - 17:00	15	6, 7, 8, 10, 11, 12, 14, 15, 16, 21, 24, 26, 31, 32, 33, 34, 35, 36, 43, 44, 45, 46	Block for all types feasible on this MRI scanner
17:00 - 18:00	6	5, 6, 7, 8, 10, 11, 12, 14	Emergency block: patient types with acces time targets less than or equal to 3 days can be booked here

(d) Blueprint calendar for MRI scanner 3 for Day 4 and 9 in the biweekly cycle. This blueprint calendar is based on Rijnstate's blueprint calendar for Thursdays in weeks with odd and even week numbers, respectively. In the weeks with even week numbers, the MRI scanner is operated till 21.00 p.m. In our biweekly cycle that is on Day 9.

Friday (Day 5 and Day 10 in the biweekly cycle)

Time	#Slots	Feasible patient types	Remarks
08:00 - 09:00	6	6, 7, 10, 11, 15, 16, 21, 24, 31, 32, 33, 34, 43, 44, 45	
09:00 - 09:40	4	27 (6, 31 / 1)	*one-stop-shop* patient type
09:40 - 10:40	6	28, 29 (10, 11, 43, 44 / 1)	*one-stop-shop* patient types
10:40 - 11:10	3	16, 33	
11:10 - 11:50	4	17 (6, 31 / 1)	*one-stop-shop* patient type
11:50 - 12:30	4	6, 7, 8, 10, 11, 12, 14, 15, 16, 21, 24, 26, 31, 32, 33, 34, 35, 36, 43, 44, 45, 46	Block for all types feasible on this MRI scanner
12:30 - 14:00	9		
14:00 - 15:00	6	5, 6, 7, 8, 10, 11, 12, 14	Emergency block: patient types with acces time targets less than or equal to 3 days can be booked here
15:00 - 16:00	6	30 (16, 33 / 1)	*one-stop-shop* patient type
16:00 - 17:00	6	6, 7, 8, 10, 11, 12, 14, 15, 16, 21, 24, 26, 31, 32, 33, 34, 35, 36, 43, 44, 45, 46	Block for all types feasible on this MRI scanner
17:00 - 18:00	6	5, 6, 7, 8, 10, 11, 12, 14	Emergency block: patient types with acces time targets less than or equal to 3 days can be booked here

(e) Blueprint calendar for MRI scanner 3 for Day 5 and 10 in the biweekly cycle. This blueprint calendar is based on Rijnstate's blueprint calendar for Fridays in weeks with odd and even week numbers, respectively.

Figure C.2: Blueprint calendars for MRI scanner 3 used for the validation of the simulation model.

C.3 Patient types definition for the MDP

Table C.2: Patient types definition for the MDP and some attributes. The second column shows how we have merged patient types based on the blueprint calendars (from Table C.1) into patient types for the MDP. Note that a HIX code may belong to multiple patient types. In that case the difference is in the access time target (abbreviated as AT in the header). Today is day 0 and tomorrow is counted as day 1. The third column shows the service time, in minutes, that is reserved for each patient type. The last columns shows whether a MRI scanner is suitable for the type of MRI examination (✓) or not (✗).

Patient type	Patient types merged	HIX code(s)	AT (in days)	Duration (in min.)	Compatibility MRI scanner ...		
					1	2	3
1	1	emergency; machine 1 & 2: 1390, 1390D, 1390A, 1390E, 2090, 2090A, 2990, 3190PB, 8490, 9090 L/R, 9190, 9290, 9491 L/R, ARTI, MWEKE	0	30	✓	✓	✗
2	2	emergency; only on machine 1: 1390C, 1390H, 1390K, 2290, 4090 L/R, 4090 AL/AR, 4292 L/R, 6990	0	30	✓	✗	✗
3	3	emergency; only on machine 2: 3090T, 3190, 3190D, 3290, 3390, 3390D, 3390P, 3490, 3690, 5090, 6990P, 7090, 7090A, 7490, 7690, 7790, 9090A, 9491 EL/ER, MRV	0	30	✗	✓	✗
4	4	emergency; only on machine 2: 5190/5192, 5191/5191R, 7690P	0	60	✗	✓	✗
5	5	emergency; only machine 3: 3090T, 3190, 3290, 3390, 3390P, 3490, 5090, 7090, 7091, 7490, 7690, 7790, 8490, 9090 L/R, 9190, 9190E, 9290, 9491, MWEKE	0	30	✗	✗	✓
6	6	1390, 1390D	3	20	✓	✓	✓
7	7 & 8	1390A, 2090, 3190PB; 3390P, 8490, 9190	3	30	✓	✓	✓
8	9	2290, 2990	3	30	✓	✓	✗
9	10 - 12	3190, 3190D, 3390, 3390D, 5090, 7090, 7490, 7690, 7790	3	30	✗	✓	✓
10	13	3690	3	30	✗	✓	✗
11	14	7091	3	20	✗	✗	✓
12	15 & 16	4090 L/R; 9090 L/R; 9290; MWEKE, 9491 L/R	5	30	✓	✓	✓
13	17 & 19 - 21	NEURO, DAG, TIA/TIAS, 1390C	5	20	✓	✓	✗
14	18	2090A	5	30	✓	✓	✗
15	22 & 23	1390H, 6990M	5	30	✓	✗	✗

(Continued on next page)

Table C.2 – *Continued*

Patient type	Patient types merged	HIX code(s)	AT (in days)	Duration (in min.)	Compatibility MRI scanner ...		
					1	2	3
16	24	3090T, 3290, 3490	5	30	✗	✓	✓
17	25	7690P	5	40	✗	✓	✗
18	26	9190E	5	30	✗	✗	✓
19	27 - 29	TIAZ, LWK, CWK	5	20	✗	✗	✓
20	30 & 31	AKP L/R, 1390, 1390D	10	20	✓	✓	✓
21	32 - 36	1390A, 1390E, 2090, 3190PB; 3390P,4090 L/R, 8490,9090 L/R, 9190, 9290, 9491 L/R, MWEKE	10	30	✓	✓	✓
22	37 & 38	2290, 2990, ARTI	10	30	✓	✓	✗
23	39	1390-3190 combi; 1390D-3190D combi; 1390-3090T combi; 1390-3690 combi; 3190-3190 combi; 3190-3390 combi; 3190D-3390D combi; 3190-3190PB combi; 3290-3390 combi; 3390-3390P combi; 3090T-3490 combi; 7690-7790 combi	10	40	✓	✓	✗
24	40 - 42	1390K, 4090 AR/AL, 6990, 9090 AL/AR, 4292 L/R	10	30	✓	✗	✗
25	43 - 46	3090T, 3190, 3190D, 3290, 3390, 3390D, 3490, 5090, 7090, 7490, 7690, 7790	10	30	✗	✓	✓
26	47	5190/5192, 5191/5191R	10	60	✗	✓	✗
27	48 - 51	3690, 6990P, 7090A, 9090A, MRV, ARTIP	10	30	✗	✓	✗

C.4 Blueprint calendars for MDP patient types

Monday (Day 1 and Day 6 in the biweekly cycle)

Time			#Slots	Feasible patient types	Remarks
08:00	-	09:10	7	12, 21	
09:10	-	09:50	4	6, 20	
09:50	-	10:20	3	7, 12, 12, 21,	
10:20	-	10:50	3	15	
10:50	-	11:20	3	24	
11:20	-	12:00	4	13	
12:00	-	12:30	3	22	
12:30	-	13:30	6	12, 21, 24	
13:30	-	14:20	5	13	
14:20	-	15:20	6	1, 2, 6, 7, 8	Emergency block: patient types with acces time targets less than or equal to 3 days can be booked here
15:20	-	16:20	6	6, 7, 8, 12, 12, 13, 15, 20, 21, 22, 22, 24	Block for all types feasible on this MRI scanner
16:20	-	17:00	4	12, 21	
17:00	-	18:00	6	1, 2, 6, 7, 8	Emergency block: patient types with acces time targets less than or equal to 3 days can be booked here

(a) Blueprint calendar for MRI scanner 1 for Day 1 and 6 in the biweekly cycle. This blueprint calendar is based on Rijnstate's blueprint calendar for Mondays in weeks with odd and even week numbers, respectively.

Tuesday (Day 2 (till 18:00) and Day 7 (till 21:00) in the biweekly cycle)

Time			#Slots	Feasible patient types	Remarks
08:00	-	08:30	3	12, 21	
08:30	-	09:00	3	7, 12, 21	
09:00	-	10:00	6	15 (24 / 5)	
10:00	-	10:30	3	24	
10:30	-	11:00	3	1, 2, 6, 7, 8	Emergency block: patient types with acces time targets less than or equal to 3 days can be booked here
11:00	-	12:30	9	6, 7, 8, 12, 13, 15, 20, 21, 22, 22, 24	Block for all types feasible on this MRI scanner
12:30	-	12:50	2	13	
12:50	-	13:30	4	14	
13:30	-	14:30	6	1, 2, 6, 7, 8	Emergency block: patient types with acces time targets less than or equal to 3 days can be booked here
14:30	-	15:00	3	22	
15:00	-	15:30	3	7, 15, 21	
15:30	-	16:00	3	6, 20	
16:00	-	16:30	3	7, 12, 21	
16:30	-	17:00	3	6, 7, 8, 12, 13, 15, 20, 21, 22, 22, 24	Block for all types feasible on this MRI scanner
17:00	-	18:00	6	1, 2, 6, 7, 8	Emergency block: patient types with acces time targets less than or equal to 3 days can be booked here
18:00	-	19:30	9	6, 7, 8, 12, 20, 21, 22	
19:30	-	21:00	9	6, 20	

(b) Blueprint calendar for MRI scanner 1 for Day 2 and 7 in the biweekly cycle. This blueprint calendar is based on Rijnstate's blueprint calendar for Tuesdays in weeks with odd and even week numbers, respectively. In the weeks with even week numbers, the MRI scanner is operated till 21.00 p.m. In our biweekly cycle that is on Day 7.

Wednesday (Day 3 (till 18:00) and Day 8 (till 21:00) in the biweekly cycle)

Time			#Slots	Feasible patient types	Remarks
08:00	-	08:30	3	12, 21	
08:30	-	09:10	4	20	
09:10	-	10:10	6	15	
10:10	-	10:40	3	24	
10:40	-	11:40	6	13	
11:40	-	12:10	3	22	
11:50	-	12:30	4	6, 7, 8, 12, 12, 13, 15, 20, 21, 22, 22, 24	Block for all types feasible on this MRI scanner
12:30	-	13:00	3	6, 7, 8, 15, 20, 21, 22	
13:00	-	13:30	3	1, 2, 6, 7, 8	Emergency block: patient types with acces time targets less than or equal to 3 days can be booked here
13:30	-	15:00	9	6, 7, 8, 12, 12, 13, 15, 20, 21, 22, 22, 24	Block for all types feasible on this MRI scanner
15:00	-	16:30	9	12, 21, 24	
16:30	-	17:00	3	12, 21	
17:00	-	18:00	6	1, 2, 6, 7, 8	Emergency block: patient types with acces time targets less than or equal to 3 days can be booked here
18:00	-	19:30	9	6, 7, 8, 12, 20, 21, 22	
19:30	-	21:00	9	6, 12, 20, 21	

(c) Blueprint calendars for MRI scanner 1 for Day 3 and Day 8 in the biweekly cycle. Blueprints are based on Rijnstate's blueprint calendars for Wednesdays in weeks with odd and even week numbers, respectively. In the weeks with even week numbers, the MRI scanner is operated till 21.00 p.m. In our biweekly cycle that is on Day 8.

Thursday, odd weeknumbers (Day 4 in the biweekly cycle)

Time			#Slots	Feasible patient types	Remarks
08:00	-	08:30	3	12, 21	
08:30	-	12:00	21	24	
12:00	-	12:30	3	15	
12:30	-	13:30	6	7, 12, 21	
13:30	-	14:00	3	1, 2, 6, 7, 8	Emergency block: patient types with acces time targets less than or equal to 3 days can be booked here
14:00	-	15:00	6	6, 7, 8, 12, 20, 21, 22	
15:00	-	15:30	3	7, 12, 21	
15:30	-	16:30	6	6, 7, 8, 12, 13, 15, 20, 21, 22, 22, 24	Block for all types feasible on this MRI scanner
16:30	-	17:00	3	12, 21	
17:00	-	18:00	6	1, 2, 6, 7, 8	Emergency block: patient types with acces time targets less than or equal to 3 days can be booked here

(d) Blueprint calendar for MRI scanner 1 for Day 4 in the biweekly cycle. This blueprint calendar is based on Rijnstate's blueprint calendar for Thursdays in weeks with odd week numbers, respectively. In the weeks with even week numbers, the MRI scanner is operated till 21.00 p.m. In our biweekly cycle that is on Day 9.

Thursday, even weeknumbers (Day 9 in the biweekly cycle)

Time	#Slots	Feasible patient types	Remarks
08:00 - 08:30	3	12, 21	
08:30 - 10:30	12	6, 7, 8, 15, 20, 21, 22	
10:30 - 11:30	6	24	
11:30 - 12:00	3	15	*one-stop-shop* patient type
12:00 - 12:30	3	7, 12, 12, 21,	
12:30 - 13:00	3	6, 7, 8, 12, 20, 21, 22	
13:00 - 14:00	6	1, 2, 6, 7, 8	Emergency block: patient types with acces time targets less than or equal to 3 days can be booked here
14:00 - 15:30	9	7, 12, 12, 21,	
15:30 - 16:30	6	6, 7, 8, 12, 12, 13, 15, 20, 21, 22, 22, 24	Block for all types feasible on this MRI scanner
16:30 - 17:00	3	6, 7, 8, 12, 20, 21, 22	
17:00 - 18:00	6	1, 2, 6, 7, 8	Emergency block: patient types with acces time targets less than or equal to 3 days can be booked here
18:00 - 19:30	9	6, 7, 8, 12, 20, 21, 22	
19:30 - 21:00	9	12, 21	

(e) Blueprint calendar for MRI scanner 1 for Day 9 in the biweekly cycle. This blueprint calendar is based on Rijnstate's blueprint calendar for Thursdays in weeks with even week numbers.

Friday (Day 5 and Day 10 in the biweekly cycle)

Time	#Slots	Feasible patient types	Remarks
08:00 - 09:00	6	7, 12, 21	
09:00 - 10:00	6	13	
10:00 - 10:30	3	24	
10:30 - 11:00	3	15	
11:00 - 11:30	3	6, 7, 8, 15, 20, 21, 22	
11:30 - 12:00	3	13	
12:00 - 12:30	3	6, 7, 8, 12, 13, 15, 20, 21, 22, 22, 24	Block for all types feasible on this MRI scanner
12:30 - 14:00	9		
14:00 - 15:00	6	1, 2, 6, 7, 8	Emergency block: patient types with acces time targets less than or equal to 3 days can be booked here
15:00 - 16:00	6	12, 21, 24	
16:00 - 17:00	6	6, 7, 8, 12, 13, 15, 20, 21, 22, 22, 24	Block for all types feasible on this MRI scanner
17:00 - 18:00	6	1, 2, 6, 7, 8	Emergency block: patient types with acces time targets less than or equal to 3 days can be booked here

(f) Blueprint calendar for MRI scanner 1 for Day 5 and 10 in the biweekly cycle. This blueprint calendar is based on Rijnstate's blueprint calendar for Fridays in weeks with odd and even week numbers, respectively.

Figure C.3: Blueprint calendars for MRI scanner 1 used for simulating the FAS-wB and M-wB policies in the comparison to the AOP in Section 6.3.

Monday (Day 1 and Day 6 in the biweekly cycle)

Time			#Slots	Feasible patient types	Remarks
08:00	-	08:30	3	9, 25	
08:30	-	09:30	6	6, 7, 8, 9, 10, 12, 16, 20, 21, 22, 25, 27	
09:30	-	10:30	6	13	
10:30	-	11:00	3	7, 9, 16, 21, 25	
11:00	-	12:00	6	4, 26	
12:00	-	12:30	3	6, 7, 8, 9, 10, 12, 16, 17, 20, 21, 22, 25, 26, 27	Block for all types feasible on this MRI scanner
12:30	-	13:00	3	1, 3, 4, 6, 7, 8, 9, 10	Emergency block: patient types with acces time targets less than or equal to 3 days can be booked here
13:00	-	13:30	3	7, 9, 16, 21, 25	
13:30	-	16:30	18	4, 26	
16:30	-	17:00	3	6, 7, 8, 9, 10, 12, 16, 17, 20, 21, 22, 25, 26, 27	Block for all types feasible on this MRI scanner
17:00	-	18:00	6	1, 3, 4, 6, 7, 8, 9, 10	Emergency block: patient types with acces time targets less than or equal to 3 days can be booked here

(a) Blueprint calendar for MRI scanner 2 for Day 1 and 6 in the biweekly cycle. This blueprint calendar is based on Rijnstate's blueprint calendar for Mondays in weeks with odd and even week numbers, respectively.

Tuesday (Day 2 (till 18:00) and Day 7 (till 21:00) in the biweekly cycle)

Time			#Slots	Feasible patient types	Remarks
08:00	-	08:30	3	9, 25	
08:30	-	09:00	3	6, 7, 8, 9, 10, 12, 16, 20, 21, 22, 25, 27	
09:00	-	09:40	4	13	
09:40	-	12:30	17	4, 26	
12:30	-	15:30	18	4, 26	
15:30	-	16:30	6	1, 3, 4, 6, 7, 8, 9, 10	Emergency block: patient types with acces time targets less than or equal to 3 days can be booked here
16:30	-	17:00	3	7, 8, 9, 10, 12, 21, 22, 25, 27	
17:00	-	17:30	3	1, 3, 4, 6, 7, 8, 9, 10	Emergency block: patient types with acces time targets less than or equal to 3 days can be booked here
17:30	-	18:00	3	27	
18:00	-	18:30	3	9, 16, 25	
18:30	-	19:30	6	7, 8, 9, 10, 12, 21, 22, 25, 27	
19:30	-	20:30	6	7, 9, 16, 21, 25	
20:30	-	21:00	3	6, 9, 16, 20, 22, 25	

(b) Blueprint calendar for MRI scanner 2 for Day 2 and 7 in the biweekly cycle. This blueprint calendar is based on Rijnstate's blueprint calendar for Tuesdays in weeks with odd and even week numbers, respectively. In the weeks with even week numbers, the MRI scanner is operated till 21.00 p.m. In our biweekly cycle that is on Day 7.

Wednesday (Day 3 (till 18:00) and Day 8 (till 21:00) in the biweekly cycle)

Time			#Slots	Feasible patient types	Remarks
08:00	-	08:30	3	12, 21	
08:30	-	09:10	4	20	
09:10	-	10:00	5	6, 7, 8, 9, 10, 12, 13, 16, 17, 20, 21, 22, 22, 25, 26, 26, 27	Block for all types feasible on this MRI scanner
10:00	-	10:30	3	26	
10:30	-	11:30	6	13	
11:30	-	12:00	3	22	
12:00	-	12:30	3	27	
12:30	-	13:00	3	6, 7, 8, 20, 21, 22	
13:00	-	14:00	6	1, 3, 4, 6, 7, 8, 9, 10	Emergency block: patient types with acces time targets less than or equal to 3 days can be booked here
14:00	-	14:30	3	22	
14:30	-	16:00	9	12, 21	
16:00	-	16:30	3	6, 7, 8, 9, 10, 12, 13, 16, 17, 20, 21, 22, 22, 25, 26, 26, 27	Block for all types feasible on this MRI scanner
16:30	-	17:00	3	6, 20	
17:00	-	17:30	3	1, 3, 4, 6, 7, 8, 9, 10	Emergency block: patient types with acces time targets less than or equal to 3 days can be booked here
17:30	-	18:00	3	27	
18:00	-	18:30	3	9, 16, 25	
18:30	-	19:30	6	7, 8, 9, 10, 12, 21, 22, 25, 27	
19:30	-	20:30	6	7, 9, 16, 21, 25	
20:30	-	21:00	3	6, 9, 16, 20, 22, 25	

(c) Blueprint calendars for MRI scanner 2 for Day 3 and Day 8 in the biweekly cycle. Blueprints are based on Rijnstate's blueprint calendars for Wednesdays in weeks with odd and even week numbers, respectively. In the weeks with even week numbers, the MRI scanner is operated till 21.00 p.m. In our biweekly cycle that is on Day 8.

Thursday (Day 4 (till 18:00) and Day 9 (till 21:00) in the biweekly cycle)

Time			#Slots	Feasible patient types	Remarks
08:00	-	08:30	3	9, 25	
08:30	-	09:00	3	6, 7, 8, 9, 10, 12, 16, 20, 21, 22, 25, 27	
09:00	-	10:00	6	13	*one-stop-shop* patient type
10:00	-	10:30	3	27	
10:30	-	11:30	6	7, 9, 16, 21, 25	
11:30	-	12:30	6	7, 9, 17, 22, 25, 27	
12:30	-	13:10	4	14	
13:10	-	14:40	9	1, 3, 4, 6, 7, 8, 9, 10	Emergency block: patient types with acces time targets less than or equal to 3 days can be booked here
14:40	-	15:10	3	6, 7, 8, 9, 10, 12, 13, 16, 17, 20, 21, 22, 22, 25, 26, 26, 27	Block for all types feasible on this MRI scanner
15:10	-	16:10	6	7, 9, 17, 22, 25, 27	
16:10	-	16:40	3	6, 7, 9, 20, 21, 25	
16:40	-	17:00	2	6, 20	
17:00	-	17:30	3	1, 3, 4, 6, 7, 8, 9, 10	Emergency block: patient types with acces time targets less than or equal to 3 days can be booked here
17:30	-	19:30	12	9, 16, 25	
19:30	-	20:30	6	7, 9, 16, 21, 25	
20:30	-	21:00	3	6, 9, 16, 20, 22, 25	

(d) Blueprint calendar for MRI scanner 2 for Day 4 and Day 9 in the biweekly cycle. This blueprint calendar is based on Rijnstate's blueprint calendar for Thursdays in weeks with odd and even week numbers, respectively. In the weeks with even week numbers, the MRI scanner is operated till 21.00 p.m. In our biweekly cycle that is on Day 9.

Friday (Day 5 and Day 10 in the biweekly cycle)

Time	#Slots	Feasible patient types	Remarks
08:00 - 08:30	3	9, 25	
08:30 - 09:30	6	6, 7, 8, 9, 10, 12, 16, 20, 21, 22, 25, 27	
09:30 - 10:10	4	13	
10:10 - 11:10	6	7, 9, 17, 22, 25, 27	
11:10 - 12:30	8	6, 7, 8, 9, 10, 12, 13, 16, 17, 20, 21, 22, 22, 25, 26, 26, 27	Block for all types feasible on this MRI scanner
12:30 - 14:00	9		
14:00 - 15:00	6	7, 9, 17, 22, 25, 27	
15:00 - 15:30	3	1, 3, 4, 6, 7, 8, 9, 10	Emergency block: patient types with acces time targets less than or equal to 3 days can be booked here
15:30 - 17:00	9	6, 7, 8, 9, 10, 12, 13, 16, 17, 20, 21, 22, 22, 25, 26, 26, 27	Block for all types feasible on this MRI scanner
17:00 - 17:30	3	1, 3, 4, 6, 7, 8, 9, 10	Emergency block: patient types with acces time targets less than or equal to 3 days can be booked here
17:30 - 18:00	3	6, 7, 8, 9, 10, 12, 13, 16, 17, 20, 21, 22, 22, 25, 26, 26, 27	Block for all types feasible on this MRI scanner

(e) Blueprint calendar for MRI scanner 2 for Day 5 and 10 in the biweekly cycle. This blueprint calendar is based on Rijnstate's blueprint calendar for Fridays in weeks with odd and even week numbers, respectively.

Figure C.4: Blueprint calendars for MRI scanner 2 used for simulating the FAS-wB and M-wB policies in the comparison to the AOP in Section 6.3.

Monday (Day 1 and Day 6 in the biweekly cycle)

Time	#Slots	Feasible patient types	Remarks
08:00 - 10:00	12	6, 7, 9, 12, 13, 16, 20, 21, 25	
10:00 - 11:30	9	12, 21	
11:30 - 12:30	6	6, 7, 9, 11, 12, 13, 16, 18, 20, 21, 25	Block for all types feasible on this MRI scanner
12:30 - 13:00	3	6, 7, 9, 12, 13, 16, 20, 21, 25	
13:00 - 14:00	6	5, 6, 7, 9, 11	Emergency block: patient types with acces time targets less than or equal to 3 days can be booked here
14:00 - 17:00	18	6, 7, 9, 11, 12, 13, 16, 18, 20, 21, 25	Block for all types feasible on this MRI scanner
17:00 - 18:00	6	5, 6, 7, 9, 11	Emergency block: patient types with acces time targets less than or equal to 3 days can be booked here

(a) Blueprint calendar for MRI scanner 3 for Day 1 and 6 in the biweekly cycle. This blueprint calendar is based on Rijnstate's blueprint calendar for Mondays in weeks with odd and even week numbers, respectively.

Tuesday (Day 2 and Day 7 in the biweekly cycle)

Time			#Slots	Feasible patient types	Remarks
08:00	-	09:30	9	6, 7, 9, 12, 13, 16, 20, 21, 25	
09:30	-	11:10	10	19	
11:10	-	12:30	8	6, 7, 9, 11, 12, 13, 16, 18, 20, 21, 25	Block for all types feasible on this MRI scanner
12:30	-	13:30	6	5, 6, 7, 9, 11	Emergency block: patient types with acces time targets less than or equal to 3 days can be booked here
13:30	-	15:00	9	6, 7, 9, 11, 12, 13, 16, 18, 20, 21, 25	Block for all types feasible on this MRI scanner
15:00	-	16:30	9	6, 7, 9, 12, 13, 16, 20, 21, 25	
16:30	-	17:00	3	12, 21	
17:00	-	17:30	3	9, 21, 25	
17:30	-	18:00	3	5, 6, 7, 9, 11	Emergency block: patient types with acces time targets less than or equal to 3 days can be booked here

(b) Blueprint calendar for MRI scanner 3 for Day 2 and 7 in the biweekly cycle. This blueprint calendar is based on Rijnstate's blueprint calendar for Tuesdays in weeks with odd and even week numbers, respectively.

Wednesday (Day 3 and Day 8 in the biweekly cycle)

Time			#Slots	Feasible patient types	Remarks
08:00	-	08:30	3	6, 7, 9, 12, 13, 16, 20, 21, 25	
08:30	-	10:30	12	6, 7, 9, 11, 12, 13, 16, 18, 20, 21, 25	Block for all types feasible on this MRI scanner
10:30	-	11:30	6	5, 6, 7, 9, 11	Emergency block: patient types with acces time targets less than or equal to 3 days can be booked here
11:30	-	11:50	2	19	
11:50	-	12:30	4	13	
12:30	-	13:30	6	5, 6, 7, 9, 11	Emergency block: patient types with acces time targets less than or equal to 3 days can be booked here
13:00	-	15:00	12	6, 7, 9, 11, 12, 13, 16, 18, 20, 21, 25	Block for all types feasible on this MRI scanner
15:30	-	17:00	9	12, 21	
17:00	-	18:00	6	5, 6, 7, 9, 11	Emergency block: patient types with acces time targets less than or equal to 3 days can be booked here

(c) Blueprint calendar for MRI scanner 3 for Day 3 and 8 in the biweekly cycle. This blueprint calendar is based on Rijnstate's blueprint calendar for Wednesdays in weeks with odd and even week numbers, respectively.

Thursday (Day 4 and Day 9 in the biweekly cycle)

Time			#Slots	Feasible patient types	Remarks
08:00	-	09:00	6	6, 7, 9, 12, 13, 16, 20, 21, 25	
09:00	-	10:40	10	19	
10:40	-	11:10	3	12, 21	
11:10	-	11:50	4	13	
11:50	-	12:30	4	6, 7, 9, 11, 12, 13, 16, 18, 20, 21, 25	Block for all types feasible on this MRI scanner
12:30	-	13:30	6	5, 6, 7, 9, 11	Emergency block: patient types with acces time targets less than or equal to 3 days can be booked here
13:30	-	14:30	6	20	
14:30	-	17:00	15	6, 7, 9, 11, 12, 13, 16, 18, 20, 21, 25	Block for all types feasible on this MRI scanner
17:00	-	18:00	6	5, 6, 7, 9, 11	Emergency block: patient types with acces time targets less than or equal to 3 days can be booked here

(d) Blueprint calendar for MRI scanner 3 for Day 4 and 9 in the biweekly cycle. This blueprint calendar is based on Rijnstate's blueprint calendar for Thursdays in weeks with odd and even week numbers, respectively.

Friday (Day 5 and Day 10 in the biweekly cycle)

Time			#Slots	Feasible patient types	Remarks
08:00	-	09:00	6	6, 7, 9, 12, 13, 16, 20, 21, 25	
09:00	-	10:40	10	19	
10:40	-	11:10	3	12, 21	
11:10	-	11:50	4	13	
11:50	-	12:30	4	6, 7, 9, 11, 12, 13, 16, 18, 20, 21, 25	Block for all types feasible on this MRI scanner
12:30	-	14:00	9		
14:00	-	15:00	6	5, 6, 7, 9, 11	Emergency block: patient types with acces time targets less than or equal to 3 days can be booked here
15:00	-	16:00	6	20	
16:00	-	17:00	6	6, 7, 9, 11, 12, 13, 16, 18, 20, 21, 25	Block for all types feasible on this MRI scanner
17:00	-	18:00	6	5, 6, 7, 9, 11	Emergency block: patient types with acces time targets less than or equal to 3 days can be booked here

(e) Blueprint calendar for MRI scanner 3 for Day 5 and 10 in the biweekly cycle. This blueprint calendar is based on Rijnstate's blueprint calendar for Fridays in weeks with odd and even week numbers, respectively.

Figure C.5: Blueprint calendars for MRI scanner 3 used for simulating the FAS-wB and M-wB policies in the comparison to the AOP in Section 6.3.

D

Additional results

D.1 Additional results for Rijnstate

Table D.1: Percentages of patients of patient types 1-5 (access time target of 0 days) examined within a certain number of days in the Rijnstate instance of Section 6.3.

AOP					
Patient Type	**(cumulative) % examined on day ...**				
	0	1	3	5	10
1	80.25 ± 1.36	96.47 ± 0.61	100.00 ± 0.00	100.00 ± 0.00	100.00 ± 0.00
2	83.81 ± 2.53	$98,34 \pm 1.00$	100.00 ± 0.00	100.00 ± 0.00	100.00 ± 0.00
3	73.50 ± 1.83	94.05 ± 2.32	100.00 ± 0.00	100.00 ± 0.00	100.00 ± 0.00
4	47.61 ± 3.96	60.33 ± 0.67	100.00 ± 0.00	100.00 ± 0.00	100.00 ± 0.00
5	79.69 ± 1.44	95.87 ± 0.83	100.00 ± 0.00	100.00 ± 0.00	100.00 ± 0.00
FAS-nB					
Patient Type	**(cumulative) % examined on day ...**				
	0	1	3	5	10
1	17.52 ± 2.55	100.00 ± 0.00	100.00 ± 0.00	100.00 ± 0.00	100.00 ± 0.00
2	6.33 ± 0.84	100.00 ± 0.00	100.00 ± 0.00	100.00 ± 0.00	100.00 ± 0.00
3	10.88 ± 1.63	100.00 ± 0.00	100.00 ± 0.00	100.00 ± 0.00	100.00 ± 0.00
4	4.31 ± 2.55	100.00 ± 0.00	100.00 ± 0.00	100.00 ± 0.00	100.00 ± 0.00
5	42.66 ± 5.76	100.00 ± 0.00	100.00 ± 0.00	100.00 ± 0.00	100.00 ± 0.00

(Continued on next page)

Table D.1 – ***Continued***

			M-nB		
Patient Type	**(cumulative) % examined on day ...**				
	0	1	3	5	10
1	17.68 ± 0.76	100.00 ± 0.00	100.00 ± 0.00	100.00 ± 0.00	100.00 ± 0.00
2	6.79 ± 0.18	100.00 ± 0.00	100.00 ± 0.00	100.00 ± 0.00	100.00 ± 0.00
3	11.02 ± 0.51	100.00 ± 0.00	100.00 ± 0.00	100.00 ± 0.00	100.00 ± 0.00
4	5.83 ± 0.59	100.00 ± 0.00	100.00 ± 0.00	100.00 ± 0.00	100.00 ± 0.00
5	42.82 ± 1.78	100.00 ± 0.00	100.00 ± 0.00	100.00 ± 0.00	100.00 ± 0.00
			FAS-wB		
Patient Type	**(cumulative) % examined on day ...**				
	0	1	3	5	10
1	66.59 ± 1.96	93.97 ± 0.28	100.00 ± 0.00	100.00 ± 0.00	100.00 ± 0.00
2	46.25 ± 3.70	97.69 ± 1.48	100.00 ± 0.00	100.00 ± 0.00	100.00 ± 0.00
3	53.03 ± 1.00	97.33 ± 0.32	100.00 ± 0.00	100.00 ± 0.00	100.00 ± 0.00
4	15.43 ± 1.76	44.44 ± 0.34	100.00 ± 0.00	100.00 ± 0.00	100.00 ± 0.00
5	61.29 ± 2.42	92.96 ± 2.43	100.00 ± 0.00	100.00 ± 0.00	100.00 ± 0.00
			M-nB		
Patient Type	**(cumulative) % examined on day ...**				
	0	1	3	5	10
1	68.75 ± 0.79	93.97 ± 0.82	100.00 ± 0.00	100.00 ± 0.00	100.00 ± 0.00
2	48.77 ± 2.91	97.69 ± 0.81	100.00 ± 0.00	100.00 ± 0.00	100.00 ± 0.00
3	56.17 ± 1.10	97.74 ± 0.32	100.00 ± 0.00	100.00 ± 0.00	100.00 ± 0.00
4	15.80 ± 1.02	95.49 ± 1.00	100.00 ± 0.00	100.00 ± 0.00	100.00 ± 0.00
5	68.27 ± 0.88	92.96 ± 0.95	100.00 ± 0.00	100.00 ± 0.00	100.00 ± 0.00

Table D.3: Booking distribution across MRI scanners per patient type in the Rijnstate instance of Section 6.3.

Patient type	**AOP**			**FAS-nB**			**M-nB**			**FAS-wB**			**M-wB**		
	1	**2**	**3**	**1**	**2**	**3**	**1**	**2**	**3**	**1**	**2**	**3**	**1**	**2**	**3**
1	86.59 ± 2.02	13.41 ± 2.02	0.00 ± 0.00	79.69 ± 1.45	20.31 ± 1.45	0.00 ± 0.00	79.92 ± 1.66	20.08 ± 1.66	0.00 ± 0.00	80.21 ± 2.23	19.79 ± 1.48	0.00 ± 0.00	80.21 ± 2.09	19.79 ± 1.18	0.00 ± 0.00
6	48.74 ± 2.11	28.25 ± 1.08	23.01 ± 1.88	60.88 ± 2.80	12.99 ± 1.58	26.13 ± 1.39	60.39 ± 2.97	13.09 ± 1.64	26.52 ± 1.38	72.31 ± 0.88	12.81 ± 0.48	14.88 ± 1.01	72.38 ± 1.14	12.77 ± 0.31	14.85 ± 1.02
7	68.09 ± 1.29	20.60 ± 0.77	11.31 ± 0.81	55.69 ± 1.21	12.31 ± 0.52	31.99 ± 1.19	55.69 ± 1.00	12.07 ± 0.59	32.25 ± 1.34	46.54 ± 0.50	22.12 ± 0.94	31.34 ± 0.54	46.73 ± 0.46	22.45 ± 0.90	30.82 ± 0.48
8	74.53 ± 0.81	25.47 ± 0.95	0.00 ± 0.94	77.65 ± 2.44	22.35 ± 1.74	0.00 ± 1.48	78.18 ± 2.65	21.82 ± 1.85	0.00 ± 1.29	83.21 ± 0.44	16.79 ± 1.28	0.00 ± 0.00	83.29 ± 0.48	16.71 ± 1.30	0.00 ± 0.00
9	0.00 ± 0.95	66.12 ± 0.68	33.88 ± 0.51	0.00 ± 1.21	40.77 ± 0.54	59.23 ± 1.25	0.00 ± 1.08	40.66 ± 0.37	59.34 ± 1.04	0.00 ± 0.00	46.44 ± 0.30	53.56 ± 0.52	0.00 ± 0.23	46.68 ± 0.36	53.32 ± 0.43
12	23.23 ± 0.61	40.68 ± 0.70	36.09 ± 0.50	34.05 ± 1.66	29.78 ± 1.69	36.17 ± 2.02	34.39 ± 1.79	30.06 ± 1.47	35.55 ± 1.87	61.82 ± 0.50	11.35 ± 0.33	26.83 ± 0.75	61.89 ± 0.42	11.54 ± 0.34	26.57 ± 0.68
13	31.45 ± 1.41	68.55 ± 1.41	0.00 ± 0.00	61.01 ± 0.85	38.99 ± 0.85	0.00 ± 0.00	61.19 ± 1.03	38.81 ± 1.03	0.00 ± 0.00	71.73 ± 1.28	28.27 ± 0.75	0.00 ± 0.00	72.29 ± 1.24	27.71 ± 0.62	0.00 ± 0.00
14	25.25 ± 2.28	74.75 ± 2.28	0.00 ± 0.00	59.84 ± 2.65	40.16 ± 2.65	0.00 ± 0.00	59.79 ± 2.63	40.21 ± 2.63	0.00 ± 0.00	98.18 ± 0.97	1.82 ± 0.14	0.00 ± 0.00	98.28 ± 1.00	1.72 ± 0.14	0.00 ± 0.00
16	0.00 ± 0.00	63.05 ± 2.36	36.95 ± 2.36	0.00 ± 0.00	57.35 ± 2.34	42.65 ± 2.34	0.00 ± 0.00	58.07 ± 2.16	41.93 ± 2.16	0.00 ± 0.00	36.74 ± 0.64	63.26 ± 0.59	0.00 ± 0.00	36.97 ± 0.67	63.03 ± 0.58
20	21.45 ± 0.78	6.39 ± 0.32	72.16 ± 0.90	26.61 ± 1.04	34.27 ± 1.73	39.12 ± 1.96	26.62 ± 0.85	34.82 ± 1.42	38.56 ± 1.71	53.96 ± 0.51	14.96 ± 0.33	31.08 ± 0.99	53.45 ± 0.56	15.20 ± 0.43	31.35 ± 0.98
21	24.51 ± 0.36	15.14 ± 0.71	60.35 ± 0.93	24.01 ± 1.06	30.17 ± 1.53	45.82 ± 0.92	23.96 ± 1.02	30.38 ± 1.47	45.66 ± 0.98	46.44 ± 0.43	17.29 ± 0.23	36.27 ± 0.65	46.38 ± 0.49	17.26 ± 0.16	36.36 ± 0.63
22	23.23 ± 0.46	76.77 ± 0.46	0.00 ± 0.00	67.11 ± 1.24	32.89 ± 1.24	0.00 ± 0.00	66.69 ± 1.56	33.31 ± 1.56	0.00 ± 0.00	68.00 ± 0.70	32.00 ± 0.59	0.00 ± 0.00	68.05 ± 0.82	31.95 ± 0.70	0.00 ± 0.00
23	51.45 ± 1.05	48.55 ± 0.65	0.00 ± 1.55	49.85 ± 1.53	50.15 ± 0.69	0.00 ± 1.33	49.08 ± 1.99	50.92 ± 0.90	0.00 ± 1.50	60.28 ± 1.25	39.72 ± 1.22	0.00 ± 0.00	59.60 ± 1.26	40.40 ± 1.15	0.00 ± 0.00
25	0.00 ± 0.00	28.89 ± 0.77	71.11 ± 0.77	0.00 ± 0.00	51.54 ± 1.72	48.46 ± 1.72	0.00 ± 0.00	50.96 ± 1.54	49.04 ± 1.54	0.00 ± 0.00	43.13 ± 0.19	56.87 ± 0.65	0.00 ± 0.00	43.14 ± 0.13	56.86 ± 0.66

D.2 Additional results for the heavier loaded system

Table D.4: Percentages of patients of patient types 1-5 (access time target of 0 days) examined within a certain number of days in the heavier loaded system of Section 6.4.

			AOP			
Patient Type	**(cumulative) % examined on day ...**					
	1	2	5	7	23	25
1	92.18 ± 0.93	98.13 ± 0.72	100.00 ± 0.00	100.00 ± 0.00	100.00 ± 1.50	100.00 ± 0.00
2	92.00 ± 1.40	97.45 ± 0.94	99.92 ± 0.12	100.00 ± 0.00	100.00 ± 0.52	100.00 ± 0.00
3	91.61 ± 1.33	96.11 ± 1.02	99.93 ± 0.10	100.00 ± 0.00	100.00 ± 0.35	100.00 ± 0.00
4	59.99 ± 4.19	65.63 ± 3.77	96.38 ± 0.32	100.00 ± 0.00	100.00 ± 0.71	100.00 ± 0.00
5	93.85 ± 0.86	97.39 ± 1.19	100.00 ± 0.00	100.00 ± 0.00	100.00 ± 1.43	100.00 ± 0.00
			FAS-nB			
Patient Type	**(cumulative) % examined on day ...**					
	1	2	5	7	23	25
1	0.00 ± 0.00	0.00 ± 0.00	0.00 ± 0.00	0.00 ± 0.00	37.62 ± 0.00	100.00 ± 0.00
2	0.00 ± 0.00	0.00 ± 0.00	0.00 ± 0.00	0.00 ± 0.00	0.97 ± 0.00	100.00 ± 0.00
3	0.00 ± 0.00	0.00 ± 0.00	0.00 ± 0.00	0.00 ± 0.00	0.88 ± 0.00	100.00 ± 0.00
4	0.00 ± 0.00	0.00 ± 0.00	0.00 ± 0.00	0.00 ± 0.00	0.72 ± 0.00	100.00 ± 0.00
5	0.00 ± 0.00	0.00 ± 0.00	0.00 ± 0.00	0.00 ± 0.00	71.31 ± 0.00	100.00 ± 0.00
			M-nB			
Patient Type	**(cumulative) % examined on day ...**					
	1	2	5	7	23	25
1	99.92 ± 0.12	100.00 ± 0.00	100.00 ± 0.00	100.00 ± 0.00	100.00 ± 0.00	100.00 ± 0.00
2	99.13 ± 0.45	100.00 ± 0.00	100.00 ± 0.00	100.00 ± 0.00	100.00 ± 0.00	100.00 ± 0.00
3	99.65 ± 0.17	100.00 ± 0.00	100.00 ± 0.00	100.00 ± 0.00	100.00 ± 0.00	100.00 ± 0.00
4	97.46 ± 1.39	100.00 ± 0.00	100.00 ± 0.00	100.00 ± 0.00	100.00 ± 0.00	100.00 ± 0.00
5	99.92 ± 0.13	100.00 ± 0.00	100.00 ± 0.00	100.00 ± 0.00	100.00 ± 0.00	100.00 ± 0.00

Table D.2: Booking distribution across MRI scanners per patient type in the heavier loaded system of Section 6.3.

AOP			
Patient Type	% booked on MRI scanner ...		
	1	2	3
1	42.58 ± 4.30	57.42 ± 4.30	0.00 ± 0.00
6	43.53 ± 1.41	26.28 ± 1.08	30.19 ± 1.07
7	62.30 ± 1.28	16.49 ± 0.74	21.21 ± 0.57
8	80.36 ± 1.17	19.64 ± 0.66	0.00 ± 1.40
9	0.00 ± 1.48	40.88 ± 0.49	59.12 ± 0.83
12	29.78 ± 0.71	31.36 ± 1.06	38.86 ± 0.77
13	44.79 ± 0.66	55.21 ± 0.66	0.00 ± 0.00
14	39.41 ± 1.98	60.59 ± 1.98	0.00 ± 0.00
16	0.00 ± 0.00	56.44 ± 0.52	43.56 ± 0.52
20	29.73 ± 1.30	8.08 ± 0.30	62.20 ± 0.85
21	31.46 ± 1.34	16.90 ± 0.41	51.64 ± 0.86
22	34.19 ± 1.60	65.81 ± 1.60	0.00 ± 0.00
23	53.28 ± 0.78	46.72 ± 1.65	0.00 ± 2.38
25	0.00 ± 0.00	39.48 ± 0.66	60.52 ± 0.66
FAS-nB			
Patient Type	**% booked on MRI scanner ...**		
	1	2	3
1	84.76 ± 2.66	15.24 ± 2.66	0.00 ± 0.00
6	30.46 ± 0.61	25.11 ± 0.84	44.43 ± 1.31
7	27.41 ± 0.98	11.83 ± 1.17	60.76 ± 1.23
8	69.44 ± 1.86	30.56 ± 1.63	0.00 ± 1.05
9	0.00 ± 1.63	17.81 ± 0.30	82.19 ± 1.38
12	41.87 ± 1.99	13.99 ± 0.51	44.14 ± 1.91
13	72.62 ± 1.52	27.38 ± 1.52	0.00 ± 0.00
14	79.38 ± 1.44	20.62 ± 1.44	0.00 ± 0.00
16	0.00 ± 0.00	58.82 ± 2.36	41.18 ± 2.36
20	46.46 ± 1.52	23.32 ± 1.47	30.22 ± 1.91
21	44.75 ± 2.01	29.19 ± 1.82	26.06 ± 1.08
22	56.98 ± 1.31	43.02 ± 1.31	0.00 ± 0.00
23	40.93 ± 1.85	59.07 ± 1.78	0.00 ± 3.18
25	0.00 ± 0.00	57.86 ± 1.78	42.14 ± 1.78

(Continued on next page)

Table D.2 – ***Continued***

M-nB			
textbfPatient Type	**% booked on MRI scanner ...**		
	1	2	3
1	89.20 ± 1.34	10.80 ± 1.34	0.00 ± 0.00
6	31.50 ± 0.75	23.06 ± 1.28	45.44 ± 1.57
7	33.15 ± 1.49	13.38 ± 0.85	53.47 ± 2.02
8	69.08 ± 1.21	30.92 ± 1.37	0.00 ± 1.96
9	0.00 ± 0.62	28.06 ± 0.67	71.94 ± 1.01
12	39.07 ± 1.11	19.62 ± 0.41	41.30 ± 1.03
13	67.34 ± 1.10	32.66 ± 1.10	0.00 ± 0.00
14	72.36 ± 1.93	27.64 ± 1.93	0.00 ± 0.00
16	0.00 ± 0.00	57.32 ± 2.03	42.68 ± 2.03
20	37.98 ± 0.36	26.84 ± 1.41	35.18 ± 1.25
21	34.47 ± 1.37	29.02 ± 0.83	36.51 ± 1.09
22	64.44 ± 0.77	35.56 ± 0.77	0.00 ± 0.00
23	40.26 ± 2.25	59.74 ± 2.89	0.00 ± 1.74
25	0.00 ± 0.00	55.38 ± 0.83	44.62 ± 0.83