Master Thesis Applied Mathematics

Optimization of patient planning at a cardiology outpatient clinic

Femke Boelens
Stochastic Operations Research
University of Twente

March 2022

UNIVERSITY OF TWENTE.

gelre ziekenhuizen

UT supervisor:
Prof. Dr. Richard Boucherie

Gelre supervisors:
Ir. Carly Overmars
Dr. Ir. Nardo Borgman

Graduation committee:
Prof. Dr. Richard Boucherie (UT)
Dr. Ir. Anne Zander (UT)
Dr. Antonios Antoniadis (UT)
Ir. Carly Overmars (Gelre)

Preface

This thesis, titled 'Optimization of patient planning at a cardiology outpatient clinic' has been written in order to complete the master Applied Mathematics at the University of Twente. The research started in September 2021 and was completed in March 2022.. During this period, I did research on capacity planning in the cardiology outpatient clinic of Gelre Ziekenhuizen in Zutphen. For this research, I dived into the planning process of the outpatient clinic and I developed mathematical models to optimize this process. For me, applying my mathematical skills and the methods I learned during my study time to a hospital setting was a perfect final project of the master Applied Mathematics. I look back at a period that I really enjoyed and in which I learned a lot.

First of all, I want to thank Gelre Ziekenhuizen for the opportunity to execute my thesis there. For this, I want to thank Hendrik-Jan for bringing me into contact and Arnout for welcoming me in the capacity management team. From the start, I felt really welcome in the hospital, for which I have to thank Nardo in particular. Thanks for smooth start of my research and introducing me to all relevant people in the hospital, but also for your enthusiasm and supervision within the hospital in the first half of my research. This brings me to the next person of Gelre Ziekenhuizen I would like to thank personally, which is Carly. Thank you for your supervision within the hospital in the second half of my research. I really appreciated your interest in what I was doing, but also for the feedback and advice you gave me throughout the process. Furthermore, I would like to thank Sanne, Arjanne and Jamie from the capacity management team too for welcoming me in the team. I also appreciated your enthusiasm about my research and the nice working atmosphere on the days I worked in the hospital. From the cardiology outpatient clinic, I would like to thank Marielle, Amber, Esther, and Pepijn for all the information they provided me with in order to get familiar with the planning process of the outpatient clinic.

A great thanks goes to Richard for his supervision. Thank you for the useful weekly discussions. Your insights, ideas and critical view were helpful during this research. Furthermore, I would like to thank you for offering me a workplace at the university and putting me in contact with the CHOIR group. My thanks also goes to the other members of CHOIR for welcoming me in the group. I enjoyed the group meetings and am looking forward to also do my PhD within this group. I especially would like to thank the people of Zilverling 4006. Eline, Stef and Sander, thanks for the nice working environment, coffee breaks and the advice you gave me during my research.

Furthermore, I would like to thank Anne and Antonios for being members of my graduation committee. Thank you for taking the time to read this thesis and assess it. Anne, I would also

like to thank you for the useful feedback in the last phase of my research.

This thesis also marks the end of my study time in Enschede. The friends that I made during these five and a half years made this time even better. Therefore, I would like to thank Annemarie, Jarco. Jente, Leander, Lotte, Lucas, Nienke, Tessa and Wisse for studying together, but also for all the fun we had in the past five and a half years. Also during this research, meeting all of you was a good distraction. Of the above people, I want to thank Jarco, Lotte and Wisse for reading (parts of) this thesis and providing me with feedback.

Finally, I want to thank my parents. During this research, you were of great support and always interested in what I was doing. I enjoyed staying with you in the part of the week I worked at the hospital. Thanks for all the support and always being there for me.

I hope you enjoy reading this thesis!

Femke Boelens, Enschede, March 2022

Abstract

We develop mathematical models to optimize and support the patient planning at a cardiology outpatient clinic, both in the long-term and the short-term. For the long term, we develop a method for scheduling patients that need a follow-up consult. We do this in such a way that also new patients that will be referred to the outpatient clinic have acceptable access times. We model the arrival and planning process of new patients as a non-stationary infinite horizon Markov Decision Process (MDP). To be able to solve it, we truncate the problem to a finite horizon problem. We approximate the state values at the horizon by the state values of a stationary infinite horizon MDP. We show that the error of this truncation and approximation is bounded. For the admission planning of follow-up patients we develop a stochastic Integer Linear Programming (ILP) model, in which we include the MDP model. The output of the model is an optimal division of capacity between new patients and follow-up patients and optimal decisions regarding the planning of follow-up patients. Based on the output of the model, we develop a method that generates an advice for a cardiologist on the week they should plan a patient that needs a follow-up consult. We perform a simulation to test the performance of our long-term planning method. We compare our method to a myopic method. Simulation results show that using our method yields lower access times than using the myopic method. Furthermore, for the short-term planning, the MDP that models the arrival and planning process of new patients can be used to find the optimal number of appointment slots that need to be reserved for new patients.

Contents

Chapter 1

Introduction

At an outpatient clinic, care is delivered to patients who do not need an overnight stay at the hospital [1]. Patients that go to an outpatient clinic are often referred by the general practitioner. Besides patients that are new to the outpatient clinic, it is also common that patients come back for a follow-up consultation after a predetermined period. For the access time of new patients, standards are set by the Dutch government, which are aggregated in the so-called the Treeknorm. According to these standards, patients that are referred to an outpatient clinic should be seen by a doctor within four weeks [2]. Dividing capacity among the different patients groups in such a way that all patients have acceptable access times is a challenging problem. This is also a problem that is encountered at the cardiology outpatient clinic of Gelre Ziekenhuizen, which is the inspiration of our research. We therefore formulate the following research question.

"How can patient planning at a cardiology outpatient clinic be optimized using mathematical methods in such a way that all patients have acceptable access times?"

Hence, the scope of this thesis is to develop a method that optimizes the process of patient planning and that can be used as a support system for planning decisions. However, before we develop such a method, we first analyze the current situation and planning strategy. Furthermore, we need to have an overview of methods that can be used to model the planning of patients and optimize decision making. Therefore, we formulate the following sub-questions which we aim to answer in this thesis, which we divide into three different categories.

Inventarization of current practise

1. *What is the current planning strategy at the cardiology outpatient clinic of Gelre Ziekenhuizen?*
2. *How well does the current planning strategy perform in terms of access times?*

Mathematical models and methods

3. *Which Operations Research methods that are known in literature can we use to model and optimize the planning process at an outpatient clinic?*
4. *How can we model the arrival and planning process of new patients?*
5. *How can we optimally schedule follow-up patients in advance, while also taking into ac-*

count new patients?

Application and practical results

6. *How well does the method of patient scheduling that we develop perform in terms of access times?*

By answering these sub-questions we can develop a method that is of relevance in multiple ways. Firstly, it produces an approach for dividing the available capacity among the new and follow-up patients. This division is relevant as it can serve as a basis for the weekly blueprint schedule. A blueprint schedule is a template of a schedule in which for each appointment slot it is specified which type of appointment should be planned. Secondly, the output of the method can produce an advice for the long-term planning of a follow-up patient. Therefore, the output can be of relevance in supporting planning decisions. Thirdly, capacity problems can be noticed while applying the method to make the capacity division. It is very relevant to notice these capacity problems far in advance, as it then is more probable they could be prevented.

This thesis is organized as follows. In Chapter 2 we analyze the current planning method of the cardiology outpatient clinic of Gelre Ziekenhuizen. In addition, we execute a data analysis to assess the performance of the current planning method and movivate our research direction. This chapter answers sub-questions 1 and 2. The goal of Chapter 3 is to provide an overview of literature that is relevant for our research. We discuss mathematical methods that are used in appointment planning in healthcare, which answers sub-question 3. Furthermore, we position our contribution in the literature. In Chapter 4 we answer sub-question 4 by formulating a Markov Decision Process (MDP) that models the arrival an planning process of patients that are new to the outpatient clinic. In this chapter we also discuss an approach to solve the MDP, which includes truncating the horizon and approximating the state values at the end of the horizon. An analysis is performed on the error that may arise by this approach. In Chapter 5 we propose a method for the long-term planning of follow-up patients, which are the patients that are known to the clinic and need another consultation after a predetermined period. We formulate this problem as an stochastic Integer Linear Program (ILP). In this model also the MDP of Chapter 4 is used to take the arrival of new patients into account. The model produces an optimal division of capacity among new and follow-up patients and the output can be used to support decisions on planning follow-up patients in advance. We also provide a solution approach for solving this model and discuss the usability of the output of the model. This chapter should answer sub-question 5. To get an indication of the method developed in Chapter 5 in practice and answer the last sub-question, we perform a case study on the cardiology outpatient clinic at Gelre Ziekenhuizen in Chapter 6. We do this by setting up a simulation in which we test our method against a myopic method and compare the access times that are realized in the simulation. In Chapter 7 we discuss the relevant aspects of implementation of our method in the hospital, as well as limitations of our research. We conclude our research in Chapter 8, in which we also answer our research question.

Chapter 2

Analysis of the current situation of the cardiology outpatient clinic

In this chapter, we analyse the cardiology outpatient clinic of Gelre Ziekenhuizen in Zutphen. We start with providing general information about the outpatient clinic, after which we discuss the current planning strategy. Furthermore, we execute a data analysis to assess the performance of the outpatient clinic. We conclude this chapter by motivating our research direction based on the current situation of the outpatient clinic.

2.1 The cardiology outpatient clinic

Gelre Ziekenhuizen is a hospital group consisting of two hospitals in The Netherlands, one in Zutphen and one in Apeldoorn. Furthermore, Gelre Ziekenhuizen has outpatient clinics in Dieren, Doesburg, Epe and Lochem. The two hospitals of Zutphen and Apeldoorn merged in 1999 and currently have more than 3,500 employees to take care of the 350,000 people in the catchment area. They deliver general care for the whole region, but also supra regional care for some medical specialisms [3]. The cardiology outpatient clinic that is the inspiration of this thesis is located in Zutphen.

At the the cardiology outpatient clinic mainly consultations take place with either cardiologists or the physician assistant (PA). Patients that come to the outpatient clinic can be divided into three groups: new patients, follow-up patients and semi-urgent patients. New patients are the patients that have been referred to the outpatient clinic by the general practitioner (GP) or that have an internal or external referral. Internal referrals are referrals from another department of Gelre Ziekenhuizen and external referrals are referrals from other hospitals. Follow-up patients are patients that need to come back on a regular basis for one or multiple tests at the heart function department and a consultation after a predetermined period, for example each half year or each year. Patients that are also characterized as follow-up patients are patients that start a heart rehabilitation program on the bike at the heart function department. After their first session and their last session, they need a consult with a cardiologist. The semi-urgent patients are patients that either were hospitalized and need to have a consultation to check them up, or patients that are known to the outpatient clinic and develop complaints for which they need to be seen earlier than the planned appointment.

The cardiology outpatient clinic in Zutphen consists of 5 cardiologists, one physician assistant, one nurse practitioner in training and several planning assistants. In general, patients also have to do one or multiple diagnostic tests at the heart function department. The tests performed at the heart function department are an electrocardiogram (ECG), placing a monitor to make an ECG for 24 hours (holter), an echocardiogram (echo), a cardiac catheterization and an exercise test on a bike (ergo). Immediately after the test the consultation can take place, as there is no time needed for analysing the result before the consultation. The decision on which diagnostic tests are needed for the new patient is made by the cardiologists by filling in a triage form after the referral is received. On this form they also indicate which components need to be analysed in the blood tests and if the patient can have a consult with the PA instead of a cardiologist.

Besides doing consultations at the outpatient clinic, the cardiologists and the PA also need to work in the cardiology clinic. There should always be one cardiologist in the clinic for patients that are hospitalized and one cardiologist available for emergency patients. Often these tasks are done by the same cardiologist for a whole week. Moreover, each week, the PA is one half-day present in the clinic. Furthermore, cardiologists are present at the outpatient clinics in Lochem and Dieren a few times a week. Other medical activities of the cardiologists include cardiac catheterization (CAG) and peer consultations (ICC) at other departments of the hospital. In addition, trans-esophageal echocardiographs (TEE) have to be made and pacemakers need to be implanted and replaced. Apart from medical activities, cardiologists have administrative tasks and meetings with colleagues, for example to discuss patients. Furthermore, each week supervision sessions between a cardiologist and the PA take place. One of the cardiologists also works for the management of the hospital and another cardiologist works every week on implementing the lean philosophy in the cardiology department.

2.2 The current planning strategy

The planning process of the outpatient clinic includes planning consultations and the corresponding diagnostics at the heart function department. Since the cardiologists and the PA also need to work in the clinic and have other activities as indicated in the previous section, the schedules of the clinic and the outpatient clinic depend on each other. Based on the critical processes that need to be guaranteed, the division of capacity between the clinic and outpatient clinic is determined. The critical processes that are leading in the planning process are the following.

- *For the clinic:*
  - Every day, one cardiologist should be present in the clinic every morning and manage urgent care in the afternoon;
  - Every day, one half-day there should be a CAG programme (by one cardiologist);
  - Every day, it should be possible to make a TEE (by one cardiologist);
  - Three half-days a week, it should be possible to implant pacemakers (by one cardi-

ologist).

– *For the outpatient clinic:*
   – Every day, new patients should be seen.

The cardiologist that is present in the clinic and manages urgent care changes every week. So each week, one of the cardiologists needs to be present in the clinic in the morning and manage urgent care in the afternoon every day. This cardiologist is also the one that can be called during the night to come and the one that has to be in the clinic in the weekend.

One of the critical processes of the cardiology department is related to the outpatient clinic. Namely, it should be possible to have consultations with new patients every day. The reason that the outpatient clinic gives priority to new patients over follow-up patients, is that the medical situation of a new patient is not known yet. In contrast, the medical situation of a follow-up patient is known, and they just need to come back for supervision by the cardiologist. In order to achieve the critical process regarding new patients, every day at least one cardiologist has a 'new patient consultation hour'. This is a half-day in which consultations take place, of which the main part is for new patients. A blueprint for such a morning and afternoon is given in Table 2.1. In this blueprint, the appointment slots containing *NP* are for new patients and the appointment slots containing *FP* are for follow-up patients (controle patiënten in Dutch). Every day, either one of the two takes place. If the above mentioned critical processes are guaranteed, 'follow-up consultation hours' and consultation hours in the outpatient clinics in Lochem and Dieren can be planned if there is still capacity left; that is, if there are still cardiologists available after guaranteeing all critical processes. The 'follow-up consultation hour' is a half-day of consultations of which a significant part is for follow-up patients. A blueprint of a typical morning and afternoon in which there is a 'follow-up consultation hour' is given in Table 2.2.

Table 2.1: Example blueprint for 'new patient consultation hours'.

(a) Morning

08:30	NP
08:45	NP after echo
09:00	FP
09:15	FP after ergo
09:30	NP after echo
09:45	NP after ergo
10:00	FP
10:15	FP after ergo rehabilitation
10:30	NP*
10:45	FP
11:00	Semi-urgent
11:15	Semi-urgent
11:30	FP after ergo for patient *
11:45	NP after echo
12:00	FP after ergo

(b) Afternoon

13:00	FP after ergo rehabilitation
13:15	NP after echo
13:30	NP after echo
13:45	FP after ergo
14:00	FP after ergo rehabilitation
14:15	TEE or FP
14:30	NP after ergo
14:45	NP
15:00	FP after ergo
15:15	Semi-urgent
15:30	NP after echo
15:45	NP after ergo
16:00	Semi-urgent
16:15	NP after ergo

Table 2.2: Example blueprint 'follow-up patient consultation hours'.

(a) Morning

08:30	Semi-urgent
08:45	NP
09:00	FP
09:15	FP after ergo
09:30	NP
09:45	FP after ergo
10:00	FP
10:15	FP after ergo rehabilitation
10:30	NP
10:45	FP
11:00	Semi-urgent
11:15	Semi-urgent
11:30	FP after ergo
11:45	FP
12:00	FP after ergo

(b) Afternoon

13:00	Semi-urgent
13:15	FP after ergo revalidaton
13:30	FP after ergo
13:45	NP
14:00	FP after ergo revalidaton
14:15	FP
14:30	FP after ergo
14:45	NP
15:00	Semi-urgent
15:15	FP after ergo
15:30	Semi-urgent
15:45	FP after ergo
16:00	NP
16:15	FP after ergo

Since October 2017, the outpatient clinic is organized as a one-stop-shop, meaning that it aims to have diagnostics at the heart function department and consultation on the same day, preferably even consecutively. This way, patients have to visit the outpatient clinic only once instead of multiple times. In order to facilitate this, the schedules for consultation hours are aligned with the schedule of the heart function department, where the diagnostic tests take place. As can be seen in the blueprints represented in Tables 2.1 and 2.2, slots are reserved for patients that need to do diagnostic tests. We need to mention that every patient needs to make an ECG before a consult. The *NP* and *FP* slots in the blueprints are for new and follow-up patients, respectively, that only need an ECG before the consultation. The slots for patients that need to make an echocardiogram or a bike test are indicated by *after echo* and *after ergo*, respectively. The *semi-urgent* slots are for patients that for example need to be seen by a cardiologist after having been hospitalized or patients that are known and develop complaints. Furthermore, the slots *FP after ergo rehabilitation* are for patients that either start or end their heart rehabilitation trajectory, which goes along with a consultation with a cardiologist.

To summarize, the planning process starts with the critical processes of the whole cardiology department that should be satisfied. For the outpatient clinic, this means that every week, there are five 'new patient consultation hours'. If there is still capacity left after satisfying these critical processes, 'follow-up consultation hours' are scheduled. Hence, the number of patients that can be helped per week is dependent on the availability of the cardiologists and the PA. The scheduling of patients to specific appointment slots depends on the type of patient they are and which diagnostic tests they need before their consultation.

2.3 Data analysis: performance of the outpatient clinic

In this section, we perform a data analysis in order to assess the performance of the planning of the cardiology outpatient clinic and detect bottlenecks. In Section 2.3.1 we describe the available data and in the subsequent sections we analyse access times of patients, the influence of the patient demand on it and the blueprint compliance. We analyse these factors over two years on a monthly basis.

2.3.1 The data

The data is retrieved from SAS Enterprise Guide, the data warehouse of Gelre Ziekenhuizen. In order to detect effects over years, it is desirable to have data from multiple years. The cardiology outpatient clinic is a one-stop-shop since October 2017. Hence, data that originates from 2017 and earlier is not useful for evaluating the current planning performance. Since the years 2020 and 2021 were during the covid pandemic, we do not consider them to be representative. Therefore, we retrieve data from 2018 and 2019 from SAS EG.

In our data analysis, the first data set we use concerns the referrals to the cardiology outpatient clinic by the GP. GPs can refer patients to the outpatient clinic in the program Zorgdomein. For each referral, we retrieve the Zorgdomein referral number, the patient number, and the referral date. The second data set concerns the appointments. Of all the appointments in the data set, we are interested in consultations with either a cardiologist or the PA. A consultation can be a first consultation or a follow-up consultation. A first consultation is an appointment with type *NP* (new patient consult), *NPZOEK* (new patient consult that is linked to an echocardiogram), *NP BF* (new patient consult for the phycisian assistant), *NPDYSP* (new patient consult that is linked to a consult at the lung outpatient clinic), and *NPERGO* (new patient consult that is linked to an ergo test). Appointment types that are categorised as a follow-up consultation are *FP*, *FPF* (follow-up consult that is linked to an ergo test), *FPFHRE* (a consultation after cardiac rehabilitation on a bike), and *TC FP*. For all appointments that are either a first or a second consultation, we retrieve the patient number, the appointment date, the appointment type, the cardiologist or PA, the location, and any remarks on the appointment. Examples of remarks are the tests that a patient needs to have before the consultation or the follow-up period. The location can either be Zutphen, Lochem or Dieren. The third data set that we use is the consultation hour data set. This data set contains all appointments in the consultation hours. For every appointment in each consultation hour, we retrieve the date and time, the cardiologist (the PA misses in this data set), the scheduled appointment type, the realized appointment type, and the location.

In the referral data set, we find 3155 unique referrals to the cardiology outpatient clinic of Zutphen, of which we can link 2815 to a consultation with a cardiologist at the cardiology outpatient clinic in Zutphen, Lochem or Dieren. A reason for the fact that not all referrals can be linked, is that a patient after the referral has been hospitalized and thus did not go to the outpatient clinic. Another option is that patients get a referral to the cardiology outpatient clinic for an echo or ergo test in the hospital at the heart function department but get the result at the

GP and hence do not have a consultation at the outpatient clinic in Zutphen.

In 2018, at the outpatient clinic 9412 consultations took place. In 2019, this number was 9958. Out of the consultations approximately 21% are new patient consultations, 68% are follow-up patient consultations, and 11% are consultations for semi-urgent patients. Hence, follow-up patients form the largest patient group.

2.3.2 Access times for new patients

The access time for new patients is defined to be the time period between the referral and the first consult at the outpatient clinic. In The Netherlands, standards for the access time are set, which is called the Treeknorm. For an outpatient clinic this Treeknorm prescribes that 80% of the patients must have a first consult within three weeks after the referral and 100% within four weeks. We analyse the performance regarding meeting this standard. For all patient referrals by the GP in the data set that are linked to a consultation, we calculate the access time in days. Note that we cannot do this for internal or external referrals, but only for referrals by the GP, as we only have data of these referrals. Hereafter, we determine which percentage of these patients has an access time of at most 21 days (three weeks) and of at most 28 days (four weeks). These percentages for the years 2018 and 2019 are shown in Table 2.3. The access time of each patient is linked to the month of the referral.

Table 2.3: Percentage of the patients that have access times within three weeks and four weeks in 2018 and 2019.

	Percentage within 3 weeks	**Percentage within 4 weeks**
2018	54.3 %	74.0 %
2019	34.3 %	57.8 %

Since we are not only interested in the mean performance, but also in the fluctuations over the year, we plot the percentages of patients having access times within three weeks and within four weeks per month in the years 2018 and 2019. The result is shown in Figure 2.1. In this figure, we also show the target percentage according to the Treeknorm. Comparing the results to the Treeknorm, we conclude that the Treeknorm is not satisfied.

2.3.3 Access time for follow-up patients

Follow-up patients are patients that need to be seen by a cardiologist after a predetermined period. Cardiologists indicate after how many months they want to see the patient again. The patient is then put on a planning list for the week he or she needs to come back. When this week approaches, the patient is actually scheduled. We define the access time for follow-up patients to be the time period between the indicated day by the cardiologist and the actual date of the consult. We only do this for the patients for whom the follow-up period is indicated in the remarks field in the patient data. Hence, access times are analysed only for a part of the

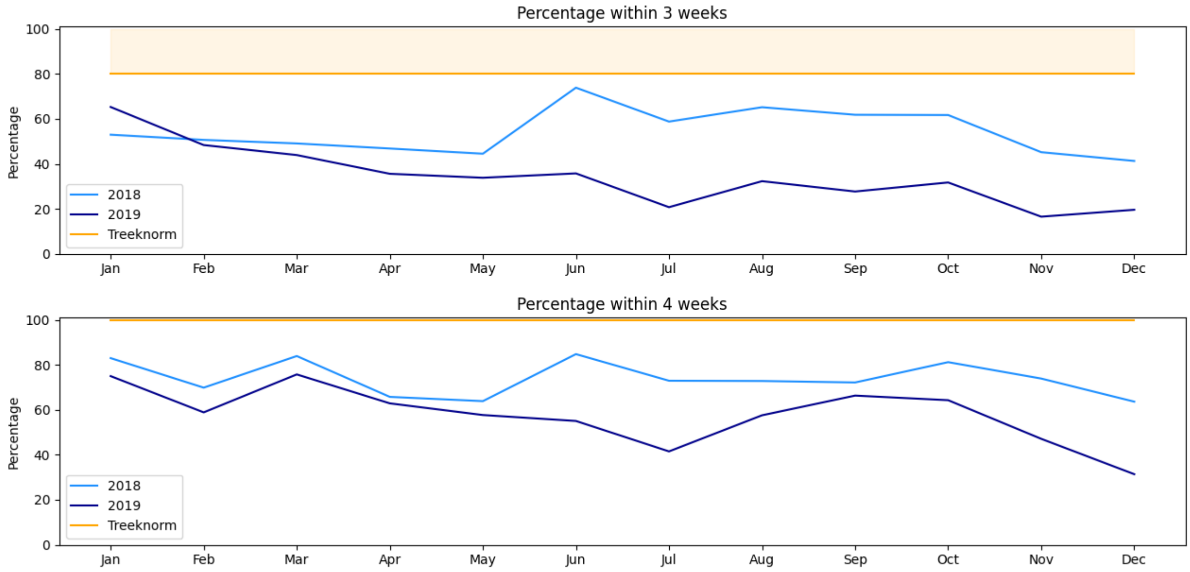

Figure 2.1: Percentage of the patients that have access times within three weeks and four weeks over 2018 and 2019.

follow-up patients. However, by doing this we still can get an indication of the access times and how they change over the year.

For follow-up patients it is less important to have a low access time than for new patients from a medical point of view, as explained before in this chapter. Hence, when a patient for example needs to have a follow-up consultation half a year after their previous consultation, it is also acceptable if this patient has a consultation for example after five or seven months. However, from a patient satisfaction point of view, it is desired that the follow-up period to the next consultation does not differ too much from the follow-up period that was mentioned by the cardiologist. Therefore, we also attach value to the access times of follow-up patients.

In Figure 2.2 we plot the percentage of follow-up patients with access times lower than four weeks. We choose four weeks in order to make the graph similar to the graphs we made for the new patients. Note that for the first part of 2018 we do not have access times for follow-up patients, as the smallest follow-up period that we indicated is 3 months and we need to couple the appointments to the appointment before, which we cannot do if this previous appointment was in 2017 because of data availability.

2.3.4 Influence of patient demand and available capacity on access times

In Figures 2.1 and 2.2 we observe some peaks and troughs in the access times of patients. In this section, we investigate the influence of the patient demand and available capacity on the access times.

For the patient demand, we distinguish between the demand of new patients and follow-up patients and we normalize the demand to the number of days per month. That is, we divide the demand by the number of days in the corresponding month, after which we multiply the result by 30. For the new patient demand, we consider the referrals by the GP, as these referral dates can

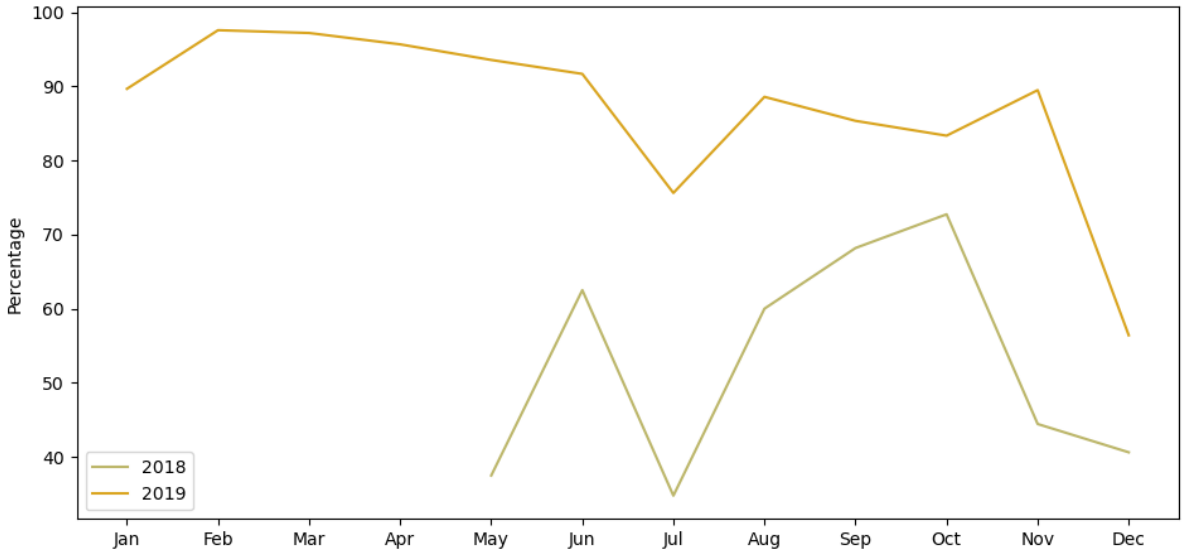

Figure 2.2: Percentage of follow-up patients that have access times within four weeks over 2018 and 2019.

be retrieved from the data. For the follow-up patient demand, we consider the same follow-up patients as for the calculation of the access times. Follow-up periods that are most common in the data are $3, 6$, and 12 months. Although not all follow-up demand is then taken into account, we still can get an impression of its influence on the access times. For the capacity, we take the number of consultation hours and normalize it the same way as the patient demand. We only have consultation hour data and thus of the available capacity of 2019. As measure for the access times of both patient groups, we again take the percentage of the patients that has an access time lower than four weeks. It is good to realize the relation of this measure to the access times. Namely, a higher percentage of patients having access time lower than four weeks, means that the access time at that moment is low. The results of this analysis are shown in Figure 2.3. For the access times we refer to Figures 2.1 and 2.2.

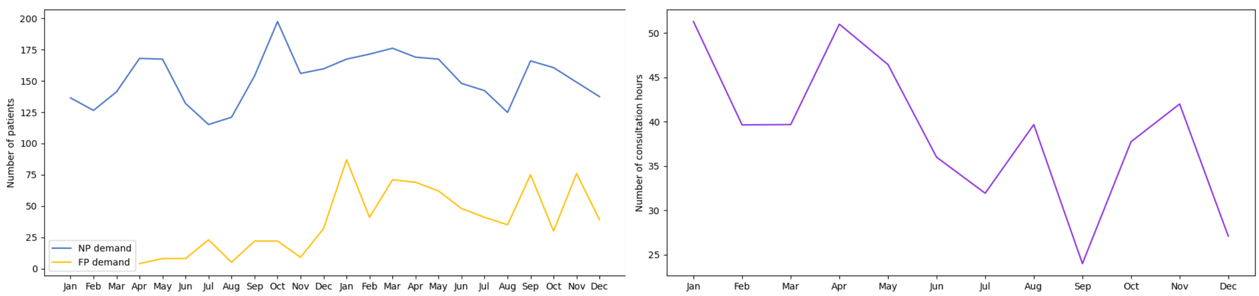

(a) Patient demand for 2018 and 2019.

(b) Capacity for 2019.

Figure 2.3: Patient demand and capacity.

Although it seems like that there are way less follow-up patients than new patients, this is not the case in reality, as we did only take into account the follow-up patients with a known follow-up period. The reason why there are three missing months of follow-up patients and that there are less follow-up patients in 2018 than in 2019 is that we need to link the follow-up appointments to previous appointments. Hence, the follow-up demand of 2018 for example does not include follow-up patients having a follow-up period of a year.

In March and April we see that less patients have access times lower than four weeks. At the same time we observe a peak in new patient demand. Moreover, in October 2018 we also observe a peak in new patient demand and we observe that in November the access time for follow-up patients becomes longer. Furthermore, in January 2019 there is a peak in follow-up patient demand, while the access time for new patients a month later is decreased. Moreover, we see troughs in the percentage having access time lower than four weeks in July and December 2019 and also observe less capacity than average in these months. However, in September, there is also a trough in capacity, while the access times do not seem to suffer from this. This may be caused by the trough in new patient and follow-up patient demand in August.

In both years we observe that the new patient demand is lower than average in January, July, August and December, and higher than average in April, May. In addition, in October 2018 and in September 2019 it is higher than average. We also observe a peak in follow-up demand in July 2018, which can be a consequence of the peak in new patient demand three months earlier. Furthermore, the peak of follow-up patient demand in January 2019 could be a consequence of both the peak in follow-up demand in July 2018 and the peak in new patient demand in October 2018. Another consequence of the new patient demand in October 2018 can be seen in the follow-up patient demand in March 2019 and in November 2019. The peak in March 2019 in its turn may have caused the peak of follow-up patient demand in September 2019.

Moreover, the trough in follow-up patient demand may be a consequence of the trough in new patient demand half a year before. And in its turn this trough may cause the trough in August 2019. However, we see that in August there is more capacity than in September, while the follow-up patient demand has a peak in September. A similar capacity problem can be observed in the summer of 2018. We see a peak of follow-up patient demand in July, while the percentage of follow-up patients with an access time lower than four weeks is low and also the new patient demand is low. Since it is in the middle of the summer, we suspect that in July 2018 there was not enough capacity to serve all follow-up patients. However, in June we see that patients were helped faster than average.

2.3.5 Blueprint adherence

In this section we investigate whether the blueprint schedule actually matches the real schedule. For this, we investigate how many new patients are planned on slots for follow-up patients and vice versa. We are also interested in the number of empty appointment slots per month. We only have this information for 2019. The results are shown in Figure 2.4.

In general, it happens a lot that patients are scheduled on the ‘wrong’ appointment slots. It is also interesting to discuss the relation of patient demand and capacity on it and whether we

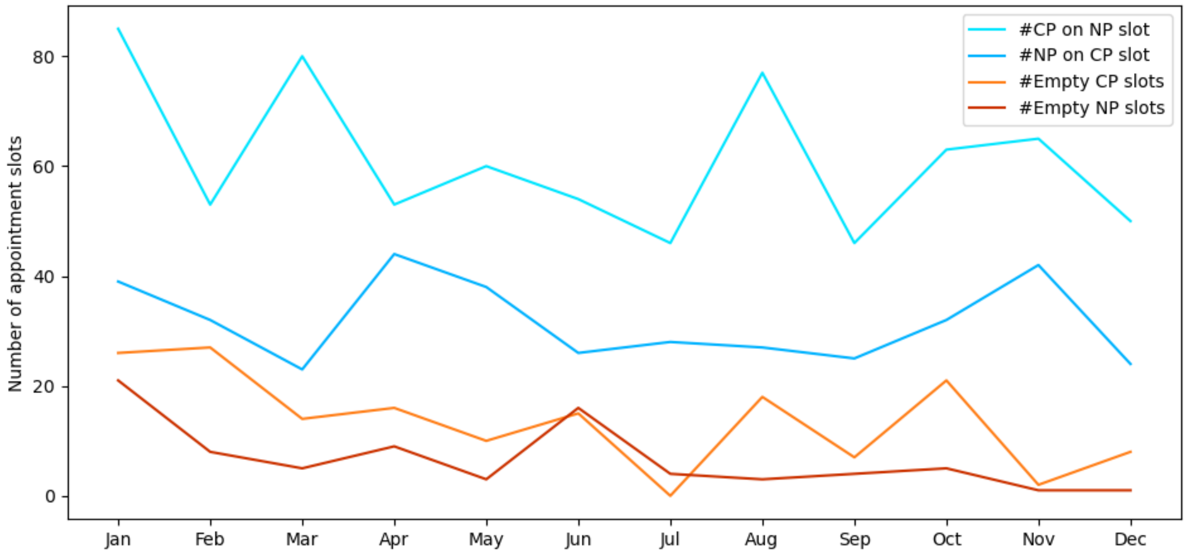

Figure 2.4: Blueprint compliance in 2019.

observe yearly fluctuations in the blueprint compliance. We observe that in January, March and August there were many follow-up patients scheduled on slots reserved for new patients in the blueprint schedule. That this was the case in January and March can be explained by the fact that the follow-up demand was higher than average in these months. In August there where really few new patients, so that may be the reason why slots were offered to follow-up patients. Moreover, in January we also see a lot of empty new patient appointment slots. Therefore, it seems like the division of time slots between new and follow-up patients did not match the patient demand of both groups in January. However, it is remarkable that also many follow-up appointment slots became empty.

Furthermore, we observe that in April many new patients were scheduled on slots for follow-up patients. In April there was a peak in new patient demand, which could be a reason. Furthermore, in November many new patients were scheduled on follow-up patient slots and vice versa. What is also interesting, is that in October there was a peak of empty follow-up patient slots, which clearly is a result of the trough in demand. However, in September and October the new patient demand was high. Hence, the number of slots for new and follow-up patients in the blueprint did not match the demand. Besides the peak in January of empty new patient slots, we also see a peak in June.

2.4 Conclusions and research direction

From the analysis of the current planning strategy of the outpatient clinic, we can conclude that in the blueprint schedule there are slots for semi-urgent patients, new patients and follow-up patients and the slots correspond to the diagnostics that patients have before the consultation. Out of the new patients and the follow-up patients, new patients have more priority. As a consequence, the goal of the outpatient clinic is to offer every week the same fixed amount of new patient consultation slots. Because of this, the number of follow-up patient consultations that can be done is dependent on the available capacity.

From the data analysis, we observe that this way of dividing appointment slots between new patients and follow-up patients might not be optimal, as throughout the year 2019 there were periods in which many new patients were scheduled on follow-up appointment slots in the blueprint but also vice versa in other periods. Furthermore, we observed relations between patient demand and access times, as well as between the available capacity and access times. Moreover, yearly patterns may be recognized in the demand of new patients. However, more clearly is the predictability of follow-up patient demand based on earlier patient demand of both new and follow-up patients. Since approximately two-third of the consultations is a follow-up consultation, this predictability is very useful.

Since dividing slots among the patient groups seems to have room for improvement and the patient demand of the groups is predictable, the focus of this thesis lies in the development of a method that optimizes decisions regarding dividing the capacity among the patient groups. Moreover, this method supports planning decisions as a consequence of the optimal division. We research if such a method results in equal access times over the year and whether it is feasible to have less peaks and troughs in the workload when using such a method.

Although three patient groups have consultations at the outpatient clinic, we mainly focus on the new patients and the follow-up patients. Throughout the thesis, we assume that part of the capacity is reserved for semi-urgent patients in order to satisfy their demand. In the method that we develop and research, we also take into account the access time norms for new patients set by the government, meaning that we aim to schedule all new patients within four weeks after their referral. Within those four weeks, we can have some freedom in scheduling, as long as the patients are scheduled within four weeks. Furthermore, we use the flexibility we have in planning follow-up patients and take into account that the actual week of consultation is not too far away from the week that is communicated by the cardiologist at the previous consultation. Hence, we aim to have a good indication of the week the follow-up consultation will take place already at the previous consultation.

Chapter 3

Literature

In this section, we give an overview of the literature that is relevant for our research. We first provide an overview of the state of the art of appointment planning in outpatient clinics. Moreover, we discuss the relevant literature of appointment planning in cardiology outpatient clinics specifically in Section 3.1.1. We also discuss the literature on patient admission planning and advance patient scheduling in more detail in Section 3.1.2. In Sections 3.2.1 and 3.2.2 we place the mathematical models and methods that we use in the literature and in Section 3.3 we conclude this chapter by summarizing our contribution to the literature.

3.1 Planning in outpatient clinics

For an overview of various aspects of appointment planning in outpatient clinics, we refer to [1] and for an overview of Operations Research and Management Science methods in healthcare, we refer to [4]. Furthermore, for an overview of optimization studies in outpatient clinics, we refer to [5]. Mathematical models and methods that are mentioned in these overviews are Computer Simulation, Heuristics, Markov Decision Processes and Markov Reward Processes, Queueing Theory, and Mathematical Programming and Stochastic Programming. Models that we use in our research, are Markov Decision Processes and Mathematical Programming, of which we study relevant literature in more detail in Sections 3.2.1 and 3.2.2, respectively.

In [4] a taxonomic classification of planning decisions in healthcare is made. A distinction is made between strategic, tactical and operational planning. In addition, the taxonomy also includes a distinction between different health care services. Outpatient clinics fall under ambulatory care services, as patients leave an outpatient clinic on the same day as they arrive. In this research, we focus on the tactical planning level of planning in outpatient clinics. Planning decisions on this level in an outpatient clinic include patient routing, capacity allocation, temporary capacity changes, the access policy, admission control and scheduling appointments. In Section 3.1.2, we discuss the relevant literature on decision making on a tactical level. Decisions on a strategical level would include dimensioning and case mix decisions. On an operational level, decisions would concern patient-to-appointment assignment and the rescheduling of patients and staff. These decisions are short-term decisions that are related to the execution of the health care delivery process. Strategic decisions can be seen as an input for tactical decisions and tactical decisions form guidelines to operational planning decisions, which are short-term decisions on the execution of the care delivery process.

Important insights that are mentioned in [1] that are of relevance for our research include trying to smooth patients over the upcoming period and trying to balance the workload. For the first, one should find patterns in patient arrivals and use these to determine the required capacity. An important aim in appointment planning is to reduce the variance in the process. Furthermore, a few challenges in healthcare planning are indicated. An example that is relevant for our research is the challenge of determining which part of the capacity should be dedicated to patients who need shorter access times and which part to patients who require an appointment in the long term.

3.1.1 Cardiology outpatient clinics

In this section we focus on literature on appointment planning in cardiology outpatient clinics. Currently, the cardiology outpatient clinic of Gelre Ziekenhuizen is organized as a one-stop shop. In [6] it is researched whether a one-stop shop cardiology outpatient clinic is clinically feasible. In a one-stop shop the aim is to perform the necessary tests on the same day as the consultation with the cardiologists so that the patient needs to visit the outpatient clinic only once. It was found that the one-stop shop policy is feasible and robust for a cardiology outpatient clinic and reduced the number of follow-up visits. Furthermore, a one-stop shop policy in a cardiology outpatient clinic is proved to be beneficial for both patients and physicians.

In [7] a model for multi-appointment scheduling in a cardiology outpatient clinic is developed. The model aims to schedule patients that need to undergo several steps in order to complete their procedure at the outpatient clinic. This is done using a linear programming approach. They are the first to research multi-appointment scheduling for cardiology outpatient clinics. Research on multi-appointment scheduling was done before for general outpatient settings. For a recent literature overview on multi-appointment planning in healthcare, we refer to [8].

In [9, Chapter 6], also research is done on multi-appointment planning and a case study is executed on a cardiology outpatient clinic. The problem of online multi-appointment scheduling is addressed. For the decisions on acceptance of newly arrived patients, a Markov Decision Process model is used and for the scheduling of patients to appointments, an Integer Linear Program is developed. In a case study on a cardiology outpatient clinic, a discrete event simulation was executed and it was found that the developed method outperformed a simple heuristic.

In summary, literature on capacity planning in cardiology outpatient clinics is focused on multi-appointment planning, where also diagnostic tests are taken into account. The multi-appointment planning models that were developed in literature focus on an optimal blueprint schedule for a general week. Hence, yearly fluctuations in the demand of different types of patients that may result in time-dependent blueprint schedules are not taken into account.

3.1.2 Tactical planning in outpatient clinics

As indicated in Section 2.4, we research a method that optimizes and supports decisions regarding a division of capacity between the new patients and the follow-up patients. This division is time-dependent, as we observe fluctuations in patient demand and capacity over the year. And in order to give an advice on the long-term planning of follow-up patients such that also new patients can be served, tactical decision making regarding capacity allocation and patient admission planning are important to our research. In this subsection, we give an overview of the literature on these tactical planning decisions.

A method for determining tactical resource and admission plans to cope with fluctuations in patient demand and capacity was developed in [10]. In this research a deterministic method in the form of a Mixed Integer Linear Program was used. In the model several patient groups are taken into account, which all have different priorities of being served within a specific time. In this model, the objective is to mimimize waiting times of the patients and the decision is on how many patients to treat per time period of a certain patient group that have waited a certain amount of time periods. In [11] the ILP from [10] is modelled as a Dynamic Programming (DP) problem and randomness is taken into account. To solve the problem, Approximate Dynamic Programming (ADP) is used. Although fluctuations in demand and capacity are taken into account, fluctuations over the year are not taken into account.

The authors of [12] formulate a bulk service queueing model to optimally assign server time slots to different patient types. Fluctuations in patient demand are taken into account in the sense that the model allows the arrival process during each of the arrival time slots to be different. The objective of the model is to decrease the waiting time of the patients.

In [13] a Stochastic Mixed Integer Programming is developed to find a weekly scheduling template for an outpatient clinic that offers multiple types of services. With their model the optimal number of appointments for each service type in each session can be determined, while balancing the workload and minimizing waiting times, server idle time and overtime. In [14], a planning model is developed to determine the required capacity in an outpatient clinic. Like in our research, this research also takes into account new patients and follow-up patients. Targets on lead-times for both patient groups are set and network flow model is formulated in the form of a Stochastic Linear Programming (SLP) model. To deal with the uncertainty in the demand of patients, chance constraints are used. In case the probability distribution of the arrival of patients is known, the model could be transformed to a deterministic model. The model finds the required capacity by minimizing the maximum required capacity to obtain the targets for lead-times for both patient groups. Another research that considers new patients and follow-up patients is [15]. They build on the model of [14] and develop a robust optimization model that provides a tactical capacity plan for an outpatient clinic. To deal with fluctuations in the patient demand, yearly, seasonal, and monthly uncertainty sets are introduced. Per uncertainty set, the expected demand and deviations are considered. A rolling horizon approach is used, and for each of the finite planning horizon the uncertainty sets corresponding to this period are adopted to model the patient demand.

Most tactical planning models for outpatient clinics in literature focus on a blueprint schedule of a general week and do not consider yearly fluctuations in patient demand and capacity. Furthermore, often same-day patients and pre-scheduled patients are taken into account, as indicated in [5]. The model that comes closest to several aspects of our model is the model presented in [15]. Their model also focuses on new patients and follow-up patients and the seasonal effects in the demand are considered. However, the objective (determine the minimum required capacity) is different from our objective. The problem of finding an optimal capacity division that may differ from week to week such that all types of patients satisfy access time targets is still open.

3.2 Models and Methods

We discuss the literature on the models that we use and also discuss the solution methods. First we review the relevant literature on non-stationary infinite horizon Markov Decision Processes, after which we discuss relevant literature on Mathematical Programming.

3.2.1 Non-stationary infinite horizon Markov Decision Processes

An approach of modelling decision making under uncertainty is modelling the process as a Markov Decision Process (MDP). Together with an optimality criterion, an MDP forms a Markov Decision Problem. An MDP is defined by a a tuple $(\mathcal{T}, \mathcal{S}, \mathcal{A}_s, p_t(\cdot|s,a), r_t(s,a))$. Here, $\mathcal{T}$ is the set of decision epochs in the horizon of the problem and $\mathcal{S}$ is the set of possible states of the system that is modelled. In each decision epoch, a decision can be made based on the state $s \in \mathcal{S}$. The set of possible decisions or actions is denoted by $\mathcal{A}_s$. Moreover, the state at the next decision epoch after being in state s and choosing action a is determined by the transition probabilities $p_t(\cdot|s,a)$, which may be time dependent. The reward that is received after choosing action a in state s is given by the reward function $r_t(s,a)$. If the problem is a minimization problem, the reward is often called the cost. For a more detailed explanation of MDPs we refer to [16].

In this thesis, we model the arrival and planning process of new patients as an MDP. This is a suitable method as we need to make decisions, namely on the planning, under uncertainty, namely that of the patient demand. In our case, the transition probabilities are time-dependent, because of the time-dependent patient demand, which we observed in Figure 2.3. Hence, the process is nonstationary. Furthermore, we are dealing with a infinite horizon, $\mathcal{T} = \{1, 2, \ldots, \infty\}$. The reason is that the arrival of patients does not stop. While planning patients, the future demand has to be taken into account and for planning the future patients, again the demand in the future has to be taken into account. Therefore, the planning horizon for patients does not have a clear end and thus is indefinite. This makes that we need to model the problem as a nonstationary infinite horizon MDP.

However, finding an optimal policy for non-stationary infinite horizon MDPs cannot be done using standard solution methods. Using standard solution methods for non-stationary MDPs such as backward induction the problem will be that there are no state values at the end of the horizon, as the horizon is infinite. Using standard solution methods for finite horizon MDPs such as value or policy iteration will cause problems due to the nonstationarity of the problem. Therefore, another solution approach than the standard approaches is needed. The authors of [17] use a linear programming approach and formulate the problem as a Countably Infinite Linear Program. Another approach to deal with such problems and which we use in this thesis is a rolling-horizon procedure. For each time step a problem is solved with a truncated, finite horizon T. The smallest truncated, finite horizon for which the optimal first decision is the same as for the infinite horizon problem is called a forecast horizon. Only the decision of the first decision epoch is actually implemented, after which the horizon is shifted one epoch and the process is repeated. This way, an infinite policy can be found recursively. In [18], an error analysis is executed and convergence results are shown in case of bounded rewards and value functions. The error bound that was found is minimized if the values of the state in the last decision epoch, which are called salvage values, are chosen to be equal to zero. In [19], the authors indicate that it is possible to find a forecast horizon, but that finding such an horizon comes with practical difficulties. Namely, all data must first be forecasted over the infinite horizon. However, an appropriate horizon length can result in near optimal solutions. The error of choosing an horizon as the truncated horizon is also analysed by the authors. A second difficulty of (finding) a large enough horizon as forecast horizon, is the computational cost due to the horizon being so large that it forms a bottleneck, as indicated in [20]. Also in this case, truncation of the horizon is desired.

For the salvage values, a common choice is to let them equal zero for all states, as was done in [18] and in [20] this is indicated as a possible choice. However, they also indicate that this choice leads to poor performance. This is also probable in our model, as with a salvage value of zero for each state, planning patients could be postponed till after the truncated horizon. Hence, a different choice for the salvage values is desired. Other possibilities that are mentioned in [20] are to learn the salvage values from the behaviour of an expert solving the MDP or from a reinforcement learning process. According to the authors, a good estimate of the salvage values will cause that truncation of the horizon does not lead to bad policies. The authors truncate a finite horizon and develop task-dependent salvage values, by training a function approximator that maps the parameters of the MDP to value functions. They analyse the error that arises by truncation and using these salvage values and they also perform numerical experiments to show how the suboptimality ratio changes with the truncated horizon. In order to determine this suboptimality ratio, the original problem with large finite horizon has to be solved, which can be computationally expensive.

In this thesis, we approximate the salvage values by the state values of a stationary MDP, with the average values of parameters of the non-stationary MDP. If we want to execute an error analysis, we can compare the Markov Reward Processes that arise from the MDPs together with their respective optimal policies. Comparison of Markov Reward Processes is for example done in [21]. Here, the author compares Markov Reward Processes with an cumulative reward criterion. In [22], the results are extended to the case of discounted rewards. The results in both articles can be used to find an error bound on the difference of two Markov Reward Processes. Furthermore, in [23] the error is analysed of having uncertainty in the transition probabilities. The error is divided in an estimation error and a policy error and a bound is found for the total error.

3.2.2 Mathematical Programming models

According to [5], deterministic models for outpatient clinic planning are almost always formulated as Integer Linear Programming (ILP) models. ILPs are programs of the form

$$\begin{array}{ll} \text{minimize} & \mathbf{c}^{\mathsf{T}}\mathbf{x} \\ \text{subject to} & A\mathbf{x} \leq \mathbf{b} \\ & \mathbf{x} \in \mathbb{Z}^n, \end{array}$$

where $\mathbf{x}$ is the decision vector, and $\mathbf{c}, \mathbf{b}$ are vectors and A is a matrix. We let $\mathcal{X}$ denote the set of all feasible solutions. To incorporate uncertainty, one can make use of stochastic programming models, in which the objective function of the ILP may involve uncertainty, which we denote by ξ. Hence the objective function may be written as a function of this uncertainty, leading to the stochastic optimization problem $\min_{\mathbf{x}\in\mathcal{X}}\{f(x) = \mathbb{E}[F(\mathbf{x}, \xi)]\}$. Such a problem is called a single-stage stochastic program [24]. In the literature, Sample Average Approximation is a well-known method for solving Stochastic Programming problems. This method was introduced in [25].

In this thesis, we model the problem of planning follow-up patients as a stochastic ILP, where the stochasticity lies in the patient demand. While planning follow-up patients, there is a certain amount of freedom in the exact week they are planned. However, they should all fit in the schedule within reasonable time from the week after which their follow-up period ends. Hence, we need to fit them all into the schedule, while satisfying the constraints that tell in which weeks they could be planned. To reach this goal, modelling the problem as an ILP is a suitable method, as in this model one can optimize a certain function while satisfying various constraints.

In literature, MDPs are often solved by formulating them as a Linear Programming (LP) model. For the LP formulation of an infinite horizon MDP we refer to [18] and for the LP formulation of a nonstationary finite horizon MDP we refer to [26]. In this thesis, we include the value function of an MDP into the objective function of an ILP. Our method of doing this is different form the standard LP formulations that we mentioned. We linearize it and as a consequence we add the bounds on the action space and state transitions to the constraints of the ILP. To the best of our knowledge, this is not done in the literature before.

3.3 Our contribution

To conclude this chapter, we discuss our contribution to the current literature. In our method we combine different methods in the field of Operations Research and apply them to the setting of a cardiology outpatient clinic. Our contribution is twofold and can be divided in a practical contribution and a scientific contribution. We first discuss the practical contribution of our research. Hereafter, we discuss the scientific contribution to the literature.

3.3.1 Practical contribution

We develop a new method for a cardiology outpatient clinic that optimizes and supports patient planning on the long term. It takes into account patients that are new to the clinic and follow-up patients that are already known and need a consultation after a predetermined period. The method provides an optimal division of capacity between these patient groups that may differ from week to week. Furthermore, our method generates an advice on the optimal week in which a follow-up patient should be planned.

The method that is developed in this thesis is new to literature on patient planning in cardiology outpatient clinics. For cardiology outpatient clinics, only multi-appointment scheduling models are developed in which also the diagnostic tests previous to the consult are taken into account. Therefore, a method that allocates capacity to different patient groups on a weekly basis is new to the literature of cardiology outpatient clinics.

3.3.2 Scientific contribution

For the general setting of an outpatient clinic models are developed that divide capacity among patient groups. With respect to the existing literature, it is new to minimize the number of patients that cannot be helped within the access time norm, instead of minimize the access time. This gives the planning process more flexibility. What is also new, is that we allow the capacity allocation between the patient groups to differ from week to week, but minimize the variability of this allocation.

Furthermore, in the nonstationary MDP that we formulate to model the arrival and planning process of new patients, we truncate the horizon of the problem. Our contribution to the literature on truncated MDPs is that we approximate the salvage values, the state values at the truncated horizon, by solving a stationary MDP that is closely related to the nonstationary MDP. In addition, in the mathematical programming model which optimally plans follow-up patients, the MDP is also taken into account. We do this by including the expected discounted costs in the objective function and add the bounds on actions and the state transitions as constraints. To the best of our knowledge, this has not been done before.

Chapter 4

Method for planning new patients

In this chapter we propose a method for planning patients that are new to the outpatient clinic. They are either referred by the GP, by another department from the hospital or by another hospital. As described in Chapter 2, norms set by the government on access times should be satisfied. However, within those norms we have the freedom to plan them, meaning that we have a certain amount of flexibility to deal with the variability of the arrival of new patients. Hence, reduce the variability in appointment slots per week as much as possible. At the same time we take the government norms for access times into account, meaning that we aim to schedule new patients within four weeks. In Section 4.1 we present a Markov Decision Process (MDP) that models the new patient arrival and planning process. The output of the model is the optimal number of appointment slots that should be in the blueprint schedule each week. Since the model is a nonstationary infinite horizon MDP, it cannot be solved using standard solution methods. We truncate the horizon of the MDP in Section 4.2. In Section 4.3 we provide an error bound analysis of the approximation error due to truncation of the horizon in Section 4.3.

4.1 Model

In this section we formulate a model that we use for determining the optimal number of time slots per week needed for new patients. The problem involves decision making under uncertainty, as the number of referred patients is uncertain. The problem of planning patients does not have a known finite horizon. When planning patients, the future demand always has to be taken into account and since the arrival of patients has no predetermined end, this means that the horizon of our problem is finite. Furthermore, the expected number of patients and capacity may change over time. Hence, we formulate the problem as a non-stationary infinite horizon MDP.

We assume that patients arrive according to a Poisson process Λ_t with time-dependent parameter λ_t for week t and that the capacity for new patients is c_t^{np} time slots per week. This capacity is determined by the total capacity and the capacity that is taken by other patient groups. As mentioned, in this thesis we focus on yearly fluctuations in patient demand. Hence, the time-dependent parameter λ_t depends on the week of the year. Furthermore, we assume that new patients are scheduled on a first come first served basis based on the week of referral.

Table 4.1: List of sets, parameters and variables in the MDP.

Set	Description
$\mathcal{T}_\infty = \{1, 2, \ldots, \infty\}$	Set of weeks in the horizon, indexed by t.
Parameter	**Description**
λ_t	The expected demand of new patients in week t.
c_t^{np}	The capacity for new patients in week t
Variable	**Description**
$s_{t,i}$	State variable, number of patients that wait for i weeks in week t. For $i = 4$, it is the number of patients waiting *at least* for four weeks, must be integer.
a_t	Decision variable, number of slots we reserve in week t for new patients, must be integer.

So if we decide on reserving slots for a certain amount of new patients, we assume that first the patients with the oldest week of referral are scheduled, then the patients with a referral that is one week newer, and so on. We now formulate this MDP by providing the decision epochs, states, actions, transition probabilities and direct costs.

4.1.1 Decision epochs

As mentioned before, the horizon is infinite. Decisions are made once a week. Hence, the set of decision epochs is the set of weeks $\mathcal{T}_\infty = \{1, 2, \ldots, \infty\}$, which we index by t. We assume that the decision for week t is made at the start of week t.

4.1.2 State space

The state $\mathbf{s}_t$ of the system in week t should include information about the number of new patients that need to be scheduled and also the time period since their referral. We refer to this time period as the access time. Therefore, the state of the system at time t takes the form

$$\mathbf{s}_t = (s_{t,0}, s_{t,1}, s_{t,2}, s_{t,3}, s_{t,4})_{t \in \mathcal{T}_\infty},$$

where $s_{t,i}$ is the number of patients that got referred i weeks ago at the beginning of week t. So this patient 'waits' for i weeks in week t. Note that $s_{t,0}$ is the number of referrals that came in between week $t-1$ and week t: these patients can be planned from week t. Furthermore, $s_{t,4}$ is the number of patients that waits at least 4 weeks. Since scheduling of patients is done on a first come first served basis, we must have that there cannot be patients waiting for $i+1$ weeks if no patients are waiting i weeks at that moment (for $i = 0, 1, 2, 3$). Hence, the state space is given by

$$\mathcal{S} = \left\{ (\mathbf{s}) \;\middle|\; \begin{array}{c} s_{t,1} = 0 \text{ if } s_{t,0} = 0 \quad \forall t \in \mathcal{T}_\infty; \\ s_{t,2} = 0 \text{ if } s_{t,1} = 0 \quad \forall t \in \mathcal{T}_\infty; \\ s_{t,3} = 0 \text{ if } s_{t,2} = 0 \quad \forall t \in \mathcal{T}_\infty; \\ s_{t,4} = 0 \text{ if } s_{t,3} = 0 \quad \forall t \in \mathcal{T}_\infty; \\ (\mathbf{s}_0, \mathbf{s}_1, \mathbf{s}_2, \mathbf{s}_3, \mathbf{s}_4) \in \mathbb{N}_0^\infty \times \mathbb{N}_0^\infty \times \mathbb{N}_0^\infty \times \mathbb{N}_0^\infty \times \mathbb{N}_0^\infty \times \mathbb{N}_0^\infty \end{array} \right\},$$

where we use the notation $\mathbb{N}_0 = \{0, 1, 2, \ldots\}$ with $\mathbb{N}_0^n$ being the n-dimensional extension.

4.1.3 Action space

At each decision epoch, the task is to decide on the number of new patients to schedule in that week. We define the vector of possible actions to be

$$\mathbf{a} = (a_t(\mathbf{s}))_{t \in \mathcal{T}_\infty},$$

with $a_t(\mathbf{s}_t)$ the number of patients that is scheduled in week t when the system is in state $\mathbf{s}_t$. Note that we do not decide on which patient is scheduled in which week, we only decide on the number of slots we need to reserve to satisfy the demand of new patients. Because of our assumption that we schedule on first come first served basis, we can derive how to schedule the patients.

The action must satisfy a few constraints. First of all, we cannot exceed the capacity c_t^{np} for new patients in week t. Therefore, we need to have $a_t(\mathbf{s}_t) \leq c_t^{np}$ for each week t in the horizon. Secondly, we cannot schedule more patients than the number of patients that are referred to the outpatient clinic and not scheduled yet and thus we do not reserve more appointment slots than the number of patients known in the system. Thirdly, the number of patients that we schedule needs to be a positive integer. Therefore, the action space for a state $\mathbf{s} \in \mathcal{S}$ is given by

$$\mathcal{A}(\mathbf{s}) = \left\{ (\mathbf{a}) \;\middle|\; \begin{array}{ll} a_t(\mathbf{s}_t) \leq \min\left\{ \sum_{i=0}^{4} s_{t,i}, c_t^{np} \right\} & \forall t \in \mathcal{T}_\infty \\ a_t(\mathbf{s}) \in \mathbb{N}_0 & \forall t \in \mathcal{T}_\infty \end{array} \right\}.$$

4.1.4 Transition probabilities

Once a decision is made, the state of the system changes. Suppose we are in state $\mathbf{s}_t$ in week t and take the action $a_t(\mathbf{s}_t))$, then the system moves to a new state $\mathbf{s}_{t+1}$ in week $t+1$. This state transition is stochastic, as it depends on the number of referred patients per week, which is a random variable. Therefore, the transition from $s_{t,0}$ to $s_{t+1,0}$ is stochastic, while transitions for the other state variables are deterministic.

From our assumption of the order of scheduling new patients, we can derive the deterministic transitions. We first schedule patients that got referred at least four weeks ago, then three weeks ago and so on. Therefore, the number of patients that got referred i weeks ago and is not scheduled in week t will be the number of patients that got referred $i+1$ weeks ago in week $t+1$. If patients that got referred three or at least four weeks ago are not scheduled in week t, which is the case if $a_t < s_{t,3} + s_{t,4}$, they form the patients that got referred at least four weeks ago in week $t+1$. Similarly, the patients that got referred two weeks ago are not scheduled in week t, which is positive if $a_t < s_{t,2} + s_{t,3} + s_{t,4}$, they are the patients that got referred three weeks ago in week $t+1$. The transitions to $s_{t+1,1}$ and $s_{t+1,2}$ are similar. Hence, given the action $a_t(\mathbf{s}_t)$ in week t, the deterministic transitions are as follows

$$s_{t+1,1}|(\mathbf{s}_t, a_t) = [s_{t,0} - [a_t - s_{t,1} - s_{t,2} - s_{t,3} - s_{t,4}]^+]^+, \tag{4.1}$$

$$s_{t+1,2}|(\mathbf{s}_t, a_t) = [s_{t,1} - [a_t - s_{t,2} - s_{t,3} - s_{t,4}]^+]^+, \tag{4.2}$$

$$s_{t+1,3}|(\mathbf{s}_t, a_t) = [s_{t,2} - [a_t - s_{t,3} - s_{t,4}]^+]^+, \tag{4.3}$$

$$s_{t+1,4}|(\mathbf{s}_t, a_t) = [s_{t,3} + s_{t,4} - a_t]^+. \tag{4.4}$$

There is only uncertainty in the transition from $s_{t,0}$ to $s_{t+1,0}$. The transition probabilities for this transition are according to the time dependent probability distribution of the number referrals Λ_t in week t. Since the number of referrals is time dependent; that is, it is distributed with time dependent parameter λ_t, the transition probabilities are time dependent.

$$\mathbb{P}(s_{t+1,0} = j | s_{t,0} = i, a_t = a) = \mathbb{P}(\Lambda_t = j),$$

which is independent of the number of arriving new patients i in the previous week. We arrive at the following time-dependent transition probabilities.

$$p_t(\mathbf{s}'|\mathbf{s}, a) = \begin{cases} P(\Lambda_t = s'_{t,0}) & \text{if s}' \text{ satisfies (4.1)-(4.4)} \\ 0 & \text{else} \end{cases}$$

for $t \in \mathcal{T}_\infty$.

4.1.5 Direct costs

Since we want to schedule patients within the standards set by the government, we incur costs whenever the standards cannot be met. If patients are scheduled within three weeks, no costs are incurred. If patients have to wait between three and four weeks, we have to incur a cost k_3 per patient as at most 20% of the patients may wait between three and four weeks (patients out of the $s_{t,2}$ patients that are planned in week t (between t and $t+1$)). These costs are incurred for the patients that already waited for two weeks and are not planned in week t. The costs are $k_3[s_{t,2} - [a - s_{t,3} - s_{t,4}]^+]^+$.

For patients that have to wait at least four weeks, much higher costs k_4 are incurred, as we do not want patients to be scheduled after four weeks. Hence, $k_4 >> k_3$. These costs are incurred for the patients that have waited for three weeks and are not planned in week t and the patients that have waited for at least four weeks and are not planned in week t. The costs are $k_4[s_{t,3} + s_{t,4} - a]^+$.

Furthermore, we do not want to have much variability in the offered number of slots. Therefore, we incur costs k_v for deviating from the mean number of slots, which is estimated to be

$\bar{a} = \frac{1}{52}\sum_{t=1}^{52}\mathbb{E}[\Lambda_t]$, the mean number of expected referrals per week. Here the mean is calculated over one year and thus 52 weeks, as we assume that yearly fluctuations are the same each year. Costs are incurred when the number of offered slots ideviates from this mean and thus $k_v|a-\bar{a}|$. So the direct costs are

$$k(\mathbf{s}, a) = k_3[s_{t,2} - [a - s_{t,3} - s_{t,4}]^+]^+ + k_4[s_{t,3} + s_{t,4} - a]^+ + k_v|a - \bar{a}|, \tag{4.5}$$

and are independent of $t \in \mathcal{T}_\infty$. Furthermore, scheduling patients on time is of more importance than minimizing variability, and thus $k_4 >> k_3 >> k_v > 0$.

4.1.6 The optimality equations

We define $v_\gamma^\pi(\mathbf{s})$ to be the expected total cost, discounted with a factor $0 < \gamma < 1$, under the policy π:

$$v_\gamma^\pi(\mathbf{s}) = \mathbb{E}_\mathbf{s}^\pi\left[\sum_{t=1}^{\infty}\gamma^{t-1}k(\mathbf{s}_t, a)\right].$$

We are interested in finding an optimal policy π^*, which is the policy that mimimizes the expected total discounted cost. In order to find an optimal policy, the following optimality equations should be solved.

$$v_t(\mathbf{s}) = \min_{a\in\mathcal{A}(\mathbf{s})}\left\{k(\mathbf{s}, a) + \gamma\sum_{\mathbf{s}'\in\mathcal{S}}p_t(\mathbf{s}'|\mathbf{s}, a)v_{t+1}(\mathbf{s}')\right\} \quad \forall t \in \mathcal{T}_\infty. \tag{4.6}$$

4.2 Solution Approach

The optimality equations (4.6) cannot be solved using standard methods due to nonstationarity and the fact that we are dealing with an infinite horizon. We therefore split up the model into two models: a nonstationary finite horizon MDP having horizon T and a stationary infinite horizon MDP that approximates the process after week $t = T$. We first discuss the truncation of the horizon, after which we present the solution approach for the two models separately.

4.2.1 Truncation of the horizon for the nonstationary infinite horizon MDP

The model to determine the number of appointment slots we need to reserve for new patients is formulated as a nonstationary infinite horizon MDP. This type of problem is not solvable in a finite number of calculations [27]. Therefore, we truncate the horizon to a finite horizon T, also known as a forecast horizon, and solve the problem using a rolling horizon approach. Until week T, we are dealing with a finite horizon nonstationary MDP, which is solvable in a finite number of calculations. To approximate the state values from week T onwards, we solve a stationary infinite horizon MDP. We let the so-called salvage values at week T be the values that we find by solving the stationary infinite horizon MDP and denote these values by $V(\mathrm{s})$ for state $\mathrm{s} \in \mathcal{S}$.

For the nonstationary finite horizon MDP, we introduce the set of decision epochs $\mathcal{T}_T = \{1, 2, \ldots, T\}$, with $T < \infty$. The state space, action space, transition probabilities and direct costs remain the same for this truncated finite horizon. Furthermore, the expected discounted total cost under the policy π is now given by

$$v_1^T(\mathbf{s}) = \mathbb{E}_\pi \left[\sum_{t=1}^{T-1} \gamma^{t-1} k(\mathbf{s}_t, a_t) + \gamma^T V(\mathbf{s}_T) \right].$$

The optimality equations also change due to truncation of the horizon:

$$\begin{aligned} v_t(\mathbf{s}) &= \min_{a \in \mathcal{A}(\mathbf{s})} \left\{ k(\mathbf{s}, a) + \gamma \sum_{\mathbf{s}' \in \mathcal{S}} p_t(\mathbf{s}'|\mathbf{s}, a) v_{t+1}(\mathbf{s}') \right\} \qquad t \in \{1, 2, \ldots, T-1\}, \\ v_T(\mathbf{s}) &= v(\mathbf{s}). \end{aligned}$$

For the stationary infinite horizon MDP, the horizon stays $\mathcal{T}_\infty$. The nonstationary elements of the nonstationary infinite horizon MDP are the new patient demand Λ_t in week t and the capacity for new patients c_t^{np} in week t. In order to make the problem stationary, we need to assume these parameters to be time-independent. Hence, we assume that for every week $t \in \mathcal{T}_\infty$ the new patient demand Λ is Poisson distributed with parameter λ. Furthermore, we assume that the capacity is c^{np} for every week $t \in \mathcal{T}_\infty$. As a consequence, the action space for a state $\mathbf{s} \in \mathcal{S}$ is

$$\mathcal{A}(\mathbf{s}) = \left\{ (\mathbf{a}) \;\middle|\; \begin{array}{ll} a_t(\mathbf{s}) \leq \min \left\{ \sum_{i=0}^{4} s_{t,i}, c^{np} \right\} & \forall t \in \mathcal{T}_\infty \\ a_t(\mathbf{s}) \in \mathbb{N} & \forall t \in \mathcal{T}_\infty \end{array} \right\}.$$

Because of time-independent patient demand, the state transition probabilities are also time-independent and are given by

$$p(\mathbf{s}'|\mathbf{s}, a) = \begin{cases} \mathbb{P}(\Lambda = s'_{t,0}) & \text{if } \mathbf{s}' \text{ satisfies (4.1)-(4.4) with } \mathbf{s}_{t+1} = \mathbf{s}' \text{ and } \mathbf{s}_t = \mathbf{s} \\ 0 & \text{else} \end{cases}$$

for all $t \in \mathcal{T}_\infty$.

The salvage values $v_T(\mathbf{s})$ for the nonstationary finite horizon MDP can be found by solving the Bellman equations of the stationary infinite horizon MDP, which are given by

$$V(\mathbf{s}) = \min_a \left\{ k(\mathbf{s}, a) + \gamma \sum_{\mathbf{s}'} p(\mathbf{s}'|\mathbf{s}, a) V(\mathbf{s}') \right\}, \tag{4.7}$$

4.2.2 Solving the stationary infinite horizon MDP

We first truncate the state space in order to make it finite. We do this by bounding the number of arriving new patients by $S^{\max}$. As a consequence, the state variables $s_{t,0}, s_{t,1}, s_{t,2}$ and $s_{t,3}$ are bounded from above by $S^{\max}$ for all $t \in \mathcal{T}$. Furthermore, we bound the number of patients waiting at least four weeks, $s_{t,4}$, by $S^{4,\max}$ for all $t \in \mathcal{T}$. In order to make bounding this state variable possible, we assume that the system is stable. The values of these bounds can be chosen such that the practical effect is negligible. As a consequence of the finite state space, the action space and direct costs are also bounded. Furthermore, the direct costs and transition probabilities in this infinite horizon problem are stationary. Since also the discount factor satisfies $0 \leq \gamma < 1$, it follows from [16, Theorem 6.2.5a] that there exists a unique optimal solution $V^*(\mathrm{s})_{\mathrm{s}\in\mathcal{S}}$ to the optimality equations (4.7). Furthermore, by [16, Theorem 6.2.10], there exists an optimal deterministic stationary policy, meaning that there exists a policy $a^\infty = (a, a, \ldots)$ that equals the optimal policy π^* and that $V^{a^\infty}(\mathrm{s})_{\mathrm{s}\in\mathcal{S}}$ solves the optimality equations (4.7).

In order to find a stationary $\epsilon-$optimal policy $(a_\epsilon)^\infty$ and an approximation to its value $V^{(a_\epsilon)^\infty}$ we use Value Iteration. Being able to find an $\epsilon-$ optimal policy means that we can find a policy as close as we want to the true optimal policy. The algorithm is given below.

Algorithm 1: The Value Iteration Algorithm [16]

1 Choose V^0, specify $\epsilon > 0$ and set $n = 0$.

2 **for** $\mathrm{s} \in \mathcal{S}$ **do**

3   Compute

$$V^{n+1}(\mathrm{s}) = \min_{a\in\mathcal{A}_\mathrm{s}}\{k(\mathrm{s}, a) + \gamma \sum_{\mathrm{s}'\in\mathcal{S}} p(\mathrm{s}'|\mathrm{s}, a)V^n(\mathrm{s}')\}. \tag{4.8}$$

4 **if** $\|V^{n+1} - V^n\| \geq \epsilon(1-\gamma)/2\gamma$ **then**

5   Update $n = n + 1$ and return to step 2.

6 **else**

7   **for** $\mathrm{s} \in \mathcal{S}$ **do**

8   Choose

$$a_\epsilon(\mathrm{s}) \in \arg\min_{a\in\mathcal{A}_\mathrm{s}}\{k(\mathrm{s}, a) + \gamma \sum_{\mathrm{s}'\in\mathcal{S}} p(\mathrm{s}'|\mathrm{s}, a)V^n(\mathrm{s}')\}$$

        and stop.

By [16, Theorem 6.3.1], a $\{V^n\}$ satisfying (4.8) converges to the optimal value V^*. Furthermore, there is a finite N such that $\|V^{n+1} - V^n\| < \epsilon(1-\gamma)/2\gamma$ holds for all $n \geq N$. And whenever this holds, the stationary policy $(a_\epsilon)^\infty)$ is $\epsilon-$optimal. In addition, $\|V^{n+1} - V^*\| < \epsilon/2$.

4.2.3 Solving the non-stationary finite horizon MDP

We solve the finite horizon MDP by backward induction, of which the algorithm is given below.

By [16, Theorem 4.5.1], the policy $\pi^* := (a_1^*, a_2^*, \ldots, a_{T-1}^*)$ such that $a_t^*(\mathrm{s}_t) \in \mathcal{A}_{\mathrm{s}_t}$ for all $\mathrm{s}_t \in \mathcal{S}$, $t \in \mathcal{T}_T$ is a Markovian deterministic policy. Furthermore, it is optimal and satisfies

$$V_1^{T,\pi^*}(\mathrm{s}) = \sup_{\pi\in\Pi^{\mathsf{HR}}}\{V_1^{T,\pi}(\mathrm{s})\}$$

$$V_t^{\pi^*}(\mathrm{s}_t) = V_t^*(\mathrm{s}_t),$$

Algorithm 2: The Backward Induction Algorithm [16]

Set $t = T$ and $v_T^*(\mathrm{s}_T) = V^*(\mathrm{s}_T)$ for all $\mathrm{s}_T \in \mathcal{S}$.
for $t = T-1, T-2, \ldots, 1$ **do**
  Compute

$$v_t^*(\mathrm{s}_t) = \min_{a \in \mathcal{A}_{\mathrm{s}_t}} \{k(\mathrm{s}_t, a) + \gamma \sum_{\mathrm{s}_{t+1} \in \mathcal{S}} p_t(\mathrm{s}_{t+1}|\mathrm{s}_t, a) v_{t+1}^*(\mathrm{s}_{t+1})\}$$

  for each $\mathrm{s}_t \in \mathcal{S}$ and set

$$a_{\mathrm{s}_t,t}^* = \arg\min_{a \in \mathcal{A}_{\mathrm{s}_t}} \{k(\mathrm{s}_t, a) + \gamma \sum_{\mathrm{s}_{t+1} \in \mathcal{S}} p_t(\mathrm{s}_{t+1}|\mathrm{s}_t, a) v_{t+1}^*(\mathrm{s}_{t+1})\}.$$

where Π^{HR} denotes the set of history dependent randomized strategies.

4.3 Error bound analysis of truncating the horizon

As explained in Section 4.2.1 we truncate the nonstationary infinite horizon problem to a nonstationary finite horizon problem with horizon T, having terminal rewards $v_T(\mathrm{s}_T)$ being equal to the state values of a stationary infinite horizon problem $V(\mathrm{s}_T)$. In this section we provide an analysis on the error that arises as a consequence of the truncation of the horizon and approximating the transition probabilities after this truncated horizon by stationary transition probabilities. First we prove that the error that arises is bounded, after which we introduce a method that may help us finding even a better bound for the error.

Theorem 4.1. *Let $(\mathcal{S}, \mathcal{A}, p_t(\cdot|\mathrm{s}, a), k(\mathrm{s}, a))$ be an infinite horizon Markov Decision Process having nonstationary transition probabilities and let $0 < \gamma < 1$ be the discount factor. Furthermore, the direct costs are stationary and bounded by some $K > 0$. Suppose that for stage $t \geq T$, we approximate the transition probabilities by stationary transition probabilities $p(\cdot|\mathrm{s}, a)$. We define the following variables.*

$$\Delta_k := \max_{\mathrm{s} \in \mathcal{S}} \max_{a_1, a_2 \in \mathcal{A}_\mathrm{s}} |k(\mathrm{s}, a_1) - k(\mathrm{s}, a_2)|.$$

$$\Delta_p := \sup_{t \in \mathcal{T}_\infty} \max_{\mathrm{s} \in \mathcal{S}} \max_{a_1, a_2 \in \mathcal{A}_\mathrm{s}} \sum_{\mathrm{s}' \in \mathcal{S}} \left| p_t(\mathrm{s}'|\mathrm{s}, a_1) - p(\mathrm{s}'|\mathrm{s}, a_2) \right|.$$

We denote the error we make by this approximation from decision epoch T by $e(T)$, which is given by

$$e(T) := \left| v_T^{\pi^*,\gamma}(\mathrm{s}_T) - v^{\tilde{\pi},\gamma}(\mathrm{s}_T) \right|,$$

where $v_T^{\pi^*,\gamma}(\mathrm{s}_T)$ is the value function of the nonstationary problem in week T when using the true optimal strategy $\pi^* = (a_t^*)_{t \in \mathcal{T}_\infty}$. Moreover, $v^{\tilde{\pi},\gamma}(\mathrm{s}_T)$ is the value of the stationary problem in decision epoch T while following the optimal strategy $\tilde{\pi} = (\tilde{a}_t)_{t \in \mathcal{T}_\infty}$ to the stationary infinite *horizon problem:*

$$e(T) \leq \gamma^T \frac{\Delta_k(1-\gamma) + K\Delta_p\gamma}{(1-\gamma)^2}.$$

Proof. The error is given by

$$\begin{aligned}
e(T) := &\left| v_T^{\pi^*,\gamma}(\mathbf{s}_T) - v^{\tilde{\pi},\gamma}(\mathbf{s}_T) \right| \\
= &\Bigg| k(\mathbf{s}_T, a_T^*) + \gamma \sum_{\mathbf{s}_{T+1}} p_T(\mathbf{s}_{T+1}|\mathbf{s}_T, a_T^*) \Big(k(\mathbf{s}_{T+1}, a_{T+1}^*) \\
&+ \gamma^2 \sum_{\mathbf{s}_{T+2}} p_{T+1}(\mathbf{s}_{T+2}|\mathbf{s}_{T+1}, a_{T+1}^*) \left(k(\mathbf{s}_{T+2}, a_{T+2}^*) + \ldots \right) \Big) \\
&- \Bigg(k(\mathbf{s}_T, \tilde{a}_T) + \gamma \sum_{\mathbf{s}_{T+1}} p(\mathbf{s}_{T+1}|\mathbf{s}_T, \tilde{a}_T) \Big(k(\mathbf{s}_{T+1}, \tilde{a}_{T+1}) \\
&+ \gamma^2 \sum_{\mathbf{s}_{T+2}} p(\mathbf{s}_{T+2}|\mathbf{s}_{T+1}, \tilde{a}_{T+1}) \left(k(\mathbf{s}_{T+2}, \tilde{a}_{T+2}) + \ldots \right) \Big) \Bigg) \Bigg| \\
\leq &\left| k(\mathbf{s}_T, a_T^*) - k(\mathbf{s}_T, \tilde{a}_T) \right| + \gamma \sum_{\mathbf{s}_{T+1}} \left| p_T(\mathbf{s}_{T+1}|\mathbf{s}_T, a_T^*) k(\mathbf{s}_{T+1}, a_{T+1}^*) - p(\mathbf{s}_{T+1}|\mathbf{s}_T, \tilde{a}_T) k(\mathbf{s}_{T+1}, \tilde{a}_{T+1}) \right| \\
&+ \gamma^2 \sum_{\mathbf{s}_{T+2}} \sum_{\mathbf{s}_{T+1}} \big| p_{T+1}(\mathbf{s}_{T+2}|\mathbf{s}_{T+1}, a_{T+1}^*) p_T(\mathbf{s}_{T+1}|\mathbf{s}_T, a_T^*) k(\mathbf{s}_{T+2}, a_{T+2}^*) \\
&\qquad - p(\mathbf{s}_{T+2}|\mathbf{s}_{T+1}, \tilde{a}_{T+1}) p(\mathbf{s}_{T+1}|\mathbf{s}_T, \tilde{a}_T) k(\mathbf{s}_{T+2}, \tilde{a}_{T+2}) \big| + \ldots.
\end{aligned}$$

The last step could be made due to the triangle inequality. We now find a bound for each separate term in the above expression for the error. We first observe that the first term is bounded by Δ_k. For the second term, we subtract and add the term $p_T(\mathbf{s}_{T+1}|\mathbf{s}_T, a_T^*)k(\mathbf{s}_T, \tilde{a}_T)$ for each possibility of state $\mathbf{s}_{T+1}$, and obtain

$$\begin{aligned}
&\gamma \sum_{\mathbf{s}_{T+1}} \left| p_T(\mathbf{s}_{T+1}|\mathbf{s}_T, a_T^*)(k(\mathbf{s}_{T+1}, a_{T+1}^*) - k(\mathbf{s}_{T+1}, \tilde{a}_{T+1})) - (p(\mathbf{s}_{T+1}|\mathbf{s}_T, \tilde{a}_T) - p_T(\mathbf{s}_{T+1}|\mathbf{s}_T, a_T^*))k(\mathbf{s}_{T+1}, \tilde{a}_{T+1}) \right| \\
&\quad \leq \gamma \Bigg(\sum_{\mathbf{s}_{T+1}} \left| p_T(\mathbf{s}_{T+1}|\mathbf{s}_T, a_T^*) k(\mathbf{s}_{T+1}, a_{T+1}^*) - k(\mathbf{s}_{T+1}, \tilde{a}_{T+1})) \right| \\
&\qquad + \sum_{\mathbf{s}_{T+1}} \left| (p(\mathbf{s}_{T+1}|\mathbf{s}_T, \tilde{a}_T) - p_T(\mathbf{s}_{T+1}|\mathbf{s}_T, a_T^*)) k(\mathbf{s}_{T+1}, \tilde{a}_{T+1}) \right| \Bigg).
\end{aligned}$$

Using the definition of Δ_k and the fact that the direct costs are bounded by K, we obtain that the second term is bounded by

$$\gamma \Bigg(\sum_{\mathbf{s}_{T+1}} \left| p_T(\mathbf{s}_{T+1}|\mathbf{s}_T, a_T^*)\Delta_k \right| + \sum_{\mathbf{s}_{T+1}} \left| (p(\mathbf{s}_{T+1}|\mathbf{s}_T, \tilde{a}_T) - p_T(\mathbf{s}_{T+1}|\mathbf{s}_T, a_T^*))K \right| \Bigg)$$

Since the sum of probabilities over the outcome space equals one and by the definition of Δ_p, we get that the second term is bounded by $\gamma(\Delta_k + K\Delta_p)$. Finding a bound for the third term is similar. We first subtract and add the term $p_{T+1}(\mathbf{s}_{T+2}|\mathbf{s}_{T+1}, a^*_{T+1})p_T(\mathbf{s}_{T+1}|\mathbf{s}_T, a^*_T)k(\mathbf{s}_{T+2}, \tilde{a}_{T+2})$ for each possibility of state $\mathbf{s}_{T+2}$. By applying the triangle inequality we then obtain

$$\gamma^2 \Bigg(\sum_{\mathbf{s}_{T+2}} \sum_{\mathbf{s}_{T+1}} \big| p_{T+1}(\mathbf{s}_{T+2}|\mathbf{s}_{T+1}, a^*_{T+1})p_T(\mathbf{s}_{T+1}|\mathbf{s}_T, a^*_T)(k(\mathbf{s}_{T+2}, a^*_{T+2}) - k(\mathbf{s}_{T+2}, \tilde{a}_{T+2}))\big|$$
$$+ \sum_{\mathbf{s}_{T+2}} \sum_{\mathbf{s}_{T+1}} \big|(p(\mathbf{s}_{T+2}|\mathbf{s}_{T+1}, \tilde{a}_{T+1})p(\mathbf{s}_{T+1}|\mathbf{s}_T, \tilde{a}_T) - p_{T+1}(\mathbf{s}_{T+2}|\mathbf{s}_{T+1}, a^*_{T+1})p_T(\mathbf{s}_{T+1}|\mathbf{s}_T, a^*_T))k(\mathbf{s}_{T+2}, \tilde{a}_{T+2})\big| \Bigg).$$

The first term in this expression is bounded by $\gamma^2\Delta_k$. The second term of this expression is bounded by

$$\gamma^2 \sum_{\mathbf{s}_{T+2}} \sum_{\mathbf{s}_{T+1}} \big|(p(\mathbf{s}_{T+2}|\mathbf{s}_{T+1}, \tilde{a}_{T+1})p(\mathbf{s}_{T+1}|\mathbf{s}_T, \tilde{a}_T) - p_{T+1}(\mathbf{s}_{T+2}|\mathbf{s}_{T+1}, a^*_{T+1})p_T(\mathbf{s}_{T+1}|\mathbf{s}_T, a^*_T))K\big|.$$

We now subtract and add the term $p(\mathbf{s}_{T+2}|\mathbf{s}_{T+1}, \tilde{a}_{T+1})p_T(\mathbf{s}_{T+1}|\mathbf{s}_T, a^*_T)$ to get that this expression is bounded by

$$\gamma^2 K \sum_{\mathbf{s}_{T+2}} \sum_{\mathbf{s}_{T+1}} \big|(p(\mathbf{s}_{T+2}|\mathbf{s}_{T+1}, \tilde{a}_{T+1})(p(\mathbf{s}_{T+1}|\mathbf{s}_T, \tilde{a}_T) - p_T(\mathbf{s}_{T+1}|\mathbf{s}_T, a^*_T))\big| +$$
$$\big|(p_{T+1}(\mathbf{s}_{T+2}|\mathbf{s}_{T+1}, a^*_{T+1}) - p(\mathbf{s}_{T+2}|\mathbf{s}_{T+1}, \tilde{a}_{T+1}))p_T(\mathbf{s}_{T+1}|\mathbf{s}_T, a^*_T))\big|,$$

which is bounded by $2\gamma^2 K\Delta_p$. Hence, the third term of the error is bounded by $\gamma^2(\Delta_k + 2K\Delta_p)$. Similarly, we can prove that the n^{th} term of the error is bounded by $\gamma^{n-1}(\Delta_k + (n-1)K\Delta_p)$. Hence, we can bound the error as follows.

$$\begin{aligned}
e(T) &\leq \Delta_k + \gamma(\Delta_k + K\Delta_p) + \gamma^2(\Delta_k + 2K\Delta_p) + \ldots + \gamma^{n-1}(\Delta_k + (n-1)K\Delta_p) \\
&= \Delta_k \sum_{n=0}^{\infty} \gamma^n + K\Delta_p \sum_{n=1}^{\infty} n\gamma^n \\
&= \Delta_k \frac{1}{1-\gamma} + K\Delta_p \frac{\gamma}{(1-\gamma)^2} \\
&= \frac{\Delta_k(1-\gamma) + K\Delta_p\gamma}{(1-\gamma)^2}.
\end{aligned}$$

To discount this error to the first week, we have to multiply by γ^T, and hence,

$$e(T) \leq \gamma^T \frac{\Delta_k(1-\gamma) + K\Delta_p\gamma}{(1-\gamma)^2},$$

which concludes the proof. □

From this theorem, we know that the error that arises from truncating the horizon to a certain point T is bounded. Hence, we may deduce that when we increase T, the error fades away. However, the bound that was found is not really tight. For example, when we would have chosen to let the salvage values equal zero, then the bound would have been $\gamma^T \frac{K}{1-\gamma}$, because for every decision epoch after T the costs $k(\mathbf{s}, a)$ are not incurred while they should have. Since these costs are bounded by K, the discounted value of the 'missed' costs is $\gamma^T \frac{K}{1-\gamma}$, which is smaller than the bound that we found. Therefore, we introduce a method that may help us to find a better bound for the error.

The nonstationary infinite horizon MDP, presented in Section 4.1, induces together with the optimal policy π^* a Markov reward process, which we denote by $(\mathcal{S}, \mathbf{p}_t, k(\mathbf{s}, a^*))$. The stationary infinite horizon MDP, presented in Section 4.2.1, also induces a Markov reward process with the policy $\tilde{\pi}$, which we denote by $(\mathcal{S}, \mathbf{p}, k(\mathbf{s}, \tilde{a}))$. From decision epoch T, the first Markov reward process is approximated by the latter. Hence, we can analyse the error we make by this approximation by comparing the two Markov reward processes after decision epoch T. For Markov reward processes with the total reward criterion results regarding the comparison of two Markov reward processes are known. We can extend the results of [21] to the discounted reward criterion, using a result from [22]. The result is in the following theorem.

Theorem 4.2. *Let $\delta_\gamma(\cdot) : \mathcal{S} \to \mathbb{R}$ be a function that depends on the discount factor $\gamma \in (0, 1)$. Suppose that for all $\mathbf{s}_T \in \mathcal{S}$ we have*

$$\left| k(\mathbf{s}_T, a_T^*) - k(\mathbf{s}_T, \tilde{a}_T) + \gamma \sum_{\mathbf{s}'} (p_T(\mathbf{s}'|\mathbf{s}_T, a_T^*) - p(\mathbf{s}'|\mathbf{s}_T, \tilde{a}_T))(v^\gamma(\mathbf{s}') - v^\gamma(\mathbf{s}_T)) \right| \leq \delta_\gamma(\mathbf{s}_T). \quad (4.9)$$

Then the error that arises if we approximate the nonstationary MDP by a stationary MDP from decision epoch T is bounded in the following way.

$$\epsilon(T) := \left| v_T^{\pi^*,\gamma}(\mathbf{s}_T) - v^{\tilde{\pi},\gamma}(\mathbf{s}_T) \right| \leq \frac{\max_{\mathbf{s}_T \in \mathcal{S}} \delta_\gamma(\mathbf{s}_T)}{1-\gamma}.$$

Proof. We can write the error in the following way.

$$\begin{aligned}
\epsilon(T) &= \left| k(\mathbf{s}_T, a_T^*) - k(\mathbf{s}_T, \tilde{a}_T) + \gamma \sum_{\mathbf{s}'} p_T(\mathbf{s}'|\mathbf{s}_T, a_T^*) V_{T+1}(\mathbf{s}') - p(\mathbf{s}'|\mathbf{s}_T, \tilde{a}_T) v(\mathbf{s}') \right| \\
&= \Bigg| k(\mathbf{s}_T, a_T^*) - k(\mathbf{s}_T, \tilde{a}_T) + \gamma \sum_{\mathbf{s}'} p_T(\mathbf{s}'|\mathbf{s}_T, a_T^*) \left(V_{T+1}(\mathbf{s}') - v(\mathbf{s}') \right) \\
&\quad + \gamma \sum_{\mathbf{s}'} \left(p_T(\mathbf{s}'|\mathbf{s}_T, \tilde{a}_T) - p(\mathbf{s}'|\mathbf{s}_T, \tilde{a}_T) \right) v(\mathbf{s}') \Bigg|.
\end{aligned}$$

The last step follows from adding and subtracting $\sum_{\mathbf{s}'} p_T(\mathbf{s}'|\mathbf{s}_T, a_T^*) v(\mathbf{s}')$. Now we recognize $V_{T+1}(\mathbf{s}') - v(\mathbf{s}')$ to be $\epsilon(T+1)$, which cannot larger than $\epsilon(T)$. By the fact that summing probabilities over the outcome space equals one and the triangle inequality, we get

$$\epsilon(T)(1-\gamma) \leq \left| k(\mathbf{s}_T, a_T^*) - k(\mathbf{s}_T, \tilde{a}_T) + \gamma \sum_{\mathbf{s}'} \left(p_T(\mathbf{s}'|\mathbf{s}_T, \tilde{a}_T) - p(\mathbf{s}'|\mathbf{s}_T, \tilde{a}_T) \right) v(\mathbf{s}') \right|.$$

We now further investigate the last term.

$$\begin{aligned}
&\sum_{\mathbf{s}'} \left(p_T(\mathbf{s}'|\mathbf{s}_T, \tilde{a}_T) - p(s'|s_T, \tilde{a}_T) \right) v(\mathbf{s}') \\
&\quad = \sum_{\mathbf{s}' \neq \mathbf{s}_T} \left(p_T(\mathbf{s}'|\mathbf{s}_T, \tilde{a}_T) - p(\mathbf{s}'|s_T, \tilde{a}_T) \right) v(\mathbf{s}') + \left(p_T(\mathbf{s}_T|\mathbf{s}_T, \tilde{a}_T) - p(\mathbf{s}_T|\mathbf{s}_T, \tilde{a}_T) \right) v(\mathbf{s}_T) \\
&\quad = \sum_{\mathbf{s}' \neq \mathbf{s}_T} \left((p_T(\mathbf{s}'|\mathbf{s}_T, \tilde{a}_T) - p(\mathbf{s}'|\mathbf{s}_T, \tilde{a}_T) \right) v(\mathbf{s}') \\
&\qquad + \left(\left(1 - \sum_{\mathbf{s}' \neq \mathbf{s}_T} p_T(\mathbf{s}_T|\mathbf{s}_T, \tilde{a}_T) \right) - \left(1 - \sum_{\mathbf{s}' \neq \mathbf{s}_T} p(\mathbf{s}_T|\mathbf{s}_T, \tilde{a}_T)) \right) \right) v(\mathbf{s}_T)
\end{aligned}$$

$$= \sum_{\mathbf{s}' \neq \mathbf{s}_T} \left(p_T(\mathbf{s}'|\mathbf{s}_T, \tilde{a}_T) - p(\mathbf{s}'|\mathbf{s}_T, \tilde{a}_T) \right) \left(v(\mathbf{s}') - v(\mathbf{s}_T) \right).$$

Since $v(\mathbf{s}_T) - v(\mathbf{s}_T) = 0$, we obtain

$$\sum_{\mathbf{s}'} \left(p_T(\mathbf{s}'|\mathbf{s}_T, \tilde{a}_T) - p(\mathbf{s}'|\mathbf{s}_T, \tilde{a}_T) \right) v(\mathbf{s}') \leq \sum_{\mathbf{s}'} \left(p_T(\mathbf{s}'|\mathbf{s}_T, \tilde{a}_T) - p(\mathbf{s}'|\mathbf{s}_T, \tilde{a}_T) \right) \left(v(\mathbf{s}') - v(\mathbf{s}_T) \right).$$

Hence,

$$\begin{aligned} \epsilon(T) &\leq \frac{\left| k(\mathbf{s}_T, a_T^*) - k(\mathbf{s}_T, \tilde{a}_T) + \gamma \sum_{\mathbf{s}'} \left(p_T(\mathbf{s}'|\mathbf{s}_T, a_T^*) - p(\mathbf{s}'|\mathbf{s}_T, \tilde{a}_T) \right) \left(v(\mathbf{s}') - v(\mathbf{s}_T) \right) \right|}{1-\gamma} \\ &\leq \frac{\max_{\mathbf{s}_T \in \mathcal{S}} \delta_\gamma(\mathbf{s}_T)}{1-\gamma}, \end{aligned}$$

which completes the proof. $\square$

Now, if we can find an expression for a function δ_γ, we can find a bound for the approximation error that we make from decision epoch T onwards. To find the influence of this approximation on the value function in the first decision epoch, we have to discount the error by multiplying it by γ^T. Hence, the error of truncating the horizon to T and approximating the process by the stationary MDP yields an error bound of $\gamma^T \frac{\max_{\mathbf{s}_T \in \mathcal{S}} \delta_\gamma(\mathbf{s}_T)}{1-\gamma}$ for a function δ_γ that satisfies (4.9). Hence, this error bound decreases for increasing T because $\gamma < 1$, which is to be expected that changes in value function farther away have less influence on the first decision epoch than changes closer by.

From (4.9), we observe that this δ_γ function is an upper bound of the absolute value expression. Finding a tight upper bound, can thus help us to sharper bound the error resulting from the approximation of the values at decision epoch T. In the remainder of this section we provide a first analysis in finding the absolute expression on the left-hand side of (4.9), which eventually can help to find a tighter bound to the error.

We need to find a worse case value for this absolute value. If a_T^* would equal $\tilde{a}_T$, then the first part of (4.9) would be zero and because the same action is taken, the same future states $\mathbf{s}'$ can be reached. Therefore, the second part would be determined by a difference in transition probabilities p_T and p.

However, (4.9) is bigger when $a_T^* \neq \tilde{a}_T$ and therefore analysis of this case could be useful for finding an expression for the δ_γ function. In this case, there is a difference in cost functions, as different actions are taken and the costs are dependent on the action, see (4.5). Moreover, as a consequence of different actions that are taken in state $\mathbf{s}_T$, we deduce from the transitions (4.1)-(4.4) that the possible next states of the reward processes are different. Let $\mathcal{S}_{a^*(\mathbf{s}_T)}$ denote the set of states that is reachable from $\mathbf{s}_T$ with action a^* and let $\mathcal{S}_{\tilde{a}(\mathbf{s}_T)}$ denote the set of states that is reachable from $\mathbf{s}_T$ with action $\tilde{a}$. Then we have

$$\begin{aligned} \mathcal{S}_{a^*(\mathbf{s}_T)} &= \{(s_0, s_1^*, s_2^*, s_3^*, s_4^*) \mid 0 \leq s_0 \leq S^{\text{max}}, s_1^* - s_4^* \text{ satisfying (4.1)-(4.4)}\} \\ \mathcal{S}_{\tilde{a}(\mathbf{s}_T)} &= \{(s_0, \tilde{s}_1, \tilde{s}_2, \tilde{s}_3, \tilde{s}_4) \mid 0 \leq s_0 \leq S^{\text{max}}, \tilde{s}_1 - \tilde{s}_4 \text{ satisfying (4.1)-(4.4)}\}. \end{aligned}$$

These sets are disjoint: $\mathcal{S}_{a^*(\mathbf{s}_T)} \cap \mathcal{S}_{\tilde{a}(\mathbf{s}_T)} = \emptyset$. Hence, we get two different summations, one in which only $p_T(\mathbf{s}'|\mathbf{s}_T, a_T^*)$ is positive and one in which only $p(\mathbf{s}'|\mathbf{s}_T, \tilde{a}_T)$ is positive. Furthermore, only the first state variable, $s_{t,0}$ is stochastic. Hence, the transition probabilities only involve the transition to this state variable, and hence the Poisson probability of the arrival of new patients. So we can rewrite the absolute value expression on the left hand side in (4.9) to

$$\Big| k(\mathbf{s}_T, a_T^*) - k(\mathbf{s}_T, \tilde{a}_T) + \gamma \sum_{\mathbf{s}' \in \mathcal{S}_{a^*(\mathbf{s}_T)}} \mathbb{P}(\Lambda_T = s_0')(v^\gamma(\mathbf{s}') - v^\gamma(\mathbf{s}_T)) \\ -\gamma\Big(\sum_{\mathbf{s}' \in \mathcal{S}_{\tilde{a}(\mathbf{s}_T)}} \mathbb{P}(\Lambda = s_0')(v^\gamma(\mathbf{s}') - v^\gamma(\mathbf{s}_T))\Big)\Big|. \tag{4.10}$$

We take a closer look at $v^\gamma(\mathbf{s}') - v^\gamma(\mathbf{s}_T)$, which is the difference in value function of two states of which one is reachable from the other state. Note that this concerns the value function of the Markov reward process $(\mathcal{S}, \mathbf{p}, k(\mathbf{s}, \tilde{a}))$. In [21] this term is called the bias-term. For a state $\mathbf{s}'$ that is reachable from $\mathbf{s}_T$ with action $\tilde{a}_T$, the bias term can be written as

$$v^\gamma(\mathbf{s}') - v^\gamma(\mathbf{s}_T) = k(\mathbf{s}', \tilde{a}_{\mathbf{s}'}) - k(\mathbf{s}_T, \tilde{a}_T) + \gamma \sum_{\mathbf{j} \in \mathcal{S}_{\tilde{a}(\mathbf{s}')}} \mathbb{P}(\Lambda = j_0) v^\gamma(\mathbf{j}) - \gamma \sum_{\mathbf{i} \in \mathcal{S}_{\tilde{a}(\mathbf{s}_T)}} \mathbb{P}(\Lambda = i_0) v^\gamma(\mathbf{i}).$$

Since $\mathbf{s}'$ is reachable from $\mathbf{s}_T$ with action $\tilde{a}_T$, $\mathbf{s}' \in \mathcal{S}_{\tilde{a}(\mathbf{s}_T)}$, we have

$$v^\gamma(\mathbf{s}') - v^\gamma(\mathbf{s}_T) = \big(1 - \gamma \mathbb{P}(\Lambda = s_0')\big) v^\gamma(\mathbf{s}') - k(\mathbf{s}_T, \tilde{a}_{\mathbf{s}_T}) - \gamma \sum_{\substack{i_0 \neq s_0' \\ \mathbf{i} \in \mathcal{S}_{\tilde{a}(\mathbf{s}_T)}}} \mathbb{P}(\Lambda = i_0) v^\gamma(\mathbf{i}).$$

Since $\mathbb{P}(\Lambda = s_0')k(\mathbf{s}_T, \tilde{a}_{\mathbf{s}_T}) = \mathbb{P}(\Lambda_T = s_0')k(\mathbf{s}_T, \tilde{a}_{\mathbf{s}_T}) = k(\mathbf{s}_T, \tilde{a}_{\mathbf{s}_T})$, the cost term of the bias-terms cancels out when substituting them in (4.10). Hence, (4.10) can be rewritten to

$$\begin{aligned} &\Big| k(\mathbf{s}_T, a_T^*) - k(\mathbf{s}_T, \tilde{a}_T) + \gamma \sum_{s_0'=0}^{S^{\max}} \mathbb{P}(\Lambda_T = s_0') \big(1 - \gamma \mathbb{P}(\Lambda = s_0')\big) v^\gamma((s_0', s_1^*, s_2^*, s_3^*, s_4^*)) \\ &\qquad - \gamma \mathbb{P}(\Lambda_T = s_0') \sum_{i_0 \neq s_0'} \mathbb{P}(\Lambda = i_0) v^\gamma((i_0, \tilde{s}_1, \tilde{s}_2, \tilde{s}_3, \tilde{s}_4)) \\ &\qquad - \mathbb{P}(\Lambda = s_0') \big(1 - \gamma \mathbb{P}(\Lambda = s_0')\big) v^\gamma((s_0', \tilde{s}_1, \tilde{s}_2, \tilde{s}_3, \tilde{s}_4)) \\ &\qquad + \gamma \mathbb{P}(\Lambda = s_0') \sum_{i_0 \neq s_0'} \mathbb{P}(\Lambda = i_0) v^\gamma((i_0, \tilde{s}_1, \tilde{s}_2, \tilde{s}_3, \tilde{s}_4)) \Big| \\ = &\Big| k(\mathbf{s}_T, a_T^*) - k(\mathbf{s}_T, \tilde{a}_T) + \gamma \sum_{s_0'=0}^{S^{\max}} \big(1 - \gamma \mathbb{P}(\Lambda = s_0')\big) \Big(\mathbb{P}(\Lambda_T = s_0') v^\gamma((s_0', s_1^*, s_2^*, s_3^*, s_4^*)) \\ &\qquad - \mathbb{P}(\Lambda = s_0') v^\gamma((s_0', \tilde{s}_1, \tilde{s}_2, \tilde{s}_3, \tilde{s}_4)) \Big) \\ &\qquad + \gamma \big(\mathbb{P}(\Lambda_T = s_0') - \mathbb{P}(\Lambda = s_0') \big) \sum_{i_0 \neq s_0'} \mathbb{P}(\Lambda = i_0) v^\gamma((i_0, \tilde{s}_1, \tilde{s}_2, \tilde{s}_3, \tilde{s}_4)) \Big|. \end{aligned}$$

Further analysis of this expression may lead to the exploration of a function $\delta_\gamma(\mathbf{s}_T)$. According to Theorem 4.1, a bound for the error of truncating the horizon at decision epoch T can be found.

Chapter 5

Method for planning follow-up patients

In this section we propose a method for planning follow-up appointments in an outpatient clinic. The method should optimize the process of and serve as a support system for planning of the follow-up patients. Follow-up patients are patients that are known to the outpatient clinic and who need a consultation after a predetermined period. Hence, this group of patients is predictable far in advance, as they follow from existing consultations. Moreover, since the follow-up periods consist of multiple months, the exact follow-up period is not that strict, as explained in Chapter 2. However, it is desirable to have a good prediction of when this consultation will take place and thus plan this patient in a week in which it is probable that the patient fits in the schedule.

We develop a method that can support planning decisions regarding these follow-up patients far in advance in which we use the planning flexibility. For a schematic overview of the method we refer to Figure 5.1 We determine a slot division between new patient appointments and follow-up patient appointments per week, such that both groups of patients fit in the schedule and have acceptable access times. Based on this division and the expectations on the number of patients, we develop a method that provides an advice to cardiologists on the week they should schedule a follow-up patient. This advice for specific weeks in the future should be available from the moment the first follow-up patients can be planned in these weeks, which is far in advance. Hence, the slot division for a specific week also should be made far in advance.

For the capacity division between new patient and follow-up patient appointments, we formulate a stochastic program in Section 5.1. In this stochastic program we also take new patients into account, using the model presented in Section 4.1. In Section 5.2, we first explain how we truncate the horizon of the infinite horizon model, after which we transform the stochastic program into a stochastic Integer Linear Program (ILP). Furthermore, we apply the Sample Average Approximation method (SAA) to our model, which can be used to deal with the stochasticity in the model when solving it. Hereafter, we transform the stochastic program to an ILP that approximates it deterministically. A method that transforms the output of the model to an advice on planning decisions is proposed in Section 5.3.

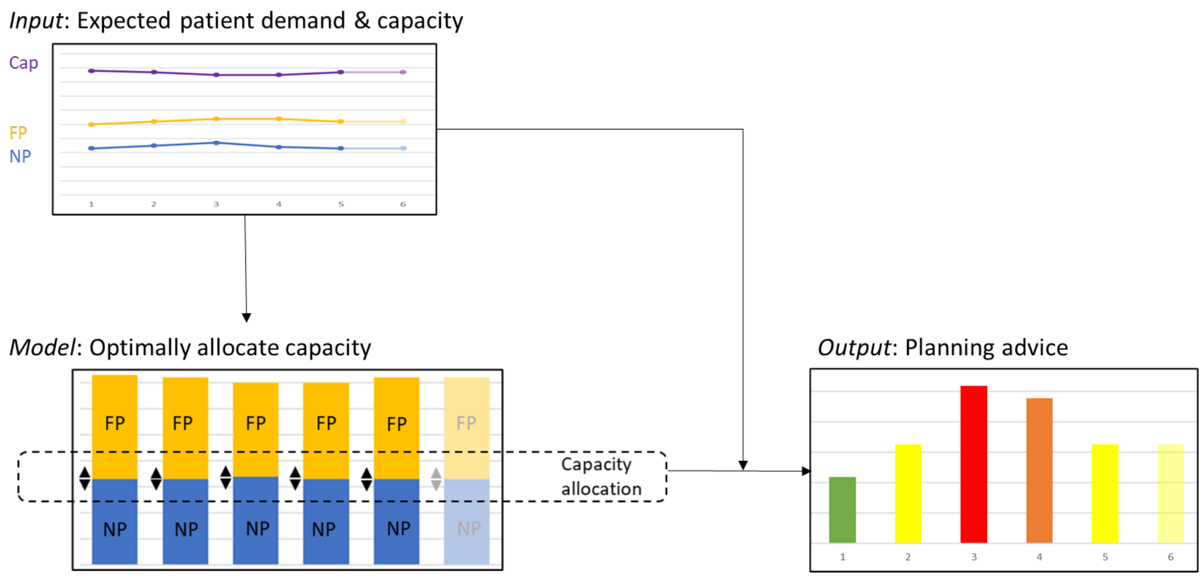

Figure 5.1: Schematic overview of the method.

5.1 Model

Since we want to give an advice on when to schedule follow-up patients, we first need to know how many slots we need to reserve per week for follow-up patients, in such a way that also new patients can have consultations in time. In this section we present a model in which we determine how to optimally schedule follow-up patients. From this, we can derive the number of appointment slots per week needed for the patients that should have a follow-up consultation at the outpatient clinic for a follow-up consultation after a predetermined time period. Although the period is predetermined, we assume, based on the situation of the cardiology outpatient clinic of Gelre, that there is a certain amount of flexibility in the precise moment a patient needs to be seen for a follow-up consultation, which we explained in Chapter 2. We can use this flexibility to schedule the follow-up patients, while taking into account that we also have to schedule new patients, who need to be scheduled soon after their referral and thus with less flexibility.

As explained in Section 4.1, planning patients is an infinite horizon problem. We need to determine how many appointment slots are needed for follow-up patients on a weekly basis. Therefore, we let a decision epoch equal a week and $\mathcal{T}_\infty = \{1, 2, \dots, \infty\}$ be the weeks in the horizon. We introduce the random variable D_t to be the demand of follow-up patients for week t. The demand consists of patients that need a follow-up consultation after a predetermined time after either having had a new consultation, a follow-up consultation, or a semi-urgent consultation. The patients for which the follow-up period ends in week t form the follow-up demand of week t. In the remainder of this report, when we talk about patients 'belonging to week t', we refer to the patients that should come back in week t. We suppose that there are m possibilities of follow-up periods, $T_f = \{t_1, t_2, \dots, t_m\}$ and that the probability of needing a follow-up consultation t_i weeks after a new consultation is q_{t_i}, after a follow-up consultation p_{t_i} and after a semi-urgent consultation is f_{t_i}. We let z_t be the number of appointment slots for follow-up patients that will be decided by our model and we introduce u to be the number of semi-urgent consultations per week, which we suppose to be time independent and thus constant throughout the year. This is a simplification that we make, as we do not take the semi-urgent patients into account. However, semi-urgent consultations may result in follow-up consultations. The demand D_t can be expressed as a sum of binomially distributed variables

$$D_t \sim \sum_{i=1}^{m} \text{Bin}(z_{t-t_i}, p_{t_i}) + \text{Bin}(a_{t-t_i}, q_{t_i}) + \text{Bin}(u, f_{t_i}), \tag{5.1}$$

where we recall that a_t is the number of new patients that we schedule in week t, which follows from solving the MDP described in Section 4.1. Note that z_{t-t_i} and a_{t-t_i} are not realizations of the number of patients scheduled, but previous decisions on reserving appointment slots for the week $t - t_i$. Since we want to give an advice on where to schedule follow-up patients and use this model to determine how many appointment slots we want to reserve for them, we have to run this model before the largest follow-up period that ends in week t begins. So in general, the weeks $t - t_i$ have not passed yet when we run the model.

As mentioned, there is flexibility in scheduling follow-up patients. We suppose we can schedule follow-up patients for week t in the weeks $t - s_1$ till week $t + s_2$. We introduce the variable $x_{t,s}$ with $t - s_1 \leq s \leq t + s_2$, which is a decision variable, representing the number of patients that belong to week t but are scheduled in week s, thus the number of patients we move from week t to week s. The total number of patients scheduled in week s equals $z_s = \sum_{s-s_2 \leq t \leq s+s_1} x_{t,s}$, as patients that belong to the weeks $s - s_2$ to $s + s_1$ can be scheduled in week s. The goal is to decide on the number of slots to reserve each week for follow-up patients, while minimizing the number of patients that cannot be scheduled. Since the demand of patients is a random variable, we want to minimize the number of follow-up patients that cannot be scheduled within the possible weeks. For a sample realization $\mathbf{d} = (d_1, \ldots, d_\infty)$ of the random vector $\mathbf{D} := (D_t)_{t \in \mathcal{T}_\infty}$ this is given by

$$d_t - \sum_{t-s_1 \leq s \leq t+s_2} x_{t,s} \tag{5.2}$$

for all weeks $t \in \mathcal{T}_\infty$.

We assume that the cost of moving patients from week t to week s depends on the number of weeks a patient is moved, $s - t$. We introduce the function $k_m(s-t)$ that represents these costs. Note that a negative value of $s - t$ means that the patient is scheduled earlier than the week they should have a follow-up consultation. We suppose that it is most desirable to plan all patients in the week they belong to, as long as this is feasible in the schedule. We also suppose that cardiologists do not always follow-up the advice of moving a patient to another week, which can, for example, have medical reasons. Therefore, we assume that a cardiologist moves a patient $s - t$ weeks with probability b_{s-t}. Then we let the cost function k_m for moving patients be defined by

$$k_m(s-t) = \begin{cases} \frac{\kappa_m}{b_{s-t}} & \text{if } s - t \neq 0 \\ 0 & \text{if } s - t = 0 \end{cases}, \tag{5.3}$$

Table 5.1: List of sets, parameters and variables in the stochastic program.

Set	Description
$\mathcal{T}_\infty = \{1, 2, \ldots, \infty\}$	Set of weeks in the horizon, indexed by t.
$T_f = \{t_1, t_2, \ldots, t_m\}$	Set of possible follow-up periods between consultations, indexed by t_i.
Parameter	**Description**
c_t	The available capacity in week t for new patients and follow-up patients.
c_t^{np}	The available capacity in week t for new patients, after allocating capacity to follow-up patients: $c_t - z_t$.
p_{t_i}	Probability of needing a follow-up consultation after t_i weeks, after having a follow-up consultation.
q_{t_i}	Probability of needing a follow-up consultation after t_i weeks, after having a new consultation.
f_{t_i}	Probability of needing a follow-up consultation after t_i weeks, after having a semi-urgent consultation.
s_1	Number of weeks we can schedule a follow-up patient before the week they should have a follow-up.
s_2	Number of weeks we can schedule a follow-up patient after the week they should have a follow-up.
k_f	Costs for not planning a follow-up patient.
κ_m	Costs for moving a follow-up patient.
Variable	**Description**
D_t	Random variable, the demand of the follow-up patients in week t.
$x_{t,s}$	Decision variable, number of slots in week s for follow-up patients that need a consult in week t, must be integer.
z_s	Capacity that is allocated to follow-up patients in week s, $z_s = \sum_t x_{t,s}$.

for some cost κ_m and we want to minimize the costs of moving patients. Hence, we want to minimize

$$\sum_{t \in \mathcal{T}_\infty} \sum_{t - s_1 \leq s \leq t + s_2} k_m (s - t) x_{t,s}.$$

Taking this in the objective, ensures that no more patients than necessary will be moved and that for the decision on moving patients also the chance on moving patients in practise is taken into account. As a consequence, as long as it fits, patients are assigned to the week they belong to. Moreover, adding these costs ensures that a solution will not include that patients from week t are moved to week s and also patients from week s are moved to week t. Furthermore, it is also ensured that patients will not be moved to a week for which the follow-up demand is not satisfied. This is because incorporating costs of moving patients will first plan patients in the week they belong instead of planning patients belonging to other weeks in this week. Moreover, by adding this cost to the objective, it will not happen that, for example, all demand is satisfied except for the whole demand of one week. Instead, indirectly the maximum demand that is not satisfied over all weeks will be minimized.

The decision variable $\mathbf{x} = (x_{t,s})_{t\in\mathcal{T}_\infty, s\in\{t-s_1,\ldots,t+s_2\}}$ should satisfy some constraints. We should have that the number of time slots for follow-up patients does not exceed the capacity in week s for follow-up patients. Since patients can be scheduled s_1 weeks before and s_2 weeks after the week t to which they belong, in week s patients belonging to the weeks $s - s_2$ until week $s + s_1$ can be scheduled. Therefore, the following constraint ensures that the capacity in week s will not be exceeded.

$$\sum_{s-s_2\leq t\leq s+s_1} x_{t,s} \leq c_s \qquad\qquad \forall s \in \mathcal{T}_\infty.$$

Moreover, we cannot schedule more patients than the patient demand. The follow-up patient demand D_t in week t is a random variable. We define $\mathbf{d} = (d_1, \ldots, d_\infty)$ to be a scenario of this demand for the whole horizon. For this scenario we should have

$$d_t \geq \sum_s x_{t,s},$$

for all weeks $t \in \mathcal{T}_\infty$.

We also want to take into account the new patient demand when scheduling follow-up patients, so we add a function to the objective function indicating how well the new patients can be scheduled. We choose this function to be the value function $v_\gamma^{\pi*}(\mathbf{s})$ of the MDP described in Section 4.1. This value function represents all discounted costs in the whole horizon $\mathcal{T}_\infty$ when planning new patients according to the policy π. The capacity for new patients can be determined from the total capacity and the capacity that is allocated to follow-up patients, which follows from the decision variables $x_{t,s}$. For week s, the capacity for new patients is given by:

$$c_s^{np}(\mathbf{x}) = c_s - \sum_t x_{t,s}.$$

Therefore, we add $c_s^{np}(\mathbf{x})$ as an argument to the value function and define $v_\gamma^{\pi*}(\mathbf{s}, \mathbf{c}^{np}(\mathbf{x}))$ to be the expected total cost when being in initial state $\mathbf{s}$ under the optimal policy π^*, having capacity $\mathbf{c}^{np}(\mathbf{x}) = (c_s^{np}(\mathbf{x}))_{s\in\mathcal{T}_\infty}$ for new patients and discount factor γ. Consequently, $v_\gamma^{\pi^*}(\mathbf{s}, \mathbf{c}^{np}(\mathbf{x}))$ represents the expected total costs under the optimal strategy, which is denoted by π^*. Hence, we also want to minimize this function.

The objective function that we minimize consists of multiple parts. Firstly, we want to minimize the expected number of follow-up patients that cannot be scheduled, which we do by minimizing the sum of (5.2) over all weeks $t \in \mathcal{T}_\infty$. We give a cost k_f for not scheduling follow-up patients. Secondly, we want to minimize the costs of moving patients from one week to another, which is given by (5.3). Third, we want to minimize the number of new patients that cannot be scheduled and the variability of scheduling new patients, which is given by $v_\gamma^{\pi^*}(\mathbf{s}, \mathbf{c}^{np}(\mathbf{x}))$, the value function of the MDP of Section 4.1. Since in this MDP we discount costs by a factor γ, we also discount the costs for not scheduling follow-up patients.

Hence, the objective function we want to minimize for a scenario $\mathbf{d} \in \mathcal{D}$ is given by the function

$$F(\mathbf{x}, \mathbf{d}) := k_f \sum_{t=1}^{\infty} \gamma^t \Big(d_t - \sum_{s \in \mathcal{T}_{s_1}^{s_2}(t)} x_{t,s} \Big) + \sum_{s=1}^{\infty} \gamma^s \sum_{t \in \mathcal{S}_{s_2}^{s_1}(s)} k_m(s-t) x_{t,s} + v_\gamma^{\pi^*}(\mathbf{s}, \mathbf{c}^{np}(\mathbf{x}))$$

For a follow-up demand scenario $\mathbf{d} = (d_1, \dots, d_\infty)$ of the random vector $\mathbf{D}$ the ILP is given by

$$\min_{\mathbf{x}} \quad k_f \sum_{t=1}^{\infty} \gamma^t \Big(d_t - \sum_{s \in \mathcal{T}_{s_1}^{s_2}(t)} x_{t,s} \Big) + \sum_{s=1}^{\infty} \gamma^s \sum_{t \in \mathcal{S}_{s_2}^{s_1}(s)} k_m(s-t) x_{t,s} + v_\gamma^{\pi^*}(\mathbf{s}, \mathbf{c}^{np}(\mathbf{x})) \tag{5.4a}$$

$$\text{s.t.} \quad \sum_{t=s-s_2}^{s+s_1} x_{t,s} \le c_s \qquad \forall s \in \mathcal{T}_\infty \tag{5.4b}$$

$$\sum_{s=t-s_1}^{t+s_2} x_{t,s} \le d_t \qquad \forall t \in \mathcal{T}_\infty \tag{5.4c}$$

$$x_{t,s} \in \mathbb{N}_0 \qquad \forall t \in \mathcal{T}_\infty,\ s \in \{t - s_1, \dots, t + s_2\}. \tag{5.4d}$$

We denote the set of outcomes of the random vector $\mathbf{D}$ by $\mathcal{D}$. For each scenario $\mathbf{d} \in \mathcal{D}$ the above ILP can be solved and an optimal decision $\mathbf{x}^*(\mathbf{d})$ can be determined. The optimal decision is scenario-dependent and thus may differ per scenario. Since the realization of D_t for all weeks t is not known yet at the moment a decision needs to be made, we need to make sure that the decision that we make is robust. We define the union of feasible solutions of the scenarios to be the set

$$\mathcal{X} := \{\mathbf{x}(\mathbf{d}) \mid \mathbf{x}(\mathbf{d}) \text{ satisfies (5.6b)-(5.6d) for } \mathbf{d} \in \mathcal{D}\}.$$

Note that this set contains for all scenarios the feasible solutions. Hence, it contains solutions that could be feasible for some of the scenarios but not for all. We want to minimize the expected objective function over the outcome space $\mathcal{D}$. Let the probability distribution of the random variable $\mathbf{D}$ be $\mathbb{P}_{\mathbf{D}}$. The stochastic program that we need to solve in order to determine an optimal decision is given by

$$\min_{\mathbf{x} \in \mathcal{X}} \ \mathbb{E}_{\mathbb{P}_{\mathbf{D}}}[F(\mathbf{x}, \mathbf{D})]. \tag{5.5}$$

5.2 Solution Approach

In this section we explain our approach of finding an optimal solution to the mathematical program (5.4). First we make sure that the number of decision variables is finite, which is not the case for the presented mathematical program. We do this by truncating the horizon, which we explain in Section 5.2.1. Moreover, the mathematical program does not have a linear objective function. However, the program can be linearized. We do this in Section 5.2.2. Furthermore, we discuss how we can apply the Sample Average Approximation (SAA) method to solve the resulting stochastic ILP. We conclude this section by presenting a deterministic ILP that approximates the stochastic ILP.

5.2.1 Truncation of the horizon for the mathematical program

The infinite horizon stochastic program of the previous section is not solvable in a finite number of calculations. Moreover, the dependencies of the random variable D_t and previous optimal decisions a_s^* of the MDP and $z_s^* = \sum_{t=s-s_2}^{s+s_1} x_{t,s}^*$ of the stochastic program on each other also cause the problem not to be solvable using existing algorithms. Therefore, we truncate the horizon of this problem. We let the new, truncated horizon be T, and we choose it large enough to have little practical impact on the decision in the first week. As with the MDP of the previous chapter, we use a rolling horizon approach when applying this method. Hence, for every week we run the model with truncated horizon T and only implement the decision of the first week, after which we shift the model horizon one week and repeat the procedure. As a result of truncation of the horizon, the set of weeks in the horizon becomes $\mathcal{T}_T = \{1, 2, \ldots, T\}$.

We note that for the first week, week 1, we are also dealing with variables for which $t < 1$, as patients that belong to an earlier week can be scheduled in the horizon $\mathcal{T}_T$ and patients belonging to a week in this horizon could have been scheduled in a week before the start of this horizon. Therefore, part of the demand of weeks $1 - s_2$ until week s_1 may be satisfied. Information on this can be obtained via previous decisions. Previous decisions are denoted by $x_{t,s}^*$, for $s < 1$. Then the part of the demand that is satisfied for such a week is $D_t - \sum_{s=t-s_1}^{0} x_{t,s}^*$. Of course, the demand for patients belonging to a week $t < 0$ should also be taken into account, as we want the demand of these weeks also to be totally satisfied. Hence, in our objective function we also have to take the demand of these weeks into account.

Furthermore, we have to make assumptions for the end of the horizon. For this, we do not allow that patients belonging to a week later than week T are scheduled before week T and that patients belonging to a week before week T are scheduled after this week. Therefore, we let $x_{t,s}$ only exist if both $s, t \leq T$. Furthermore, by $v_1^T(\mathbf{s}, \mathbf{c}^{np}(\mathbf{x}))$ we denote the value of scheduling new patients in the weeks 1 to T when starting in state $\mathbf{s}$ and having capacity $\mathbf{c}^{np}(\mathbf{x})$ for new patients. Finally, for the sake of notation we define the set $\mathcal{T}_a^b(t)$ to be $\mathcal{T}_a^b(t) := \{\min\{1, t-a\}, \ldots, \max\{T, t+b\}\}$ and $\mathcal{S}_a^b(s) := \{s-a, \ldots, \max\{T, s+b\}\}$. Hence, for a scenario $\mathbf{d} = (d_1, \ldots, d_\infty)$ of the random vector $\mathbf{D}$ the mathematical program with truncated horizon is given by

$$\min_{\mathbf{x}} \quad k_f \sum_{t=1-s_2}^{T} \gamma^t \left(d_t - \left(\sum_{s \in \mathcal{T}_{s_1}^{s_2}(t)} x_{t,s} + \sum_{s=t-s_1}^{0} x_{t,s}^* \right) \right) + \sum_{s \in \mathcal{T}_T} \gamma^s \sum_{t \in \mathcal{S}_{s_2}^{s_1}(t)} k_m(s-t) x_{t,s} + v_1^T(\mathbf{s}, \mathbf{c}^{np}(\mathbf{x})) \tag{5.6a}$$

$$\text{s.t.} \quad \sum_{t \in \mathcal{S}_{s_2}^{s_1}(s)} x_{t,s} \leq c_s \qquad \forall s \in \mathcal{T}_T \tag{5.6b}$$

$$\sum_{s \in \mathcal{T}_{s_1}^{s_2}(t)} x_{t,s} + \sum_{s=t-s_1}^{0} x_{t,s}^* \leq d_t \qquad \forall t \in \mathcal{T}_T \tag{5.6c}$$

$$x_{t,s} \in \mathbb{N}_0 \qquad t \in \mathcal{S}_{s_2}^{s_1}(s), \forall s \in \mathcal{T}_T. \tag{5.6d}$$

For each week, this ILP needs to be solved to obtain an optimal slot division and the optimal values of $x_{t,s}$. The week for which the slot division and the values of $x_{t,s}$ are determined, needs to be the first week of the horizon of the ILP. Since we use a rolling horizon approach, only the decision of the first week of the ILP is implemented, which is the week for which we wanted the slot division and optimal values of $x_{t,s}$. If the slot division of the next week needs to be determined, the planning horizon is 'rolled' one week forward. Then this week is the first week of the ILP.

5.2.2 Linearization of the mathematical program

Since the function $v_1^T(\mathbf{s}, \mathbf{c}^{np}(\mathbf{x}))$ is not linear in $\mathbf{x}$, the stochastic program (5.6) is not linear. The reason is that for finding an optimal policy, we have to enumerate all possible actions, which are bounded by $c_s^{np} = c_s - \sum_t x_{t,s}$ for week s. This enumeration depends on $\mathbf{x}$, making it not linear in $\mathbf{x}$. However, the expression can be linearized, which we show in this section. The costs per week for week $t = 1, 2, \ldots, T-1$ are given by (4.5). When looking at (4.1)-(4.4), we see that we can rewrite this to

$$k(\mathbf{s}, a) = k_4 \cdot s_{t+1,4} + k_3 \cdot s_{t+1,3} + k_v|a - \bar{a}|,$$

for an action $a \in \mathcal{A}_\mathbf{s}$. Hence, when starting in state $\mathbf{s}$, the expected total costs for the weeks 1 to T are

$$v_1^T(\mathbf{s}, \mathbf{c}^{np}(\mathbf{x})) = \sum_{t=1}^{T-1} \gamma^t \left(k_4 \cdot s_{t+1,4} + k_3 \cdot s_{t+1,3} + k_v|a_t - \bar{a}|\right) + \gamma^T V(\mathbf{s}_T)$$

when choosing action a_t in week t. The expressions for $s_{t+1,3}$ and $s_{t+1,4}$ can be found in the state transitions (4.1)-(4.4). Clearly, these expressions are not linear in $\mathbf{x}$ and hence, $v_1^T(\mathbf{s}, \mathbf{c}^{np}(\mathbf{x}))$ is not linear in the decision variable $\mathbf{x}$ of the LP. Therefore, we linearize the expressions, so that we can include the state transitions as the constraints of the LP. We start with the expression for $s_{t+1,4}$ for $t \in \mathcal{T}_T$, for which we linearize the maximum function. For this, we need to add the following constraints for $t = 1, 2, \ldots, T-1$ to the model.

$$\begin{aligned}
&s_{t+1,4} \geq 0 \\
&s_{t+1,4} \geq s_{t,3} + s_{t,4} - a_t \\
&s_{t+1,4} \leq B_4(1 - y_{t,4}^s) \\
&s_{t+1,4} \leq s_{t,3} + s_{t,4} - a_t + B_4 y_{t,4}^s,
\end{aligned}$$

for binary variable $y_{t,4}^s \in \{0, 1\}$ and $B_4 \in \mathbb{R}$ such that all possible values of $s_{t+1,4}$ are at most B_4. The first two constraints are to ensure that $s_{t+1,4}$ is bigger than or equal to the the terms inside the maximum function on the right-hand side and the last two constraints are to ensure that $s_{t+1,4}$ is less than or equal to one of them, so that $s_{t+1,4}$ will be exactly equal to the maximum function on the right hand side for all weeks t. In this expression the action a_t in week t occurs. The action space $\mathcal{A}_\mathbf{s}$ is bounded by the minimum of the number of patients that are referred and the capacity for new patients. The latter depends on the decision variable $\mathbf{x}$ of the stochastic program, namely in week t we have $c_t^{np} = c_t - \sum_{\tau \in \mathcal{S}_{s_2}^{s_1}(t)} x_{\tau,t}$. We can implement the bounds on the action space by adding the following constraints for $t = 1, 2, \ldots, T-1$ to the program.

$$a_t \leq \sum_{i=0}^{4} s_{t,i}$$
$$a_t \leq c_t - \sum_{\tau} x_{\tau,t}.$$

As a consequence, only feasible actions can be chosen and a_t becomes a decision variable. Hence, both $\mathbf{x}$ and $\mathbf{a}$ are decision variables in the model and in the objective we thus minimize over both $\mathbf{x}$ and $\mathbf{a}$.

For linearizing the transition to $s_{t+1,3}$, we can use similar techniques. We first obtain that we need to add the following constraints for $t = 1, 2, \ldots, T-1$ to the stochastic program.

$$\begin{aligned}
&s_{t+1,3} \geq 0 \\
&s_{t+1,3} \geq s_{t,2} - \max\{a_t - s_{t,3} - s_{t,4}, 0\} \\
&s_{t+1,3} \leq B_3(1 - y^s_{t,3}) \\
&s_{t+1,3} \leq s_{t,2} - \max\{a_t - s_{t,3} - s_{t,4}, 0\} + B_3 y^s_{t,3},
\end{aligned}$$

for binary variable $y^s_{t,3} \in \{0, 1\}$ and for $B_3 \in \mathbb{R}$ such that all possible values of $s_{t+1,3}$ are at most B_3 for all weeks t. There is still a maximum function in these equations. We replace this function by the auxiliary variable $r_{t+1,3}$ and add the following constraints for $t = 1, 2, \ldots, T-1$ to the model to linearize this auxiliary variable.

$$\begin{aligned}
&r_{t+1,3} \geq 0 \\
&r_{t+1,3} \geq a_t - s_{t,3} - s_{t,4} \\
&r_{t+1,3} \leq R_3(1 - y^r_{t+1,3}) \\
&r_{t+1,3} \leq a_t - s_{t,3} - s_{t,4} + R_3 y^r_{t+1,3},
\end{aligned}$$

for binary variable $y^r_{t+1,3} \in \{0, 1\}$ and $R_3 \in \mathbb{R}$ such that all possible values of $r_{t+1,3}$ are at most R_3 for all weeks t.

Since we also need to be able to find $s_{t,0}, s_{t,1}$ and $s_{t,2}$, as they occur in the above expressions, we also incorporate the transitions to the states $s_{t+1,0}, s_{t+1,1}$ and $s_{t+1,2}$ in the constraints of the LP. We linearize the expressions for $s_{t+1,1}$ and $s_{t+1,2}$ similarly to linearizing the state transitions for $s_{t+1,3}$ and $s_{t+1,4}$. For the constraint regarding $s_{t+1,0}$ we choose a robust estimate for the outcome of the random variable Λ_t and denote this by $\hat{s}_{t+1,0}$. Note that by introducing this estimate we do not incorporate stochasticity in the arrival of new patients. However, we choose the estimate in such a way that the new patients are likely to fit in the schedule. We suppose that $\mathbf{s}_1$, the state in the first week, is input in the model. Furthermore, the expression $|a_t - \bar{a}|$ in the cost function of the MDP is not linear. However, it is linearizable by introducing the variable w_t for each $t \in \mathcal{T}_T$ and adding the following constraints.

$$\begin{aligned}
&w_t \geq a_t - \bar{a} && \forall t \in \mathcal{T}_T \\
&w_t \geq -a_t + \bar{a} && \forall t \in \mathcal{T}_T.
\end{aligned}$$

Consequently, we arrive at the following ILP for a scenario $\mathbf{d}$ of the random vector $\mathbf{D}$.

$$\min_{\mathbf{x},\mathbf{a}} \quad k_f \sum_{t=1-s_2}^{T} \gamma^t \Big(d_t - \Big(\sum_{s \in \mathcal{T}_{s_1}^{s_2}(t)} x_{t,s} + \sum_{s=t-s_1}^{0} x^*_{t,s} \Big) \Big) + \sum_{s \in \mathcal{T}_T} \gamma^s \sum_{t \in \mathcal{S}_{s_2}^{s_1}(t)} k_m(s-t) x_{t,s}$$
$$+ \sum_{t \in \mathcal{T}_{T-1}} \gamma^t (k_4 \cdot s_{t+1,4} + k_3 \cdot s_{t+1,3} + k_v w_t) + \gamma^T V(\mathbf{s}_T) \tag{5.7a}$$

$$\text{s.t.} \quad \sum_{t \in \mathcal{S}_{s_2}^{s_1}(s)} x_{t,s} \le c_s \qquad \forall s \in \mathcal{T}_T \tag{5.7b}$$
$$\sum_{s \in \mathcal{T}_{s_1}^{s_2}(t)} x_{t,s} + \sum_{s=t-s_1}^{0} x^*_{t,s} \le d_t \qquad \forall t \in \mathcal{T}_T \tag{5.7c}$$
$$w_t \ge a_t - \bar{a} \qquad \forall t \in \mathcal{T}_T \tag{5.7d}$$
$$w_t \ge -a_t + \bar{a} \qquad \forall t \in \mathcal{T}_T \tag{5.7e}$$
$$a_t \le \sum_{i=0}^{4} s_{t,i} \qquad \forall t \in \mathcal{T}_T \tag{5.7f}$$
$$a_t \le c_t - \sum_{\tau} x_{\tau,t} \qquad \forall t \in \mathcal{T}_T \tag{5.7g}$$
$$s_{t+1,0} = \hat{s}_{t+1,0} \qquad \forall t \in \mathcal{T}_{T-1} \tag{5.7h}$$
$$s_{t+1,1} \ge 0 \qquad \forall t \in \mathcal{T}_{T-1} \tag{5.7i}$$
$$s_{t+1,1} \ge s_{t,0} - r_{t+1,1} \qquad \forall t \in \mathcal{T}_{T-1} \tag{5.7j}$$
$$s_{t+1,1} \le B_1(1 - y^s_{t,1}) \qquad \forall t \in \mathcal{T}_{T-1} \tag{5.7k}$$
$$s_{t+1,1} \le s_{t,0} - r_{t+1,1} + B_1 y^s_{t,1} \qquad \forall t \in \mathcal{T}_{T-1} \tag{5.7l}$$
$$r_{t+1,1} \ge 0 \qquad \forall t \in \mathcal{T}_{T-1} \tag{5.7m}$$
$$r_{t+1,1} \ge a_t - \sum_{i=1}^{4} s_{t,i} \qquad \forall t \in \mathcal{T}_{T-1} \tag{5.7n}$$
$$r_{t+1,1} \le R_1(1 - y^r_{t+1,1}) \qquad \forall t \in \mathcal{T}_{T-1} \tag{5.7o}$$
$$r_{t+1,1} \le a_t - \sum_{i=1}^{4} s_{t,i} + R_1 y^r_{t+1,1} \qquad \forall t \in \mathcal{T}_{T-1} \tag{5.7p}$$
$$s_{t+1,2} \ge 0 \qquad \forall t \in \mathcal{T}_{T-1} \tag{5.7q}$$
$$s_{t+1,2} \ge s_{t,1} - r_{t+1,2} \qquad \forall t \in \mathcal{T}_{T-1} \tag{5.7r}$$
$$s_{t+1,2} \le B_2(1 - y^s_{t,2}) \qquad \forall t \in \mathcal{T}_{T-1} \tag{5.7s}$$
$$s_{t+1,2} \le s_{t,1} - r_{t+1,2} + B_2 y^s_{t,2} \qquad \forall t \in \mathcal{T}_{T-1} \tag{5.7t}$$
$$r_{t+1,2} \ge 0 \qquad \forall t \in \mathcal{T}_{T-1} \tag{5.7u}$$
$$r_{t+1,2} \ge a_t - \sum_{i=2}^{4} s_{t,i} \qquad \forall t \in \mathcal{T}_{T-1} \tag{5.7v}$$
$$r_{t+1,2} \le R_2(1 - y^r_{t+1,2}) \qquad \forall t \in \mathcal{T}_{T-1} \tag{5.7w}$$

$$r_{t+1,2} \leq a_t - \sum_{i=1}^{4} s_{t,i} + R_2 y^r_{t+1,2} \quad \forall t \in \mathcal{T}_{T-1} \tag{5.7x}$$

$$s_{t+1,3} \geq 0 \quad \forall t \in \mathcal{T}_{T-1} \tag{5.7y}$$

$$s_{t+1,3} \geq s_{t,2} - r_{t+1,3} \quad \forall t \in \mathcal{T}_{T-1} \tag{5.7z}$$

$$s_{t+1,3} \leq B_3(1 - y^s_{t,3}) \quad \forall t \in \mathcal{T}_{T-1} \tag{5.7aa}$$

$$s_{t+1,3} \leq s_{t,2} - r_{t+1,3} + B_3 y^s_{t,3} \quad \forall t \in \mathcal{T}_{T-1} \tag{5.7ab}$$

$$r_{t+1,3} \geq 0 \quad \forall t \in \mathcal{T}_{T-1} \tag{5.7ac}$$

$$r_{t+1,3} \geq a_t - s_{t,3} - s_{t,4} \quad \forall t \in \mathcal{T}_{T-1} \tag{5.7ad}$$

$$r_{t+1,3} \leq R_3(1 - y^r_{t+1,3}) \quad \forall t \in \mathcal{T}_{T-1} \tag{5.7ae}$$

$$r_{t+1,3} \leq a_t - s_{t,3} - s_{t,4} + R_3 y^r_{t+1,3} \quad \forall t \in \mathcal{T}_{T-1} \tag{5.7af}$$

$$s_{t+1,4} \geq 0 \quad \forall t \in \mathcal{T}_{T-1} \tag{5.7ag}$$

$$s_{t+1,4} \geq s_{t,3} + s_{t,4} - a_t \quad \forall t \in \mathcal{T}_{T-1} \tag{5.7ah}$$

$$s_{t+1,4} \leq B_4(1 - y^s_{t,4}) \quad \forall t \in \mathcal{T}_{T-1} \tag{5.7ai}$$

$$s_{t+1,4} \leq s_{t,3} + s_{t,4} - a_t + B_4 y^s_{t,4} \quad \forall t \in \mathcal{T}_{T-1} \tag{5.7aj}$$

$$x_{t,s} \in \mathbb{N}_0 \quad \forall t, s \in \mathcal{T}_T \tag{5.7ak}$$

$$a_t \in \mathbb{N}_0 \quad \forall t \in \mathcal{T}_T \tag{5.7al}$$

$$w_t \in \mathbb{N}_0 \quad \forall t \in \mathcal{T}_T \tag{5.7am}$$

$$s_{t+1,i} \in \mathbb{N}_0 \quad \forall t \in \mathcal{T}_{T-1},\ 0 \leq i \leq 4 \tag{5.7an}$$

$$r_{t+1,i} \in \mathbb{N}_0 \quad \forall t \in \mathcal{T}_{T-1},\ 0 \leq i \leq 4 \tag{5.7ao}$$

$$y^j_{t+1,i} \in \{0,1\} \quad \forall t \in \mathcal{T}_T,\ 0 \leq i \leq 4,\ j \in \{s,r\}. \tag{5.7ap}$$

As a consequence of the adjustments to and reformulation of the mathematical program, we need to rewrite the stochastic program (5.5). We define the function F_T to be

$$F_T(\mathbf{x}, \mathbf{a}, \mathbf{d}^T) = k_f \sum_{t=1-s_2}^{T} \gamma^t \Big(d_t - \left(\sum_{s \in \mathcal{T}^{s_2}_{s_1}(t)} x_{t,s} + \sum_{s=t-s_1}^{0} x^*_{t,s} \right) \Big) + \sum_{s \in \mathcal{T}_T} \gamma^s \sum_{t \in \mathcal{S}^{s_1}_{s_2}(t)} k_m(s-t) x_{t,s}$$
$$+ \sum_{t \in \mathcal{T}_{T-1}} \gamma^t (k_4 \cdot s_{t+1,4} + k_3 \cdot s_{t+1,3} + k_v w_t) + \gamma^T V(\mathbf{s}_T).$$

We rewrite the stochastic program to the stochastic ILP

$$\min_{(\mathbf{x},\mathbf{a}) \in \mathcal{X}} \ \mathbb{E}_{\mathbb{P}_{\mathbf{D}^T}} [F_T(\mathbf{x}, \mathbf{a}, \mathbf{D}^T)], \tag{5.8}$$

where $\mathbf{D}^T := (D_t)_{t \in \mathcal{T}_T}$. Here, $\mathcal{X}$ represents the set of solutions $(\mathbf{x}, \mathbf{a})$ that satisfy the constraints (5.7b)-(5.7ap).

5.2.3 Sample Average Approximation

In order to deal with stochasticity in the model due to the random variable D_t, representing the demand of follow-up patients in week t, we use the Sample Average Approximation (SAA) method, which was introduced in [25]. With this method we approximate the expectation in the objective function by a sample average estimate derived from random Monte Carlo samples. Therefore, we need to incorporate the sampling into the objective. Moreover, we need to adjust the constraints that involves the random variable, which we sample. We first explain how we can sample the random variable D_t. Thereafter, we discuss the sample average approximation method, including the adjustments to the stochastic program (5.8) due to the sampling. The resulting model is called the SAA problem.

The distribution of the follow-up demand D_t of week t is given by (5.1). Hence, we generate samples of D_t by adding samples from the separate binomial distributions. We note that the binomial distributions depend on previous optimal decisions. However, this gives a problem if the previous optimal decisions are for a week that falls within the planning horizon $\mathcal{T}_T$. This is the case for the demand from the first week after week $1+\min_i\{t_i\}$. Namely, for $t \geq 1+\min_i\{t_i\}$, we have that $t-(1+\min_i\{t_i\}) \geq 1$, which falls within $\mathcal{T}_T$. In order to be able to sample demand, we estimate the previous optimal decision a_s^* by $\hat{s}_{s+1,0}$ and z_s^* by $c_s - \hat{s}_{s+1,0}$.

If we generate N random sample scenario's $\mathbf{d}^1, \mathbf{d}^2, \ldots, \mathbf{d}^N$, the value of the objective with sample scenario $\mathbf{d}^n$ and solution $\mathbf{x}$ can be determined by $F(\mathbf{x}, \mathbf{a}, \mathbf{d}^n)$. And the value of the objective using a sample of length N over the whole horizon can be approximated by

$$\hat{f}_N(\mathbf{x}, \mathbf{a}) = \frac{1}{N} \sum_{n=1}^{N} F_T(\mathbf{x}, \mathbf{a}, \mathbf{d}^n) \tag{5.9}$$

for a solution $\mathbf{x} \in \mathcal{X}$. Furthermore, in the SAA problem the constraint (5.7c) is

$$\sum_{s \in \mathcal{T}_{s_1}^{s2}(t)} x_{t,s} + \sum_{s=t-s_1}^{0} x_{t,s}^* \leq \frac{1}{N} \sum_{n=1}^{N} d_t^n \qquad \forall t \in \mathcal{T}_T. \tag{5.10}$$

For the SAA method, we introduce M to be the number of replications, N the number of scenarios in the sampled problem and N' be the sample size to estimate the objective value for a feasible solution $\mathbf{x}$. The objective of the SAA problem is (5.9) and the constraints of the SAA problem are the constraints of (5.7), where the constraint (5.7c) is replaced by (5.10). The SAA method is given by Algorithm 3.

5.2.4 Integer Linear Programming approximation to the Stochastic Program

In this section, we formulate an ILP that approximates the stochastic ILP (5.5). Since the ILP is deterministic, it is computationally less expensive than the SAA method.

Instead of deciding based on an expectation of the objective function, we replace the scenario d_t of week t by its expectation $\mathbb{E}[D_t]$ for all weeks t. This has consequences for both the objective and one of the constraints. The expectation of D_t for week t can easily be determined by

$$\mathbb{E}[D_t] = \sum_{i=1}^{m} z_{t-t_i} \cdot p_{t_i} + a_{t-t_i} \cdot q_{t_i} + u \cdot f_{t_i}.$$

Algorithm 3: The Sample Average Approximation Algorithm [25]

1 **for** $m = 1, 2, \ldots, M$ **do**
 (a) Generate N sample scenarios.
 (b) Solve the SAA problem. Let $[\hat{\mathbf{x}}_N^m, \hat{\mathbf{a}}_N^m]$ be the solution vector and $\hat{v}_N^m$ be the optimal objective value.
 (c) Generate independent random samples $\mathbf{d}^1, \mathbf{d}^2, \ldots, \mathbf{d}^{N'}$. Evaluate the estimated upper bound of the true optimal objective value $\hat{f}_{N'}(\hat{\mathbf{x}}_N^m, \hat{\mathbf{a}}_N^m)$ using (5.9) and its estimated variance

$$S^2_{\hat{f}_{N'}(\hat{\mathbf{x}}_N^m, \hat{\mathbf{a}}_N^m)} = \frac{1}{N'(N'-1)} \sum_{n=1}^{N'} \left[F_T(\hat{\mathbf{x}}_N^m, \hat{\mathbf{a}}_N^m, \mathbf{d}^n) - \hat{f}_{N'}(\hat{\mathbf{x}}_N^m, \hat{\mathbf{a}}_N^m) \right]^2.$$

2 Evaluate the estimated statistical lower bound of the optimal objective value

$$\bar{v}_N^M = \frac{1}{M} \sum_{m=1}^{M} \hat{v}_N^m.$$

and its estimated variance

$$S^2_{\bar{v}_N^M} = \frac{1}{M(M-1)} \sum_{m=1}^{M} \left[\hat{v}_N^m - \bar{v}_N^M \right]^2.$$

3 For each solution $\hat{\mathbf{x}}_N^m$, $m = 1, 2, \ldots, M$, estimate the optimality gap by

$$\hat{f}_{N'}(\hat{\mathbf{x}}_N^m, \hat{\mathbf{a}}_N^m) - \bar{v}_N^M.$$

and its estimated variance

$$S^2_{\bar{v}_N^M} + S^2_{\hat{f}_{N'}(\hat{\mathbf{x}}_N^m)}.$$

4 Choose the solution $(\mathbf{x}, \mathbf{a})$ with the smallest optimality gap.

Furthermore, the stationary state values $V(\mathbf{s})$ for all $\mathbf{s}$ can be determined on beforehand by applying value iteration, which is given in Algorithm 1. As a result, we obtain the following ILP.

$$\begin{aligned}
\min_{\mathbf{x},\mathbf{a}} \quad & k_f \sum_{t=1-s_2}^{T} \gamma^t \Big(\mathbb{E}[D_t] - \Big(\sum_{s \in \mathcal{T}_{s_1}^{s_2}(t)} x_{t,s} + \sum_{s=t-s_1}^{0} x^*_{t,s} \Big) \Big) + \sum_{s \in \mathcal{T}_T} \gamma^s \sum_{t \in \mathcal{S}_{s_2}^{s_1}(t)} k_m(s-t) x_{t,s} \\
& + \sum_{t \in \mathcal{T}_{T-1}} \gamma^t (k_4 \cdot s_{t+1,4} + k_3 \cdot s_{t+1,3} + k_v w_t) + \gamma^T V(\mathbf{s}_T) \qquad (5.11a)
\end{aligned}$$

$$\text{s.t.} \quad \sum_{t \in \mathcal{S}_{s_2}^{s_1}(s)} x_{t,s} \le c_s \qquad \forall s \in \mathcal{T}_T \qquad (5.11b)$$

$$\sum_{s \in \mathcal{T}_{s_1}^{s_2}(t)} x_{t,s} + \sum_{s=t-s_1}^{0} x^*_{t,s} \le \mathbb{E}[D_t] \qquad \forall t \in \mathcal{T}_T \qquad (5.11c)$$

Constraints (5.7g)-(5.7ap). (5.11d)

5.3 Generating planning advice based on the output of the model

From the stochastic program presented in this chapter, the optimal decision on scheduling follow-up patients is $x^*_{t,s}$ for all weeks t and $s \in \{t-s_1, \dots, t+s_2\}$, representing the optimal number of follow-up patients that need to be moved from week t to week s. Furthermore, the optimal decision a^*_s represents the optimal number of appointment slots that need to be reserved for new patients for week s. The application of the output of the model is threefold. Firstly, an optimal slot division throughout the year can be derived from this output. The optimal number of appointment slots for follow-up patients in a week s can be determined by $z^*_s = \sum_{t \in \mathcal{S}_{s_2}^{s_1}} x^*_{t,s}$ and the optimal number of new patient appointment slots is a^*_s.

Secondly, capacity problems can be noticed far in advance, because this model is run far before the week on which a decision is made. The objective of the SILP minimizes, among other things, the number of patients that cannot be scheduled. Hence, if there are capacity problems in a specific week, the model does not plan patients that do not fit. After solving the SILP, one can easily derive the weeks in which capacity problems will arise. These are the weeks in which the demand is not satisfied. We can find weeks in which the demand is not satisfied by determining for every week t

$$\mathbb{E}[D_t] - \sum_{s \in \mathcal{T}_{s_1}^{s_2}} x^*_{t,s}.$$

If this number is positive, then it is expected that the demand of week t is not satisfied. From this information, the outpatient clinic knows that in or around week t, extra capacity is needed. After making sure that there is extra capacity, the model can be run again with this higher capacity. On the other hand, weeks with more capacity than needed can be noticed far in advance, which could also be useful information for the outpatient clinic.

Thirdly, based on the optimal decision $x^*_{t,s}$ we could also generate a planning advice. The number $x^*_{t,s}$ represents the optimal number of patients that needs to be moved from week t to week s, in the case the follow-up demand equals its expectation. Hence, of the patients for which the follow-up period ends in week t, it is optimal to plan $x^*_{t,s}$ of them in week s. We now explain a method for providing advice to a cardiologist for planning a follow-up consultation for a specific patient. As indicated before, we suppose that there is a chance that a cardiologist will not follow the advice to move a patient, for example due to medical reasons. Therefore, earlier in this chapter we introduced a probability b_{s-t} that a doctor really moves a patient $s-t$ weeks if it is advised. Hence, the number of patients that actually will be moved from week t to week s, which we denote by the random variable $X_{t,s}$, has a binomial distribution with parameters the number of times A an advice should be given to move a patient from week t to week s and the probability b_{s-t}. We can determine the required number of times this advice should be given such that in the end $x^*_{t,s}$ patients are moved. So we want to determine A such that

$$x^*_{t,s} = \mathbb{E}[X_{t,s}] = A \cdot b_{s-t}.$$

Therefore, the expected number of times the advice for moving is required for moving $x^*_{t,s}$ patients from week t to week $s \neq t$ is $\frac{x^*_{t,s}}{b_{s-t}}$. Consequently, the number of times we give the advice to plan a week t patient in week t is $\mathbb{E}[D_t] - \sum_{s \neq t} \frac{x^*_{t,s}}{b_{s-t}}$. We note that this number can be negative if the probability of moving is low.

The advice can be generated in the following way. Each week $t - s_1 \leq s \leq t + s_2$ gets a color indicating how desirable it is to plan the patient in that week. Based on these colors, the cardiologist can decide on the week to schedule the patient. We introduce the variable $g_{t,s}$ to be a measure for how desired it is to schedule a week t patient in week s. In this variable, we take the following two aspects into account.

We let the advice for scheduling a week t-patient be dependent on the required number of times the advice for moving should be given. Since the advice may change after every patient that is planned, we take into account the patients that already have been planned. We define $\hat{x}_{t,s}$ to be the number of patients that have been moved from week t to week s so far. Hence, this variable is updated every time a patient is moved from week t to week s. Then the number of times an advice for moving a patient from week t to week s still needs to be given, relative to the total demand of week t, is

$$\frac{x^*_{t,s} - \hat{x}_{t,s}}{b_{s-t}\big(\mathbb{E}[D_t] - \hat{x}_{t,s}\big)}.$$

The number of times we have to give an advice to plan a week t patient in week t, relative to the total demand of week t, is then

$$1 - \sum_{s \neq t} \frac{x^*_{t,s} - \hat{x}_{t,s}}{b_{s-t}\big(\mathbb{E}[D_t] - \hat{x}_{t,s}\big)}.$$

In addition, we let the advice for scheduling a week t patient be dependent on the number of free slots in the weeks around it. The number of free slots in a week s can be determined by $c_s - a^*_s - z^*_s$. If the number of free slots in a week is high, then it is more desirable to plan a patient in this week. Hence, we define the advice variable $g_{t,s}$ to be

$$g_{t,s} = \begin{cases} \frac{x^*_{t,s} - \hat{x}_{t,s}}{b_{s-t}\left(\mathbb{E}[D_t] - \hat{x}_{t,s}\right)} \cdot (c_s - a^*_s - z^*_s) & \text{for } s \in \mathcal{T}^{s_2}_{s_1}(t), s \neq t \\ 1 - \sum_{\substack{s \in \mathcal{T}^{s_2}_{s_1}(t) \\ s \neq t}} g_{t,s} \cdot (c_t - a^*_t - z^*_t) & \text{for } s = t. \end{cases}$$

The scale of colors of the weeks is from red (really undesired to schedule a patient in this week) to green (really desired to schedule a patient in this week). Based on values of $g_{t,s}$ colors are given to weeks. For an example of how such an advice looks like, we refer to Figure 5.1.

Chapter 6

Numerical results

In this chapter we present numerical results. In the first section, we test the influence of truncation of the horizon of the nonstationary MDP presented in Chapter 4 and the choice of the salvage values on the optimal decision in the first epoch. Furthermore, in the second section, we present the numerical results on the case of the cardiology outpatient clinic of Gelre Ziekenhuizen in Zutphen. We do this by setting up a simulation study. We assess the performance of our method of planning follow-up patients that was presented in Chapter 5. Although the method is on planning follow-up patients, it also takes into account the new patients. Therefore, by this simulation study we can assess the effect of using this method for both the new and follow-up patients.

6.1 Numerical examples of truncation of the horizon of the MDP

In this section, we numerically investigate the influence of the truncation of the horizon described in Section 4.2.1, and letting the salvage values thus equal the stationary state values, on the decision in the first time epoch. We also investigate the case of all salvage values being equal to zero. We execute this numerical investigation by considering an example for which we solve the truncated nonstationary MDP using Algorithm 2 and determine the optimal decision in the first epoch for several values of the truncated horizon T. We note that this example can only provide us an indication of the influence of the horizon and the salvage values on the initial decision. The horizon lengths that we test are $T = 4$ to $T = 15$. The input parameters for the stationary MDP are reported in Table 6.1.

The stationary MDP is solved using the Value Iteration Algorithm presented in 1, which was implemented in Python 3.9.9. We initialize the values of all states to be zero and with $\epsilon = 10$. The algorithm converges in nine iterations, which took around seven hours on a 3.0GHz Intel(R) Xeon(R) Gold 5217 CPU. In Figure 6.1 we plot the convergence of the values of eight random states and indeed observe convergence. The values of two of the eight states do not seem to change in these nine iterations. Hence, we take a closer look at the convergence of the values of state $(8, 0, 0, 0, 0)$ and $(10, 5, 9, 3, 0)$ in Figure 6.2 and observe that the values of these states also have changed with the iterations and have converged.

In our first test example the initial state equals $(8, 10, 0, 0, 0)$. Hence, the initial situation in-

Table 6.1: Input parameters for the stationary MDP.

Parameter	Value
λ	8
c^n	9
γ	0.8
k_4	1 000 000
k_3	500 000
k_v	0.1

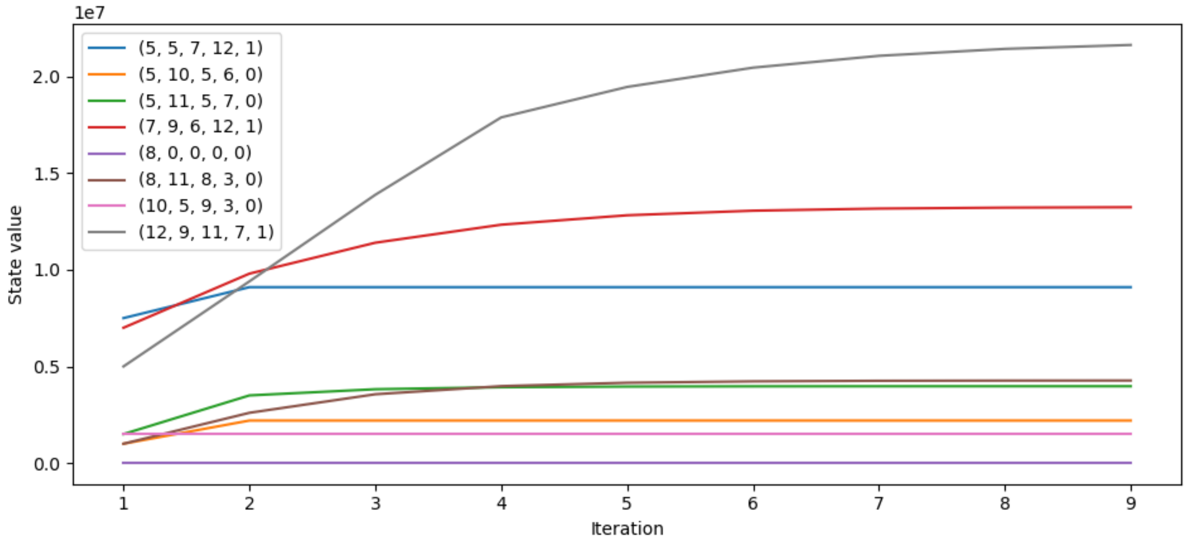

Figure 6.1: Convergence of the values of eight of the states.

cludes that eight patients just got referred and there are ten patients that got referred a week ago but are not planned yet. In our second example the initial state equals $(6, 10, 0, 0, 0)$. Furthermore, we let $c_1^n = 10$, instead of 9. We increase the capacity of the first week in this second example to test whether planning patients is postponed for a shorter horizon. We test the influence of the horizon length and the salvage values of the states at the end of the horizon on the optimal decision in the first decision epoch. We start with a horizon of four epochs, meaning that we only take λ_t and c_t^n into account for $t \leq 4$. The salvage values resulting from our method are denoted by $v^\gamma(\mathrm{s}_4)$ for all possible states s_4. We also test the influence of salvage values equal to zero for all states. The initial optimal decision for the two examples is reported for the different horizon lengths and the two choices of salvage values in Table 6.3. Testing all horizons took approximately a day on a 3.0GHz Intel(R) Xeon(R) Gold 5217 CPU.

Since the original nonstationary infinite horizon MDP is not solvable, the actual optimal first decision cannot be determined. However, from the results in Table 6.3 we may deduce that the optimal first decision for the first example is to schedule 8 patients and for the second example to schedule 10 patients. We do not expect that the optimal decisions for the two examples are different than the ones we found, as it is to be expected that from a certain point in time the

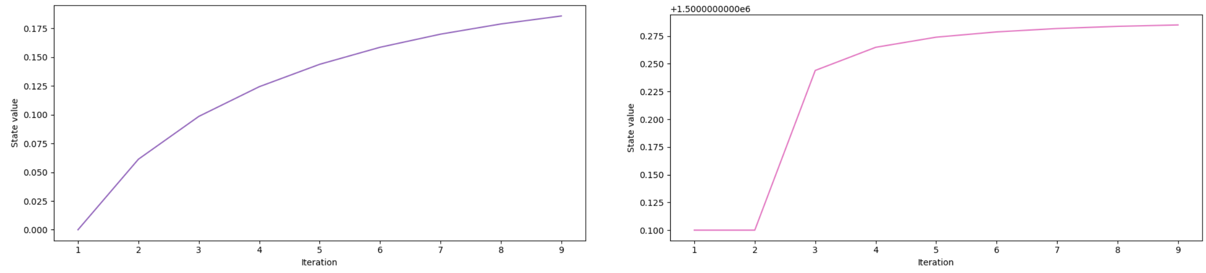

(a) State $(8,0,0,0,0)$.

(b) State $(10,5,9,3,0)$.

Figure 6.2: Convergence of the values of states $(8,0,0,0,0)$ and $(10,5,9,3,0)$.

Table 6.2: The demand parameter and the capacity for new patients for fifteen weeks.

Week t	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
λ_t	6	6	8	8	10	10	8	8	6	6	8	8	10	10	8
c_t^n	8	8	8	8	10	8	8	8	8	6	6	8	8	8	8

exact time dependent information after this point does not influence on the optimal decision of the first week. And hence, the optimal decision in the first epoch does not change anymore by increasing the horizon. We observe that when choosing the salvage values to be the stationary state values, this decision can be determined with nonstationary MDPs having shorter horizons than when choosing the salvage values to be zero. In the first example, the optimal first decision of planning 8 patients already can be determined with a horizon of four epochs.

Table 6.3: The optimal decision in the first epoch for different horizon lengths T and different choices for the salvage values.

T	Optimal first decision for salvage value $v^\gamma(\mathbf{s}_T)$	Optimal first decision for salvage value 0
4	8	2
5	8	2
6	8	4
7	8	6
8	8	8
9	8	8
10	8	8
11	8	8
12	8	8
13	8	8
14	8	8
15	8	8

(a) Initial state $(8,10,0,0,0)$

T	Optimal first decision for salvage value $v^\gamma(\mathbf{s}_T)$	Optimal first decision for salvage value 0
4	7	2
5	6	2
6	10	2
7	10	4
8	10	8
9	10	9
10	10	10
11	10	10
12	10	10
13	10	10
14	10	10
15	10	10

(b) Initial state $(6,10,0,0,0)$, $c_1^n = 10$

6.2 Simulation study to test the performance of the ILP

In this section, we present numerical results on the performance of the method presented in Chapter 5. We assess the performance of the method in terms of access times of patients. The simulation is, as the model, on a weekly level. Hence, the access times of patients are also measured in weeks. During the simulation, we keep track of access times of both the new patients and the follow-up patients. We compare the performance of our method to another method of dividing the appointment slots between patients groups and planning them. We call the policy that results from this method the myopic policy. We start with explaining how we simulate both the myopic policy and the policy that follows from our method. Thereafter, we present the input parameters for the simulation study, which we base on the situation of the cardiology outpatient clinic of Gelre Ziekenhuizen. In Section 6.2.2 we explain the setup of our simulation and in Section 6.2.3 we report the results.

In the myopic policy we try to imitate the current planning strategy as much as possible, meaning that we determine the appointment slot division per week in the following way. The number of appointment slots for the new patients is equal to the mean number of new patients arriving per week and is the same in all weeks. The remainder of the capacity is for follow-up patients. As a consequence, in weeks with less capacity, less follow-up patients can have consultations, while in weeks with more capacity, more follow-up patients can have consultations. Furthermore, patients that become a follow-up patient added to the demand of the week in which their follow-up period ends. The planning of both the new patients and follow-up patients is according to the first come first serve principle, meaning that the patients who do not fit in the week they arrive, are the first patients that are scheduled in the next week. Hereafter, patients that arrive in this next week are scheduled. With follow-up patients arriving in a week we mean that in this week their follow-up period ends. The access time of follow-up patients is the time difference between the week of their arrival and the week of their consultation.

The policy that follows from the method presented in Chapter 5 includes that the slot division per week may differ. The optimal slot division per week follows from solving the model. Namely, it is optimal to plan a_s^* slots for new patients and $z_s^* = \sum_t x_{t,s}^*$ slots for follow-up patients. It may be the case that the number of slots for new patients and the number of slots for follow-up patients is less than the total capacity. Since we use a robust estimate for the new patient demand, we choose to give this potential rest capacity to follow-up patients. Furthermore, for the planning of follow-up patients an advice is generated according to the method described in Section 5.3. If the advice is followed, which happens with probability b_{s-t} for moving a patient from week t to week s, then this patient is added to the follow-up demand of week s. Otherwise, the patient is added to the follow-up demand of week t. The access times of follow-up patients is then the time difference between the week to which the patient is added and the week of their actual consultation. Planning of both the new and follow-up patients is again done according to the first come first serve principle. The first difference with the myopic strategy is the slot division. The second difference is the week of which the follow-up patients are added to the follow-up demand.

We implement the model in Python 3.9 [28]. The deterministic ILP 5.11 presented in Chapter 5 is solved using the solver of Gurobi [29]. The ILP can be solved in 0.04 seconds on an AMD Ryzen™ 5650U PRO processor, which yields the optimal solution for one week. When implementing the ILP, due to technical reasons, the state variables also needed to become decision variables as their value depends on the values of the decision variables $\mathbf{x}$ and $\mathbf{a}$. However, $V(\mathbf{s}_T)$ is not linear in $\mathbf{s}_T$. In order to tackle this problem, we perform linear regression to approximate the state value based on the values of the state variables. For a state $\mathbf{s} = (s_0, s_1, s_2, s_3, s_4)$ this linear regression model is of the form

$$V(\mathbf{s}) = \theta_0 + \theta_1 s_0 + \theta_2 s_1 + \theta_3 s_2 + \theta_4 s_3 + \theta_5 s_4 + \epsilon_{\mathbf{s}},$$

where $\theta_i,\ 0 \leq i \leq 5$ are coefficients and $\epsilon_{\mathbf{s}}$ is the error variable. We use the `LinearRegression` function of the `sklearn` package in Python, which fits the model by minimizing the residual sum of squares between the observed values and the predicted values by the linear approximation. We learn the coefficients by performing linear regression on the values of the stationary MDP of the example in Section 6.1. This results in a coefficient of determination $R^2 = 0.95$. The values of the coefficients are shown in Table 6.4.

Table 6.4: Coefficients of the linear regression model, rounded to integers.

Coefficient	**Value**
θ_0	$-65\,991\,627$
θ_1	$1\,765\,278$
θ_2	$2\,269\,834$
θ_3	$2\,789\,388$
θ_4	$2\,768\,850$
θ_5	$3\,385\,158$

Since the objective is linear, we do not take the coefficient θ_0 into account, as it does not influence the optimal decision. Hence, the objective changes to

$$\begin{aligned}\min_{\mathbf{x},\mathbf{a}} \quad & k_f \sum_{t=1-s_2}^{T} \gamma^t \Big(\mathbb{E}[D_t] - \Bigg(\sum_{s \in \mathcal{T}_{s_1}^{s_2}(t)} x_{t,s} + \sum_{s=t-s_1}^{0} x_{t,s}^* \Bigg) \Big) + \sum_{s \in \mathcal{T}_T} \gamma^s \sum_{t \in \mathcal{S}_{s_2}^{s_1}(t)} k_m(s-t) x_{t,s} \\ & + \sum_{t \in \mathcal{T}_{T-1}} \gamma^t (k_4 \cdot s_{t+1,4} + k_3 \cdot s_{t+1,3} + k_v w_t) + \gamma^T (\theta_1 s_{T,0} + \theta_2 s_{T,1} + \theta_3 s_{T,2} + \theta_4 s_{T,3} + \theta_5 s_{T,4}).\end{aligned}$$

6.2.1 Input parameters for the simulation

We want to test our method for a setting that is realistic for the cardiology outpatient clinic at Gelre Ziekenhuizen. Therefore, we base our input parameters on the data of Gelre, which is described in Section 2.3.1. The parameters that we can retrieve from this data are the demand of new patients, the capacity and the probabilities of follow-up consultations. Moreover, the values of s_1 and s_2 are chosen to be 4 weeks and 8 weeks respectively, which is also based on the conversations with cardiologists and planners of Gelre Ziekenhuizen.

In order to predict the weekly capacity over the year, the cardiology outpatient clinic categorizes every week to be either a production week, a reduction week, or a surplus week. We follow this categorization in order to determine the weekly capacity in our simulation. The majority of the weeks is categorized as a production week, which can be seen as an average week. In a production week, the cardiologists can take care of ten consultation hours in Zutphen and the PA of eight. Furthermore, in both Lochem and Dieren a consultation hour can take place. Fifteen weeks are categorized as reduction weeks in which there is less capacity than average. In such a week, the cardiologists can take care of five consultation hours in Zutphen and the PA 5.3 consultation hours. In addition, either in Lochem or Dieren a consultation hour takes place. There are six weeks categorized as surplus weeks. In these weeks it is to be expected that all cardiologists and the PA are present and that fifteen consultation hours take place in Zutphen by the cardiologists and eight by the PA, two in Lochem and one in Dieren. From the data we retrieve that in a consultation hour by a cardiologist in Zutphen, twelve patients can be helped. A PA can help eight patients per consultation hour and in Lochem and Dieren ten patients per consultation hour can be helped. Furthermore, as previously explained, we assume the weekly reserved slots for semi-urgent patients to be fixed and equal to 25 slots per week. In addition, per week 6 slots are reserved for patients that need a consultation for hart rehabilitation. Therefore, in a production week, the number of patients that can have a consultation is 173. In a reduction week, at most 82 consultations can take place and in a surplus week at most 243 consultations can take place. As assumed in Chapter 4, the arrival of new patients is according to a Poisson distribution with time dependent parameter λ_t for week t. Using the data of 2018 and 2019 we retrieve this time dependent parameter per week. The number of GP referrals and the number of internal and external referrals form the new patient demand. We plot the average number of new patients per week over 2018 and 2019 in Figure 6.3. These years are just two realizations of the demand in each week, which makes that this graph is fluctuating a lot. Therefore, we first analyze the fluctuations per month to get a more global view and then

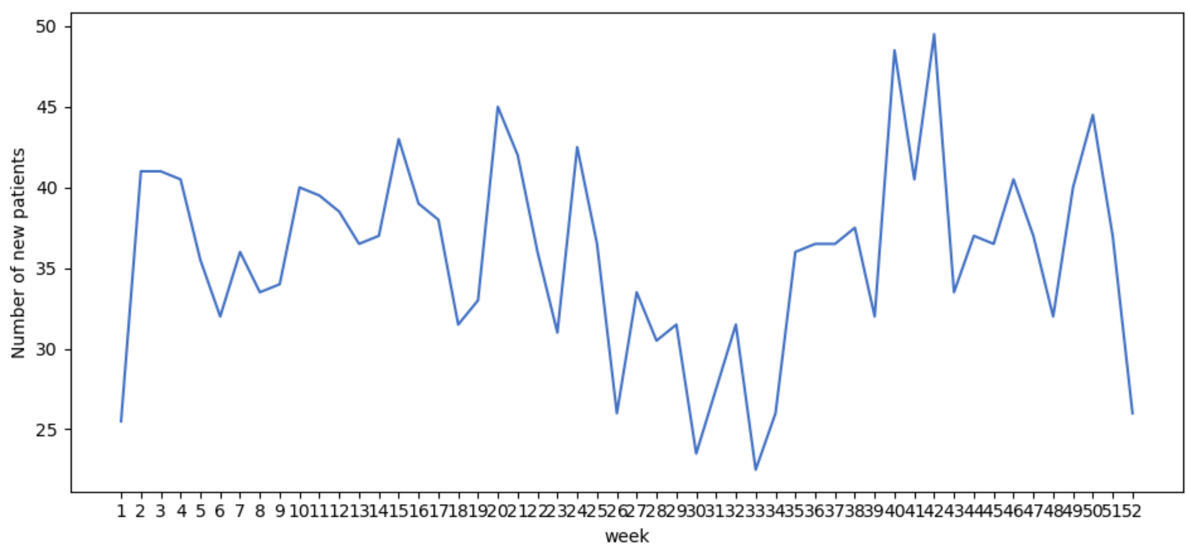

Figure 6.3: Average number of new patients per week of 2018 and 2019.

translate this to a weekly level. We normalize the patient demand in the sense that we divide it by the number of days in that month and multiply the result by 30. This way we get a better indication of the fluctiations over the year. The normalized average number of arrivals per month of 2018 and 2019 is plotted in Figure 6.4.

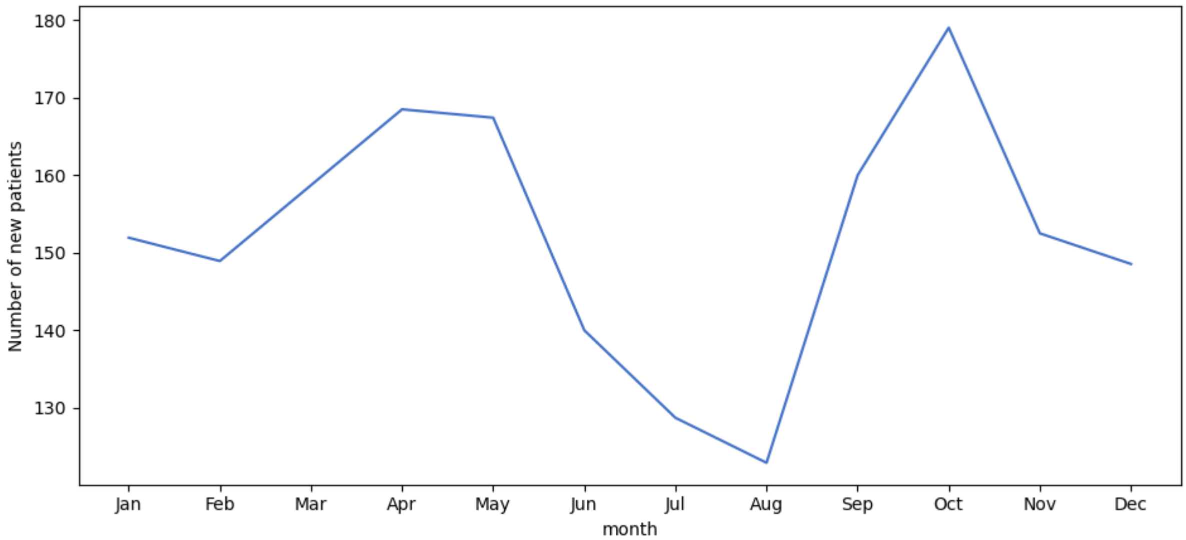

Figure 6.4: Average number of new patients per month of 2018 and 2019.

Based on Figures 6.4 and 6.3, we deduce the weekly demand parameter λ_t for each week of the year. The values of the parameter per week are in Table 6.5. As explained in Section

Table 6.5: Values of the time dependent patient demand parameter λ_t per week of the year.

Week t	λ_t	**Week** t	λ_t	**Week** t	λ_t	**Week** t	λ_t
1	30	14	45	27	30	40	45
2	42	15	45	28	30	41	45
3	42	16	45	29	30	42	45
4	42	17	37	30	30	43	45
5	37	18	37	31	30	44	42
6	37	19	45	32	30	45	42
7	37	20	45	33	30	46	42
8	42	21	45	34	30	47	42
9	42	22	37	35	42	48	42
10	42	23	37	36	42	49	42
11	42	24	30	37	42	50	42
12	42	25	30	38	42	51	42
13	42	26	30	39	45	52	30

5.2.2, we estimate the outcome of the random variable Λ_t by $\hat{s}_{t+1,0}$. We choose this estimate to be $1.02 \cdot \lambda_t$.

Furthermore, we retrieve the probability of a follow-up consultation after another consultation. Follow-up consultations may follow from all consultation types. For each consultation that took place in 2018 we search in the remarks field in the data set if the period until the next appointment is indicated. If this is not the case, we search for a possible next appointment that either took place in 2018 or 2019 and calculate the time difference between the two consecutive appointments. A histogram of the number of days between two consecutive consultations is provided in Figure 6.5. Based on this data analysis and conversations with cardiologists, we

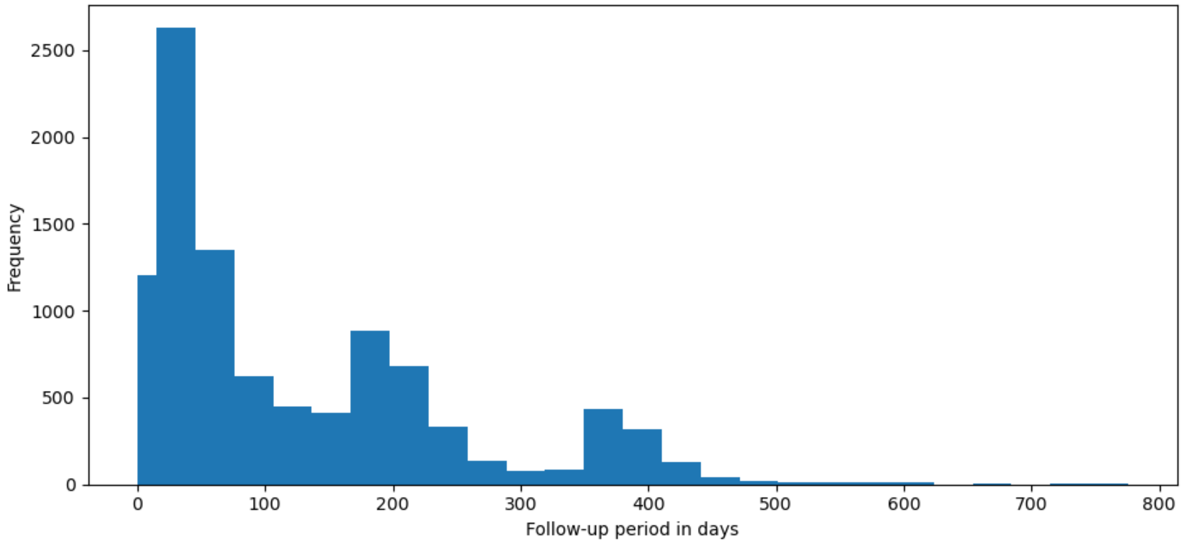

Figure 6.5: Histogram of the follow-up periods of consultations in 2018.

consider the follow-up periods to be two months, six month, and one year. Since our model is on a weekly level instead of a monthly level, we take the first follow-up period to be nine weeks, the second to be 26 weeks and the third to be 52 weeks. Every follow-up consultation after a period of at most four months is considered to be a 'two-month follow-up consultation'. For determining the probabilities of having a follow-up consultation after a specific period, every consulation after a period between five and eight months is considered to be a 'six-month follow-up consultation' and every follow-up consultation that took place at least nine months after the previous consultation is considered to be a 'twelve-month follow-up consultation'. The probabilities of having a follow-up consultation after one of these periods are given in Table 6.6.

Table 6.6: Probabilities of having a consultation per follow-up period.

Previous consultation	**Follow-up after 2 months**	**Follow-up after 6 months**	**Follow-up after 12 months**	**No follow-up**
New consultation	0.45	0.05	0.04	0.46
Follow-up consultation	0.24	0.21	0.19	0.36
Semi-urgent consultation	0.37	0.17	0.11	0.37
Rehabilitation consultation	0.33	0.14	0.10	0.43

Based on conversations with the cardiologists, the probability of moving a patient to a later week than the end of the follow-up period should be higher than the probability of moving a patient to a week earlier. We let the probability b_{s-t} of moving a follow-up patient be given by

$$b_{s-t} = \begin{cases} 0.8^{s-t} & \text{if } s > t \\ 0.6^{t-s} & \text{if } s < t \\ 1 & \text{if } s = t. \end{cases}$$

Furthermore, we set all auxiliary variables in the big-M constraints of the ILP to $1\,000$. Hence, $B_1 = R_1 = B_2 = R_2 = B_3 = R_3 = B_4 = 1\,000$. Values of the other parameters are listed in Table 6.7. For choosing the values, we also take conversations with cardiologists and planners into account. We note that this simulation is a numerical example with which we try to get an indication of the performance of the method. Regarding the costs in the objective function of the ILP, we put significantly more weight on planning patients within acceptable time than on reducing variability and moving follow-up patients. Hence, the costs for planning patients not within acceptable times are chosen to be significantly larger than the costs for variability and moving follow-up patients in this simulation.

Table 6.7: Parameter values.

Parameter	Explanation	Value
T	Horizon of the ILP	25
K_f	Costs for not planning a follow-up patient	$10\,000$
k_4	Costs for planning a new patient after four weeks	$1\,000\,000$
k_3	Costs for planning a new patient after three weeks	$500\,000$
k_v	Costs for variability in slots for new patients	0.1
κ_m	Coefficient in the cost function of moving follow-up patients	0.1
γ	Discount factor	0.9

6.2.2 Simulation setup

We first solve the ILP for all weeks for which we run the simulation. In reality, expectations about parameters might change, but in the simulation, we fix our expectations for the new patient demand and the capacity for the whole simulation length. The follow-up patient demand depends on previous decisions. Hence, we need to initialize the previous decisions in order to determine the expected follow-up patient demand. The previous number of slots for semi-urgent patients and for heart rehabilitation consultations are the same as in the simulated years, namely 25 and 6, respectively. We let the previous number of new patient slots be λ_w for all weeks before the simulation started, where w is the week of the year. We let the previous number of slots for follow-up patients be 112. On these numbers we base the follow-up demand in the first weeks for which we solve the model. After a year we totally base the expected follow-up demand on decisions that are actually taken by our method, as the longest follow-up period is one year. Hence, the decisions of the first year are not really representative.

We simulate the patient demand in the following way. We sample the number of new patients arriving in a specific week according to a Poisson distribution with time-dependent parameter.

Follow-up consultations arise from previous consultations. Hence, follow-up patient demand arises from other demand, as also indicated in (5.1). We sample follow-up patient demand in the following way. For every patient that has a consultation, there is a probability that this patient has a follow-up consultation after a certain period. Hence, for each patient, we generate a random number between zero and one, and based on the outcome, this patient becomes a follow-up patient in the future. Under the myopic policy, this patient is added to the follow-up patient demand of the concerned week in the future. However, under the policy of the method that is developed in this thesis, an advice for planning the follow-up patient is generated. As explained in Chapter 5, there is a chance that the patient will or cannot be planned in the advised week. We simulate whether or not the advise is followed by generating a random number. If this random number is higher than the probability that the advise is followed, the follow-up patient is added to the demand of the week their follow-up period ends. If the random number is lower than the probability that the advise is followed and thus the given advise is actually followed, then the patient is added to the follow-up patient demand of the advised week.

We start our simulation in the situation in which all previous new patients are scheduled. Hence the initial state of the MDP is $\mathbf{s_0} = [s_{0,0}, 0, 0, 0, 0]$, where $s_{0,0}$ is randomly generated from a Poisson distribution. Moreover, in the initial situation no follow-up patients from the past need to be scheduled in one of the simulated weeks. Hence, the variables $x_{t,s}$ for $t < 0$ and $s > 0$ are zero. Since the follow-up patient demand follows from previous patients together with the planning advice, in the first weeks of the simulation no follow-up demand could have been generated. Hence, we replace the follow-up demand for the weeks it could not have been generated by the expected demand, which is 112 for the first weeks. Since the initial situation is not very representative for a general week and since the decisions from the model are also not representative in the beginning, we warm up for five years, thus for 260 weeks, after which we start collecting results.

6.2.3 Results of the simulation

After warming up the simulated process of the arrival and planning process of patients for five years, we run the simulation for twenty years, which is $1\,040$ weeks. This took about 7 seconds. In the simulation both the strategy from our method and the myopic strategy are simulated parallel. For each patient, we determine his or her access time in terms of weeks. The access time for follow-up patients under our method is calculated as the time difference between the advised week and the actual week of the consultation. Furthermore, we determine the percentage of patients having access time lower than four weeks. For new patients we also determine the percentage of patients having access time lower than three weeks. We compare the results of our method to both the myopic strategy and the current planning strategy, for which we base the results on the data analysis performed in Chapter 2. We execute 50 simulation runs. The results regarding the access times for new patients are in Table 6.8 and the results regarding access times for follow-up patients are in Table 6.9. The access time is reported in weeks and the mean access time and a 95% confidence interval over the simulation runs are given. Furthermore, we provide the mean and 95% confidence intervals over the simulation runs for the percentage of patients having access time lower than four and lower than three weeks, respectively.

Table 6.8: Simulation results: access times for new patients.

	Treeknorm	**Current (2018/2019)**	**Developed method**	**Myopic method**
Access time (weeks)	-	3.7	0.4 ± 0.1	3.6 ± 1.5
Percentage <4 weeks	100%	75.9%	$99.9\% \pm 0.2\%$	$56.6\% \pm 23\%$
Percentage <3 weeks	80%	44.3%	$99.5\% \pm 0.3\%$	$39.1\% \pm 19\%$

Table 6.9: Simulation results: access times for follow-up patients.

	Current (2018/2019)	**Developed method**	**Myopic method**
Access time (weeks)	4.3	0.9 ± 0.04	0.9 ± 0.04
Percentage <4 weeks	73.2%	$99.6\% \pm 0.2\%$	$98.3\% \pm 0.04\%$

We also look at the capacity utilization. We consider the utilization of new patient slots and follow-up patient slots separately. We base the utilization of the current situation on the year 2019, because the needed data is available for this year. For the current situation, we do not distinguish between new patient and follow-up patient utilization. The results of comparison between the current situation, the performance of our method and the performance of the myopic method in terms of capacity utilization are shown in Table 6.10. We report the average utilization per week and also provide 95% confidence intervals over the simulation runs.

Table 6.10: Simulation results: utilization.

	Developed method	**Myopic method**
Utilization new patients	$82.4\% \pm 1\%$	$98.6\% \pm 1\%$
Utilization follow-up patients	$82.3\% \pm 1\%$	$79.2\% \pm 1\%$

To illustrate our method, we compare the slot division and patient demand under our method and under the myopic method for one simulated year in Figure 6.6. We also plot the patient demand. The new patient demand is the same in both graphs. However, due to the planning advices for follow-up patients, the follow-up patient demand is different. For follow-up patients, we plot the demand after giving the advice. That is, a patient is added to the follow-up demand of a specific week if that week is advised and the planning advice is followed. If the advice is not followed, this patient is added to the demand of the week his or her follow-up period ends. Furthermore, for this simulated year, we also determine the access time. The results for our method and the myopic method are shown in Figure 6.7.

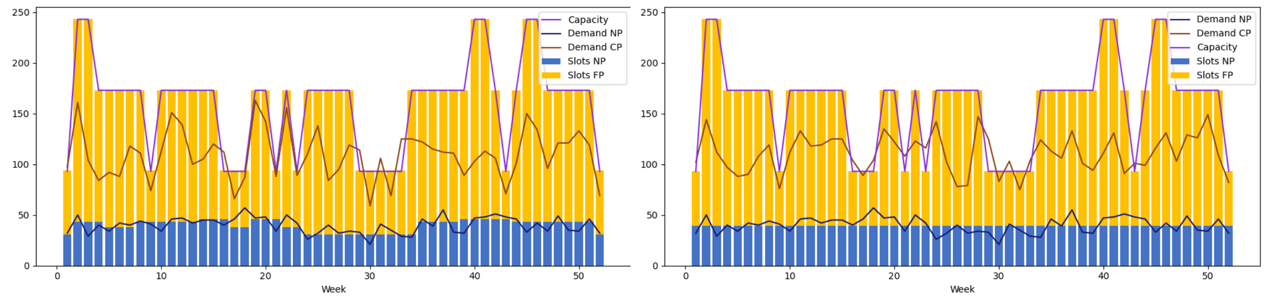

(a) Developed method

(b) Myopic method

Figure 6.6: Comparison of methods in terms of capacity division and the patient demand for one simulated year.

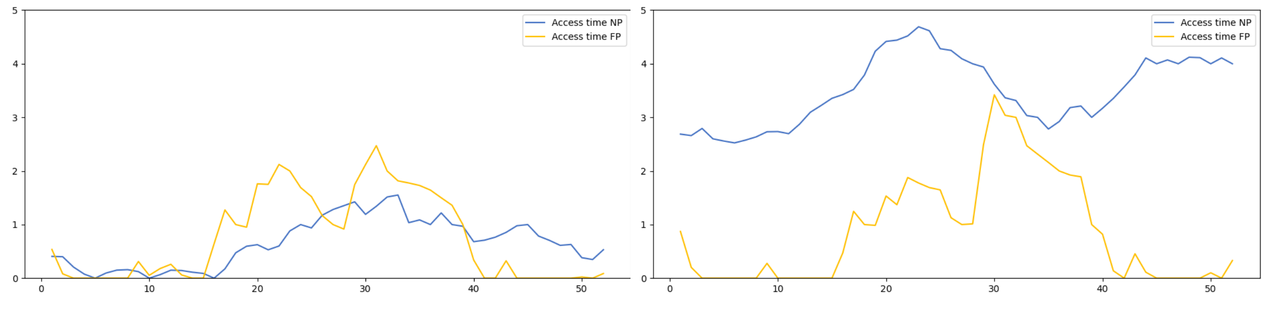

(a) Developed method

(b) Myopic method

Figure 6.7: Comparison of methods in terms of access times for one simulated year

Chapter 7

Discussion

In this chapter, we first discuss the obtained results. We consider the analysis of the error bound of truncating the horizon of the nonstationary MDP, the corresponding numerical results, and the simulation results. Furthermore, we discuss the important aspects for implementation of the method in the hospital. We close this chapter by discussing the limitations of or research and recommendations for future research.

7.1 Discussion of the results

In this research, we truncated the horizon of a nonstationary MDP and approximated the salvage values, the state values at the truncated horizon, by the state values of a similar stationary MDP. In Section 4.3 we analysed the error that arises due to this approximation. We were able to find an error bound. However, the bound was not really tight. That is, if we would determine a similar bound for the case that all salvage values are chosen to be zero, the error bound would be lower. However, it is reasonable to expect that the stationary MDP state values are a better choice for the salvage values than zero. When choosing the salvage values to be zero, there is no consequence of postponing planning decisions. The stationary MDP state values as salvage values give an indication of the value of the state in a general week. Hence, it is not encouraged to postpone planning decisions, as in a general week also patients arrive and need to be scheduled. The consequences of this reasoning can also be seen in the numerical example of Section 6.1. In the case of the nonstationary MDP state values as a choice for the salvage values the optimal decision for the first decision epoch could already be determined for shorter truncated horizons than in the case of zero salvage values. This indicates that the error in value function of the first decision epoch is smaller for the case of stationary MDP state values as salvage values than for zero salvage values. Although this is just one example, we observed similar behaviour for other examples that we did not include in this thesis.

In terms of computational costs it is very relevant to be able to determine the optimal decision already for small truncated horizons. For a case that is realistic to the hospital, in which the arrival rate is higher causing the state and action space to become larger, determining the optimal decision for the first decision epoch and thus one week takes longer than a day for an horizon of six weeks. In order to tackle the problem of large state and action space for situations in which the patient demand and capacity are larger than in the hospital case we considered, also referred to as the curses of dimensionality, one could make use of approximate dynamic programming techniques [30]. Furthermore, solving the stationary infinite horizon MDP took approximately twelve hours for the example of Section 6.1. However, this only needs to be done once as a precalculation and then can be used every week and therefore its computation time is of less relevance.

From the results of the simulation, presented in Section 6.2.3, we observe that for new patients, the access times are significantly shorter when using our method instead of the myopic strategy. With our developed method the Treeknorm is satisfied in the simulation study. Furthermore, the variability in access times is much larger when the myopic method is used. This can also be observed in Figure 6.7. For other years, the graph resulting from our developed method looks similar, but the graph resulting from the myopic method is different for each year and each run. We expect that this result is the consequence of two aspects of our method. The first aspect is that our method takes into account yearly fluctuations in patient demand. By doing this, flexibility is used to cope with the variability in patient demand. As a result, the blueprint is prepared for this variability on beforehand, which we also observe in Figure 6.6. Therefore, it would also be interesting to test the performance of a policy in which the number of slots for new patients is equal to the expected patient demand in each week. The second aspect that may be the cause of lower access times for new patients by using our method is that we reserve slots based on a robust estimate of the expected new patient demand. This way, we consciously aim for a higher utilization, of which the result also could be observed. The utilization of new patient slots under the myopic method is almost one, which in general means infinite access times because of instability of the system. It would be interesting to also test the performance of a myopic strategy under which a more robust estimate for the expected new patient demand is chosen, instead of just the expected number of arriving new patients.

From the simulation results, we also observe that the access time and utilization for follow-up patients is similar for our method and the myopic method. Under both methods, it should be possible for almost all follow-up patients to be scheduled within four weeks of the expected week. Since a robust estimate is used for new patients, follow-up patients get less capacity in our method than in case of the myopic strategy. However, this does not seem to have influence on their access times, according to the simulation results.

7.2 Implementation of the method in the hospital

In this section we discuss the aspects that are important for actual implementation of the method described in Chapter 5 in the hospital. The method can be used for the long-term tactical planning. As explained in Section 5.3, the application of the method of Chapter 5 is threefold. Namely, capacity problems can be recognized far in advance, the optimal capacity division between new patients and follow-up patients can be deduced and a method is presented for generating an advice for planning follow-up patients. The goal is to be able to give an advice for every patient that needs a follow-up consultation. Since the largest follow-up period is one year, already one year in advance the model should be solved for a specific week. The method for providing an advice from the outcome of the ILP can be applied in real-time. That is, every time a patient is planned, the advice for the next follow-up patients is adapted based on the planning of this patient.

Before using the method, some precalculations should be done. First, per week of the year the expected demand should be determined and the probabilities of having a follow-up per follow-up period and previous consultation type should be determined. For the situation of the cardiology outpatient clinic of Gelre Ziekenhuizen this was done in Section 6.2.1. Furthermore, the stationary infinite horizon MDP needs to be solved. Moreover, linear regression has to be applied on them, which is explained in Section 6.1. In principle, these precalculations only have to be done once, unless significant changes in capacity or patient demand are expected. This can happen for example when the outpatient clinic expands its capacity and/or the catchment area changes.

For each week for which an advice and optimal slot division needs to be determined, the ILP (5.11) needs to be solved. The week for which the slot division needs to be determined, has to be the first week of the ILP. In order to solve the model, the expected capacity and expected patient demand need to be known for this week and the 25 weeks after this week, as this is the horizon length of the ILP. Furthermore, one should have information on the slot division of the previous year and on the amount of satisfied expected demand in the weeks before the required week. After solving the ILP, the optimal slot division for the required week can be deduced from the optimal decision for the first week of the ILP. The ILP determines the slot division for the whole horizon of 25 weeks. However, only the decision for the first week of the ILP is actually implemented. If the slot division of the next week needs to be determined, the planning horizon is 'rolled' one week forward. Then this week is the first week of the ILP and also a forecast of capacity and demand needs to be known for one week later.

The advice that is generated in Section 5.3 is updated after every patient that is planned in a certain week in the future. To make this method work in practise, the generation of the advice should be implemented in a tool that may be part of the planning software of the hospital. Every time a cardiologist wants to plan a patient, he uses the tool for a planning advice. After the patient is planned, the planning decision will be processed, regardless of whether the given advice is actually followed. Hence, the advice will be updated after each planned follow-up patient.

To be able to give an advice for every follow-up patient, the slot division for a specific week already needs to be determined a year in advance. We are aware that this is quite far in advance. Therefore, it is not unlikely that the expectations on the capacity and the patient demand change. However, if for the first patients an advice is already given for specific weeks in the future, but the demand or capacity expectations change significantly for these weeks, it is possible to rerun the ILP of Chapter 5. The fact that some follow-up patients are already planned can easily be added to the constraints. The main change to the ILP would be that we have to add a constraint which ensures that the decision variable for the amount of patients being moved from one week to another, is at least equal to the number of patients that is already moved. Furthermore, part of the follow-up demand is known for these weeks, as there are already follow-up patients planned. Hence, we can substitute this realized demand as part of the expectation of the random variable.

The developed method is able to detect capacity problems far in advance. After solving the LP, it can be checked if the demand is satisfied for all weeks on which is decided. If the demand of a specific week is not met, then this is an indication of capacity problems around this week. Hence, the outpatient clinic has to increase capacity in this week or a week around this week. After this is done, the model can be solved again using the updated capacity.

According to the advice that can be generated, follow-up patients can be planned in specific weeks. Hence, capacity is allocated. On the shorter term, for example two or three months in advance, the MDP described in Chapter 4 can be solved in order to optimally plan the new patients. As input demand, an improved expectation of new patient demand can be used. As input capacity, the capacity that is already given to follow-up patients can be subtracted from the total capacity. By applying this method, the short-term planning can be improved relative to the resulting slot division from the long-term planning method.

7.3 Limitations and recommendations for future research

The Integer Linear Program that is developed for planning follow-up patients is deterministic. However, in the patient demand of both new patients and follow-up patients, stochasticity is involved. To deal with this stochasticity for the new patients, we included a robust estimate of their expected demand. However, for follow-up patients we included just the expected value. To this end, the Sample Average Approximation method can be used to also involve the stochasticity in the solution approach of the model, which we also discussed in this thesis. In further research, this method can be considered for solving our model. Moreover, research can be done on other methods that can deal with the stochasticity in the model and make it more robust. For example, in further research, the model could be formulated as a multiple stage stochastic programming model. In such a model, one could distinguish between follow-up patients based on their follow-up period. When first running the model, one could only make decisions for follow-up patients having the longest follow-up period. Decisions on follow-up patients having shorter follow-up periods can considered to be scenario-dependent and thus postponed.

Moreover, in further research the performance of this method could be tested in more detail and for more cases. To this end, also a sensitivity analysis could be executed. For example, more research could be done on parameter values, such as the choice of the chance a follow-up patient cannot be moved and thus how undesirable moving patients to particular weeks actually is.

Moreover, developed a method for providing a planning advice for each patient that will have a follow-up consultation in the future. More research could be done on generating this advice. Since the outcome of the model and the advice are based on the expected number of patients, further research may focus on providing advice in case less or more patients than expected need a follow-up consult.

To be able to provide an advice for each future follow-up patient, the model developed in Chapter 5 needs to be run far in advance. As indicated in Section 7.2, the model could be adapted to the already planned follow-up patients in order to rerun it with different input capacity or demand. Further research could focus on such an extension of the developed model. For example, it could be incorporated that the model is always run again if follow-up patients with the longest follow-up period already have been planned. The part of the expected patient demand that comes from these patients can be substituted by the realization and the planning decisions for these follow-up patients can be included as described in Section 7.2.

Furthermore, we truncated the horizon of infinite horizon Markov Decision Process in order to be able to solve it in a finite number of calculations. We approximated the salvage values, the state values at the truncated horizon, by the state values of a stationary infinite horizon Markov Decision Process. We proved that the resulting approximation error is bounded. However, the bound that was found is not really tight. In further research, more analysis on the resulting error should be done in order to find a better error bound, both analytically and numerically. For numerical analysis, one could execute a sensitivity analysis on the influence of the chosen salvage values on the value function and the optimal decision in the first decision epoch. Another limitation regarding the Markov Decision Process is the large computation time. If this model would be applied for examples that are larger than we considered in our research, it is necessary to look for other solution methods, such as Approximate Dynamic Programming.

Chapter 8

Conclusion

The goal of this thesis was to optimize patient planning at a cardiology outpatient clinic. To reach this goal, we developed a method, in which we consider patients that are new to the outpatient clinic and patients that need a follow-up consultation. The method allocates capacity among these patient groups in such a way that all patients have acceptable access times. For this method, we formulated two mathematical optimization models. In the first model, the arrival and planning process of patients that are new to the clinic is modelled as a Markov Decision Process. In the second model the problem of optimally planning follow-up patients is addressed and an Integer Linear Program is formulated. In this model, also the Markov Decision Process model is included, so that the arrival of new patients is also taken into account.

The applicability of the developed method for optimally planning follow-up patients is threefold. Firstly, capacity problems can be recognized far in advance. Secondly, an optimal division of capacity between new patients and follow-up patients can be deduces from the output of the model. And thirdly, an advice can be generated for planning follow-up patients on a patient level. The resulting method for capacity division and providing planning advice on a patient level is able to anticipate on variability in patient demand and capacity. In the resulting division of capacity, the variability is minimized. However, by allowing the division to be different from week to week, we act flexibly in order to anticipate on this variability in patient demand and capacity. The Markov Decision Process model can be used to optimally schedule new patients, after the follow-up patients have been planned according to the advice following from the Integer Linear Programming model. Furthermore, the developed method and models are not only applicable to a cardiology outpatient clinic, but also to other outpatient clinics.

By means of a simulation study, we tested the performance of the developed planning method. In this simulation study, we considered a setting that is close to that of the cardiology outpatient clinic. The simulation results show that with the developed method, new patients have lower access times and lower variability in access times than with a myopic method that is close to the current strategy of the hospital. Furthermore, the access times for follow-up patients are similar with both methods. Hence, on average the access time is lower with our developed method than with the myopic method with the same total capacity.

In conclusion, we developed a method that optimizes and supports patient planning at outpatient clinics in such a way that all patients have acceptable access times.

Bibliography

[1] M. E. Zonderland, *Appointment planning in outpatient clinics and diagnostic facilities*. Springer, 2014.

[2] "Wat zijn de aanvaardbare wachttijden volgens de treeknorm?" Oct 2021. [Online]. Available: https://www.promovendum.nl/zorgverzekering/vragen/wat-aanvaardbare-wachttijden-volgens-treeknorm

[3] "Over gelre ziekenhuizen." [Online]. Available: https://www.gelreziekenhuizen.nl/over-ons/

[4] 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*, vol. 1, no. 2, pp. 129–175, 2012. [Online]. Available: https://doi.org/10.1057/hs.2012.18

[5] A. Ahmadi-Javid, Z. Jalali, and K. J. Klassen, "Outpatient appointment systems in healthcare: A review of optimization studies," *European Journal of Operational Research*, vol. 258, no. 1, pp. 3–34, 2017. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0377221716305239

[6] C. Falces, J. Sadurní, J. Monell, R. Andrea Riba, M. Ylla, A. Moleiro, and J. Cantillo, "One-stop outpatient cardiology clinics: 10 years' experience," *Revista española de cardiología*, vol. 61, pp. 530–3, 06 2008.

[7] L. A. Apergi, J. S. Baras, B. L. Golden, and K. E. Wood, "An optimization model for multi-appointment scheduling in an outpatient cardiology setting," *Operations Research for Health Care*, vol. 26, p. 100267, 2020. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S2211692320300473

[8] A. G. Leeftink, I. A. Bikker, I. M. H. Vliegen, and R. J. Boucherie, "Multi-disciplinary planning in health care: a review," *Health Systems*, vol. 9, no. 2, pp. 95–118, 2020. [Online]. Available: https://doi.org/10.1080/20476965.2018.1436909

[9] A. Schneider, "Integral capacity management planning in hospitals," Ph.D. dissertation, University of Twente, Netherlands, Aug. 2020.

[10] P. Hulshof, R. Boucherie, E. Hans, and J. Hurink, "Tactical resource allocation and elective patient admission planning in care processes," *Health care management science*, vol. 16, no. 2, pp. 152–166, 2013.

[11] P. Hulshof, M. Mes, R. Boucherie, and E. Hans, *Tactical planning in healthcare using approximate dynamic programming*, ser. Memorandum. University of Twente, Department of Applied Mathematics, Sep. 2013, no. 2014.

[12] S. Creemers, J. Beliën, and M. Lambrecht, "The optimal allocation of server time slots over different classes of patients," *European Journal of Operational Research*, vol. 219, no. 3, pp. 508–521, 2012. [Online]. Available: https://ideas.repec.org/a/eee/ejores/v219y2012i3p508-521.html

[13] X. Qu, Y. Peng, N. Kong, and J. Shi, "A two-phase approach to scheduling multi-category outpatient appointments – a case study of a women's clinic," *Health care management science*, vol. 16, 03 2013.

[14] T. B. T. Nguyen, A. I. Sivakumar, and S. C. Graves, "Capacity planning with demand uncertainty for outpatient clinics," *European Journal of Operational Research*, vol. 267, no. 1, pp. 338–348, 2018. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0377221717310494

[15] N. Aslani, O. Kuzgunkaya, N. Vidyarthi, and D. Terekhov, "A robust optimization model for tactical capacity planning in an outpatient setting," *Health Care Management Science*, vol. 24, no. 1, pp. 26–40, March 2021. [Online]. Available: https://ideas.repec.org/a/kap/hcarem/v24y2021i1d10.1007_s10729-020-09528-y.html

[16] M. L. Puterman, "Markov decision processes: Discrete stochastic dynamic programming," in *Wiley Series in Probability and Statistics*, 1994.

[17] A. Ghate and R. L. Smith, "A linear programming approach to nonstationary infinite-horizon markov decision processes," *Operations Research*, vol. 61, no. 2, pp. 413–425, 2013. [Online]. Available: http://www.jstor.org/stable/23481839

[18] W. J. Hopp, "Identifying forecast horizons in nonhomogeneous markov decision processes," *Operations Research*, vol. 37, no. 2, pp. 339–343, 1989. [Online]. Available: http://www.jstor.org/stable/171513

[19] J. C. Bean, J. R. Birge, and R. L. Smith, "Aggregation in dynamic programming," *Operations Research*, vol. 35, no. 2, pp. 215–220, 1987. [Online]. Available: http://www.jstor.org/stable/170693

[20] A. massoud Farahmand, D. Nikovski, Y. Igarashi, and H. Konaka, "Truncated approximate dynamic programming with task-dependent terminal value," in *AAAI*, vol. 30, 2016. [Online]. Available: https://ojs.aaai.org/index.php/AAAI/article/view/10397

[21] N. M. van Dijk, "Error Bounds and Comparison Results: The Markov Reward Approach For Queueing Networks," in *Queueing Networks*, ser. International Series in Operations Research & Management Science, R. J. Boucherie and N. M. van Dijk, Eds. Springer, March 2011, ch. 9, pp. 397–459. [Online]. Available: https://ideas.repec.org/h/spr/isochp/978-1-4419-6472-4_9.html

[22] M. Overmars, “Analyzing the convergence of q-learning through markov decision theory,” Master’s thesis, University of Twente, 2021.

[23] A. Mastin and P. Jaillet, “Loss bounds for uncertain transition probabilities in markov decision processes,” *2012 IEEE 51st IEEE Conference on Decision and Control (CDC)*, pp. 6708–6715, 2012.

[24] L. A. Hannah, “Stochastic optimization,” in *International Encyclopedia of the Social Behavioral Sciences (Second Edition)*, second edition ed., J. D. Wright, Ed. Oxford: Elsevier, 2015, pp. 473–481. [Online]. Available: https://www.sciencedirect.com/science/article/pii/B9780080970868420106

[25] A. J. Kleywegt, A. Shapiro, and T. Homem-de Mello, “The sample average approximation method for stochastic discrete optimization,” *SIAM Journal on Optimization*, vol. 12, no. 2, pp. 479–502, 2002. [Online]. Available: https://doi.org/10.1137/S1052623499363220

[26] A. Bhattacharya and J. Kharoufeh, “Linear programming formulation for non-stationary finite-horizon, markov decision process models,” *Operations Research Letters*, vol. 45, 09 2017.

[27] G. Neustroev, M. de Weerdt, and R. Verzijlbergh, “Discovery of optimal solution horizons in non-stationary markov decision processes with unbounded rewards,” *Proceedings of the International Conference on Automated Planning and Scheduling*, vol. 29, no. 1, pp. 292–300, May 2021. [Online]. Available: https://ojs.aaai.org/index.php/ICAPS/article/view/3491

[28] “Python.” [Online]. Available: https://www.python.org/

[29] “The fastest solver,” Mar 2022. [Online]. Available: https://www.gurobi.com/

[30] W. B. Powell, *Approximate Dynamic Programming: Solving the curses of dimensionality*. John Wiley & Sons, 2007, vol. 703.