UNIVERSITY OF TWENTE.

Master Thesis

Stochastic Operations Research

Optimal time allocation of an orthopedic surgeon

Eline R. Tsai

Sint Maartenskliniek

Graduation Committee

Prof. Dr. Richard Boucherie (University of Twente, supervisor)
Dr. Ir. Werner Scheinhardt (University of Twente)
Prof. Dr. Marc Uetz (University of Twente)
Ir. Rob Vromans (Sint Maartenskliniek, supervisor)

August 29, 2017

Abstract

The goal of this project is to determine the optimal allocation of the outpatient department and operating room (OR) sessions per orthopedic surgeon. The amount of OR and OD sessions that can be distributed over the year is budgeted per surgeon. The allocation of these sessions influences the OR queue length. The optimal OR queue length is a trade-off between the waiting time of the patients in the OR queue and the amount of unused OR time. In general, hospitals aim at giving patients the best treatment as quickly as possible. Inefficient use of OR time leads to unnecessary health care costs.

The system is modeled as a Markov chain and we use stochastic dynamic programming to determine the optimal allocation of the sessions. Performance measures are translated into queue length boundaries and uncertainties regarding a planning horizon are considered. The initial model is very detailed and computationally expensive, which is why several approximations and simplifications are proposed. These simplifications cause a significant reduction in the size of the state space and the computation time of the transition probabilities. Due to the complicated structure of the system, more research on approximation methods is needed.

Acknowledgments

First of all, I want to thank my supervisor Richard Boucherie for being my mentor and providing me with advice and new ideas whenever needed. You also gave me the opportunity to go to Canada for my internship, which was a great learning experience. A special thanks goes to my supervisor Rob Vromans for his advice and his enthusiasm. You always tried to make me think more critically about the practical applicability of my work. I also want to thank Marc Uetz and Werner Scheinhardt for taking the time to read and assess my work.

The CHOIR group gave me a very warm welcome. Thanks a lot for that! I am very grateful that I was allowed to go to ORAHS with you. My very first conference, where I gave a presentation as well. It was scary and exiting at the same time. Thank you for supporting me and helping me calm down my nerves. ORAHS, the summer school, the meetings and our lunch walks motivated me and made the whole final project experience more enjoyable.

I would like to thank Ingeborg for checking some of my computations and even help me correct one. I also want to thank Linde for her ideas, support and for letting me borrow her laptop for some computations.

I am very grateful for my parents and sister. Even though there is an ocean between us, you still manage to be a good support system. Hope to see you again soon!

Finally, I want to thank Kit. I am grateful for all your love and support over the past couple of months. You have always showed a readiness to listen.

Contents

1 Introduction

The Sint Maartenskliniek (SMK) is the only hospital in the Netherlands fully specialized in movement and posture [15]. The SMK has locations in Nijmegen, Boxmeer, and Woerden. The SMK has several departments such as, the rheumatology department, the department of orthopedic surgery, sport medical department and the rehabilitation clinic. The focus of this project in on the department of orthopedic surgery.
The first section of this chapter describes the possible routes of a patient through the orthopedic chain. Then the different activities of an orthopedic surgeon, patient types, and levels of planning will be explained. The final two sections of this chapter list the goals of this project and give an outline of this thesis.

1.1 Patient pathway

Figure 1 shows a schematic depiction of the patient pathway. When patients arrive at the SMK, they first visit the outpatient department (OD) for an initial consultation (NC) with an orthopedic surgeon. Subsequently the surgeon can refer the patient either to the radiology department, to the OD for a follow-up consultation (VC) or terminate the patient's treatment at the SMK. From this point onward we will refer to an orthopedic surgeon as a surgeon. When a patient needing surgery finishes radiology he receives an OR-ticket. Having an OR-ticket, the patient is part of the queue for a screening appointment. We refer to the diagnostics process as all the visits of the patient to the radiology department and consultations at the OD in between. Patients always have to finish their diagnostic process before entering screening and they always have to go through screening before they can go into surgery. Patients who are referred to the radiology or the screening department are put on a waiting list together with patients from other surgeons, When a patient has finished his screening process successfully, i.e., the patient is fit to receive surgery, the patient enters the OR queue of the surgeon who will perform the surgery on this patient. Usually, a patient sees the same surgeon at the outpatient department (OD) and in the operating room (OR).
After surgery, patients are admitted at the recovery ward. It rarely happens that a patient cannot have surgery due to unavailability of beds in the recovery ward [16]. A patient's treatment can be terminated during any of the steps given in Figure 1, except for the OR and the recovery ward. After their stay at the recovery ward, patients always have a dismissal consultation (OC) before they finish their treatment. Here the surgeon decides whether the treatment has been successful. The surgeon can also decide that the patient needs another consultation at the OD (a VC) or that the patient needs further surgery. As said before, patients always have to go through screening before they can go into surgery. This means that these types of patients also have to go through screening again.

1.2 Activities of a surgeon

During a week a surgeon has several duties at the SMK. These duties include sessions at the OR or the OD, but also duties like administration, education and management. One session is defined as a full business day. In this thesis we focus on the OR and OD sessions of the surgeon and ignore the other activities and the corresponding planning rules. We focus on the OR and the OD sessions, because these are the activities where the surgeon is in contact with his patients. The surgeon sees patients during consultations at the OD and when he is in the OR performing

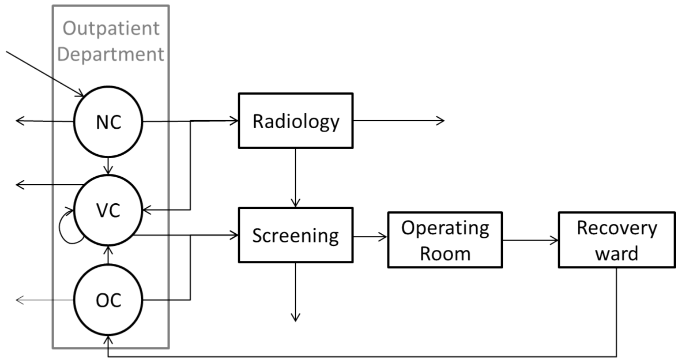

Figure 1: Patient pathway

surgery. The allocation of these sessions influences the OR queue length. The more OD sessions done in a certain period, the more patients will be in the OR queue in the weeks to follow. The more OR sessions done in a certain period, the more patients will leave the OR queue in that period.

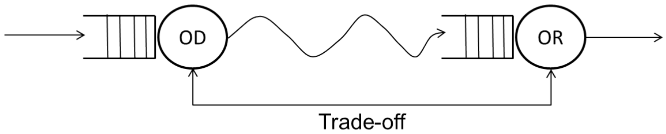

There are different types of OD consultations. Some of these were already introduced in the previous section. There are four main types of OD consultations: NC, VC, OC and POS. The different types of OD sessions will not explicitly be included in the model. Patients who switch between surgeons have a POS consultation to meet their new surgeon. The decision to switch is made by the SMK. Patients who switch have already finished screening and go from the POS consultation directly to the OR queue. One of the reasons for a patient to switch between surgeons is because the patient needs a type of surgery in which another surgeon is specialized.

1.3 Levels of planning

The SMK distinguishes between several levels of planning (see Figure 2). Decisions made at each level should correspond to decisions made at higher levels in order to realize the budget set by the SMK for the year. Furthermore, the amount of staff should correspond to the production goals, which is why the levels of staff planning and levels of session planning are connected as well.
At the beginning of the year, the *Human Resources Budget* and the *production goals* are set. The Human Resources Budget describes the amount of staff needed in a certain year. This budget corresponds to the production goals, i.e., to the amount of OR and OD sessions each surgeon has to do that year. The amount of OR and OD sessions a surgeon has to do is based on several factors and differs per surgeon. The average surgery duration of a surgeon influences the average amount of surgeries the surgeon does during an OD session. The OR hitrate (h_{OR}) is the average number of surgeries a surgeon performs per OR session. The OD hitrate is the average number of consultations

needed for one OR-ticket. If a surgeon has an OD hitrate of 7 and has consultations with 70 patients at the OD in a certain week, then we assume that 10 OR-tickets will be created due to these OD sessions. These hitrates and the amount of supporting staff affect the budgeted amount of OD and OR sessions of a surgeon.
The *long term leave planning* and the *production plan* describe per month the number of holidays the surgeon has, the number of days the surgeon is at work, and his activities.
Based on the *minimal staffing requirements* and the *updated production plan*, the amount of staff and the amount of each session per surgeon per week is known. At the moment the updated production plan is made by hand. This plan is made two months prior to the date it is used. Each updated production plan covers a period of four consecutive weeks. A good updated production plan corresponds to decisions made at a higher level and the minimal staffing requirements. Such a updated production plan also results in a balance between the waiting time of the patients in the OR queue and the amount of unused OR time slots.
Given the minimal staffing requirements, an *activity schedule* is made. A day is split up in two parts. The activity schedule shows on which days and which part of the day during that week the sessions will take place.
Based on the activity schedule an OD-grid is assigned to each of the OD sessions, i.e., the *time slot schedule* is determined. This grid assigns to each time slot a patient type (NC, VC, OC or POS). The OD-grid can differ per session depending on the amount of supporting staff the orthopedic surgeon has during a session. The schedulers who work on these levels of planning (grid planners) have no information about the patients currently in the system. For example, they do not know how many patients are on the waiting list for a VC. The *patient schedulers* have this information. They assign patients to the time slots that are reserved for their type.
The main focus of this thesis is on the updated production plan. A good production plan meets the budget set at the beginning of the year and will give the schedulers in the lower levels of planning a good standard schedule.

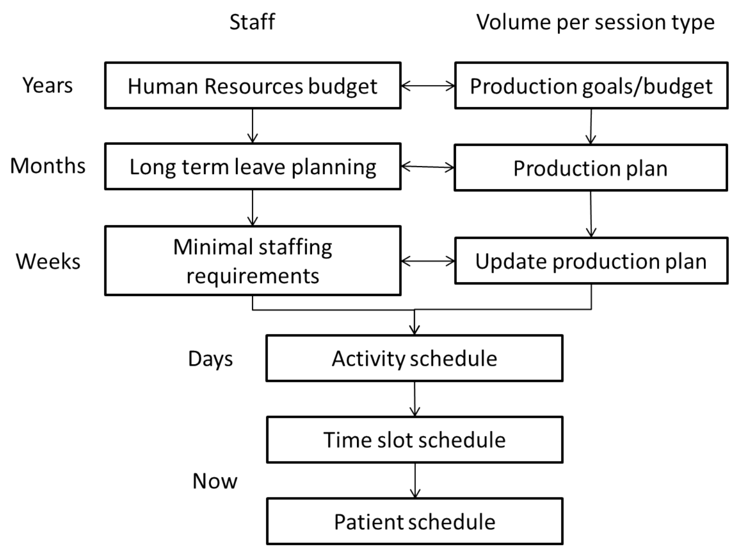

Figure 2: Different levels of patient and staff planning

1.4 Types of patients

There are six different units in the orthopedic chain: knee, spine, hip, upper extremity, child, and knee and ankle. One orthopedic surgeon can treat patients from different units. There are roughly 5 different urgency classifications for orthopedic patients (see Table 1). The internal access time to the OR is defined as the time between the moment the patient receives an OR-ticket and the moment the patient goes to the OR for surgery. Patients who require acute care usually arrive at the SMK in an ambulance and are scheduled at the OR in one of the slots reserved specially for these patients. At the SMK, a distinction is made between surgeries without an implant, surgeries for primary implants, and revisions on implants. Due to the length and complexity of revision surgeries, the SMK does not plan more than 2 of these surgeries during one OR day of a surgeon. Patients who need a revision surgery also get priority at the OD.
Furthermore, there are time slots reserved at the OR for urgent patients. If these time slots are not fully occupied two weeks prior to the OR session, these time slots will also be made available for non-urgent patients.

Urgency category	Maximum internal access time to the OR
Elective	6 months
2 Months	2 months
1 Month	1 month
Emergency	2 weeks
Acute	48 hours

Table 1: Urgency classifications

1.5 Goals

In general, the goal of the SMK is for patients to receive the best treatment as quickly as possible. Which implies that the SMK wants to minimize the external access time to the OD and the internal access times to screening, radiology, and the OR. The external access time to the OD depends on the OD capacity of the surgeons. The access time to screening depends on the timing of the screening sessions and the availability of anesthesiologists and supporting staff. The access time to radiology depends on the timing of the radiology sessions and availability of radiologists and supporting staff. The allocation of the OD and OR sessions affects the access time to the OR. The access time to the OR affects the utilization of the OR. Inefficient use of OR time leads to unnecessary health care costs, which is one of the main reasons why the SMK wants to keep the unused OR time below a certain level. We have two main performance measures:

- Minimize the waiting time of the patients in the OR queue
- Keep the amount of unused OR time below a certain level.

Our *goal* is to determine the optimal allocation of the OD and OR sessions per surgeon per week throughout the year. An allocation is optimal if it minimizes the waiting time of the orthopedic patients in the OR queue, while keeping the amount of unused OR time below a certain level. To be able to determine this optimal allocation, we need to know what the optimal queue length boundaries are. The more patients in the OR queue of a surgeon, the easier it is to fully fill the reserved OR time of this surgeon, but the longer patients have to wait before they can go into surgery. The fewer patients in the OR queue, the faster patients can receive surgery, but the more OR time remains empty.

1.6 Thesis outline

In Chapter 2 give an overview of related literature. In Chapter 3 we explain how the orthopedic chain can be modeled as a Markov chain. We explain what the relevant parameters and variables are. In this section we also derive the one-step transition probabilities. Chapter 4 explains how stochastic dynamic programming can be used to determine the optimal allocation of the operating room and outpatient department sessions.Queue length boundaries will be constructed in this chapter as well. In the final section of this chapter we address several possibilities of how uncertainties regarding a planning horizon can be incorporated in the model. The detailed model described in Chapter 3 is computationally expensive, which is why we explain several approximations in Chapter 5. Section 5.1 presents an alternative, less detailed model. In Section 5.2 and Section 5.3 we explain how we can approximate the transition probabilities in order to speed up the computations. The possibility of having a monotone policy will be explained in Section 5.4. The final section of this chapter presents some test case results. In Chapter 6 we explain how we came to the transition probabilities we mostly use in this thesis. These probabilities are partly based on data and partly on expert opinions. Finally, in Chapter 7 we conclude and discuss our research. Recommendations for future research are given as well.

2 Literature review

Orthopedic patients can visit several departments before their treatment is finished, such as the OD and the OR in the orthopedic chain, the radiology department, screening and the recovery ward. Optimal planning is required throughout each of these departments in order to provide the best treatment as quickly as possible. In the first two sections willl give a brief overview of research done in patient admission control and OR scheduling. Then we have a look at previous work done in the SMK on the allocation of OD and OR sessions. Literature on a myopic policy for our system is explored as well. In the last two sections we talk about literature on Markov chains, dynamic programming techniques and we refer to interesting research done by Hulshof et al [9]. Our system has a lot of similarities with that of [9] and we explain here what the differences are.

Patient admission control
Hospitals can profit from controlling patient admissions. Nunes et al. [10] model the control of patient admissions as a MDP. The paper considers elective admissions. Treatment patterns are used to classify the progression of the patients during predefined treatment periods. An optimal stationary policy is found by using value iteration. The authors refer to papers by Kapadia et al [7] [8], who provide a method to estimate probabilities of admissions of patients in any treatment patters. In this thesis we refer to [8] when computing transition probabilities from the data.
In the SMK, according to [16], the formats of the OD-grids used in the time slot schedule are a result of Project Gemini, 2013. Recall that these grids assign to each time slot a type of OD consultation, for example a NC. This grid results in patient admission control.

OR scheduling
A master surgery schedule (MSS) assigns blocks of OR time to different surgical specialties or patient types [5] [19]. OR scheduling is a frequently studied research topic. According to [5], many MSS approaches only consider the effects of the MSS on the OR and OR staff. However, only few approaches focus on the relation of the OR with other departments. For example, Vanberkel et al [19] and Fügener et al [5] focus on the effects of an MSS on the downstream departments. Vanberkel et al., provide the foundation for a decision support tool to relate the MSS to the workload of downstream departments. Fügener et al. expand this approach by including ICU bed requests and multiple downstream wards instead of a single downstream unit. Fügener et al. also assign costs to specific MSSs and introduce several algorithms to minimize downstream costs.

Previous study in the SMK on allocation of OD and OR sessions
The goal of this study is to find the optimal allocation of the OD and OR sessions per surgeon. In [16] the optimal allocation of OD and OR sessions in the SMK is computed per unit. The author used linear programming to find a solution but due to the complexity of this model and the computation time a heuristic was developed for practical use. This heuristic keeps track of the ratio between the number of patients from a certain unit that has already been treated and the production goal of this unit, and assigns the sessions accordingly. In [1], the optimal allocation of the OD and OR sessions in the SMK is done per unit as well. Different approaches that use an integer linear program as a basis are compared. The authors of these reports use average delays between states, for example between the OD and the OR. In doing so they do not include the variability in the delay caused by screening and diagnostics. Furthermore, because one surgeon can treat several units, it can be interesting to determine the optimal allocation of the sessions per surgeon instead.

Myopic policy: (s, S) inventory policy
A myopic policy for our system could be an (s, S) inventory policy. Here we could, for example, assign to each week an average number of OD and OR sessions. If the OR queue length is below the

lower boundary s, we can do a maximum possible amount of OR sessions until the OR queue length is at least equal to S. Tijms [17] describes the basics concepts of these policies. Our system has a random lead time, because the amount of time a patient spends in the system between a NC and before entering the OR queue is uncertain. We have a variable production rate, because the number of patients seen at the OD can differ per week. Lastly, we have a system with periodic review. One paper that addresses systems with random lead times in which orders are allowed to cross in time is the paper by Bashyam and Fu [2]. They also consider service level constraints, as it is hard to measure customer or patient satisfaction. This policy is a myopic one, because then we decide solely on the OR queue length how many OR and OD sessions to assign to a certain week. However, the number of patients in diagnostics and screening needs to be considered as well. Even though this approach is myopic, it can be used to get some insight on boundaries for the OR queue length in our approach.

Dynamic programming

A common way to model patient flow through a hospital is by using a Markov chain. Markov chains have the property that no information prior to the current time step is needed to make predictions about the state in the next time step. The basic concepts of Markov chains can be found in Ross [13]. Because our system is modeled as a Markov chain and we have a budget to stick to, we use stochastic dynamic programming (SDP) to find the optimal allocation of the OD and the OR sessions.

Due to the curse of dimensionality, as the size of the problem increases, solving a dynamic program is typically intractable [9]. Approximate dynamic programming (ADP) provides a a powerful set of strategies for problems that are difficult to solve due to their size [11]. Powell [11] gives an extensive overview of several ADP techniques. There are also several tricks we can apply to the SDP in order to reduce the computation time, for example time aggregation. Cao et al. [4] use time aggregation to solve an infinite horizon MDP. According to Cao et al. State aggregation generally results in an approximation of the original model while time aggregation does not result in loss of accuracy. The Markov property is reserved with time aggregation. In order to use the time aggregation approach proposed by Cao et al., it must be possible to divide the state space into two complementary subsets $\mathcal{S}_1$ and $\mathcal{S}_2$. These subsets must be such that actions are only taken when the system is in a state that is in $\mathcal{S}_1$ and the transition probabilities for states in $\mathcal{S}_2$ do not depend on actions.

Similar research by Hulshof et al.

The system described in this thesis has a lot of similarities with the system of Hulshof et al [9], who consider a more generic system. They decide at each decision epoch how many patients to treat from each queue that have already been waiting a certain amount of time. The authors focus on the waiting time of the patients and do not consider idle time of the queues. Since inefficient use of OR time can lead to higher health care costs, we look at the amount of unused OR time as well. The action space in [9] is restricted by the amount of available resources. The amount of these resources is fixed and can differ per week. Our system has more restrictions which makes it more difficult to find the optimal solution using their approach. We have a budgeted amount of OD and OR sessions to stick to. We also have more restrictions on the state space, mainly on the the OR queue length. These restrictions are the result of the fact that the optimal queue length is a trade-off between the waiting time of the patients and the amount of unused OR time. In this thesis we focus on the construction of an SDP for our problem. We construct an exact SDP and also explain our attempts at approximations in trying to make the SDP numerically tractable. An ADP approach similar to [9] can be used as well, but first the exact SDP has to be constructed. The output of the ADP can then be compared with the output of the exact SDP to get an idea of how accurate the approximations in the ADP are. Due to time constraints we did not manage to include the ADP approach in this thesis.

3 Markov chain

In this section we transform the orthopedic chain into a Markov chain. First we list the assumptions and simplifications made. Then we explain the possible transitions. We are also going to look at the mean performance of the system to give us an idea of how the one-step probability matrix can be computed.

3.1 Assumptions

This section explains the assumptions and simplifications made.

- The computations are made per surgeon. We do not include the total OR capacity, i.e., when scheduling the OR sessions of a certain surgeon, we do not take into account the possibility that the ORs may already be booked by other surgeons.
- Some thought is put into the ratio of the amount of time a surgeon has to do each activity. Which is why we assume that if we stick to the budgeted amount of OD and OR sessions for the year, we can ignore the other activities of the surgeons.
- We assume a fixed duration per surgery. Which means that we do not explicitly take patients into account that require a revision surgery.
- The different types of OD sessions (NC, VC and OC) will not explicitly be included. The probability that a patient goes from a NC to diagnostics may differ from the probability with which a patient goes from a VC to diagnostics. For our model we merge these probabilities to one probability. We make this assumption because we do not know exactly how many patients are planned for each of these consultations by the patient planners during the OD sessions.
- Not all types of priority patients are included. We only include the urgency labels: emergency, 1 month, 2 months and elective. Patients who require acute care are excluded because they are scheduled on one of the time slots specially reserved for these patients. The chosen number of OD and OR sessions does not affect these patients.

3.2 Network of queues

In this section we describe how the orthopedic chain can be modeled as a network of queues. Figure 3 shows a schematic depiction of this network. The number of patients seen at the OD in a certain week depends on the chosen number of OD sessions for that week. A fraction of these patients is again referred to the OD for another consultation, a fraction leaves the OD because their treatment at the SMK is terminated and a fraction is referred to radiology, so they start their diagnostic process, or to screening. The number of patients leaving the OR queue in a certain week depends on the number of OR sessions planned for that week, the number of patients in the OR queue and the rescheduling probability r_t in that week. The rescheduling probability for week t is the probability that a patient in the OR queue in week t does not want to be scheduled for the OR in that week. During the summer holidays these probabilities are usually higher compared with the rest of the year because in that period a lot of patients are on holiday.

A patient is in the state diagnostics from the first moment he is referred to the radiology department. The patient leaves this state as soon as he receives an OR ticket. The state diagnostics is a combination of radiology and OD visits. Recall that in the screening and radiology department, patients of all the surgeons are put on the same waiting lists. Furthermore, we do not know the schedule of the radiology and screening staff. These queues are considered as autonomous processes, which is why we model the screening queue and the diagnostics queue as infinite server queues. This implies that once a patient has entered one of these queues, his service starts immediately.

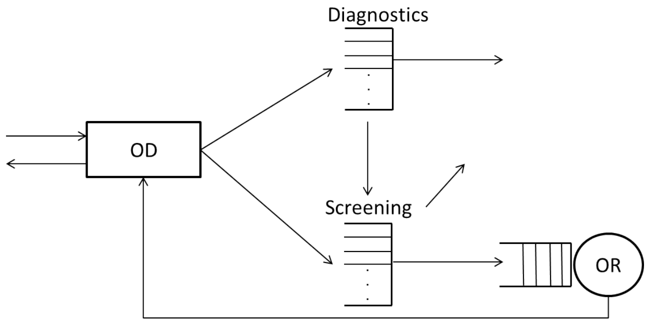

Figure 3: Adjusted patient pathway

There are several ways to model the screening and diagnostics queues. The most detailed one is by keeping track of how long how many patients have already been in either of these queues. One could also take a fixed delay caused by either of these queues, but we want to see the effect of a certain amount of OD sessions in a certain week on the OR queue in the next week and the weeks to follow. According to [3], from 2012 to 2016, patients had an average diagnostics delay of approximately 19 weeks. Which means that the probability that a patient stays one week in diagnostics can differ significantly from the probability that a patient stays 19 weeks in diagnostics. Define $p_R(i)$ to be the probability that a patient leaves diagnostics after i weeks and $p_S(i)$ to be the probability that a patient leaves screening after i weeks. The patient can then either exit the system or the patient can be referred to the next stage in the orthopedic chain. The patient stays for at least one more week in diagnostics and one more week in screening with probability $1 - p_R(i)$ and $1 - p_S(i)$ respectively. In order to have a finite state space we put boundaries on the length of stay in diagnostics and screening denoted by z_R and z_S respectively. Note that it is possible that a patient can have an OD consultation in week t, be referred to radiology in that same week and finish his diagnostics in that week. This is denoted by the probability $r_R(0)$. A similar argument holds for $r_S(0)$. Due to the structure of the diagnostics and screening queue, the time a patient spends in either of these queues is a Coxian random variable [14].

3.3 Mean performance (Model 1)

In this section we compute the conditional expectation of random variables used in the model. We compute these expected values to get some insight on the relations between the states and the corresponding transition probabilities. Furthermore, we will use them as a mnemonic for the computations of the one-step transition probabilities of the Markov chain. Also, the conditional expected queue length will later on be used in the cost function of our SDP.

Take $R_{t,i}$ and $T_{t,i}$ to be the number of patients that complete their i-th week in respectively diagnostics and screening by the end of week t. The action chosen for week t is the amount of OD and

OR sessions the surgeon will perform in that week: $a_t = (a_{t,OD}, a_{t_OR})$. Take C_t to be the number of patients that can be scheduled at the OD in week t. D_t is the number of patients that can be seen at the OR in week t. Take h_{OD} to be the average number of time slots during an OD session and recall that h_{OR} is the OR hitrate and equal to the average number of surgeries a surgeon performs per OR session. The following holds:

$$C_t = h_{OD} a_{t,OD}$$
$$D_t = h_{OR} a_{t,OR}$$

The number of time slots can differ per OD session. This amount depends on the amount of supporting staff a surgeon has during an OD session. The state of the system at the beginning of week t is denoted by $\sigma_t = (R_t, T_t, X_t)$. The number of patients in the OR queue at the beginning of week t is denoted by X_t. Take $T_t = \{T_{t,0}, T_{t,1}, ..., T_{t,z_S}\}$ and $R_t = \{R_{t,0}, R_{t,1}, ..., R_{t,z_R}\}$. Figure 4

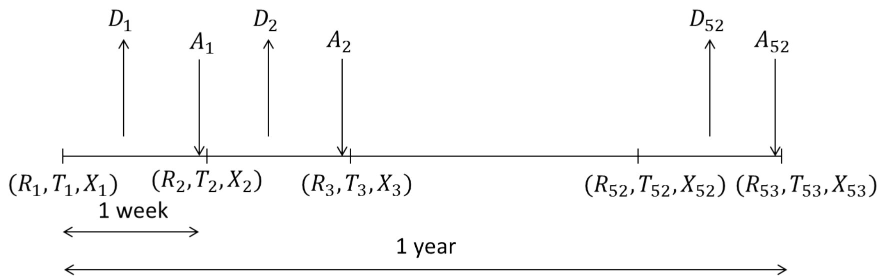

Figure 4: Timeline of arrivals (A_t) and departures (D_t) of OR queue

shows a timeline of the arrival and departure process of the OR queue. Patients who arrive at the OR queue in week t can be scheduled for the OR in week $t + 1$ the earliest. Figure 5 shows most of the relevant probabilities for this model. The rescheduling probability is not included in the figure. The probability that a patient that is seen at the OD in week t is referred to the state diagnostics is equal to $p_{OD,R}$. The probability that a patient goes directly to screening after an OC is denoted by $p_{OD,T}$. Recall that $r_R(i)$ is the probability that a patient exits diagnostics after being exactly i weeks in this state. The probability that this patient then goes to screening is denoted by $r_{RS}(i)$. A patient's treatment is terminated after i weeks of screening with probability $r_R(i)1 - r_{RS}(i)$. A patient that leaves screening after i weeks goes to the OR queue with probability $r_{SO}(i)$.

Our internal transition process follows a multinomial distribution, similar to [9]. The internal transition process is defined as transitions between the OD, diagnostics, screening and the OR. This includes arrivals to the OR queue, but not the OR queue length. Assume that patients go from state i in week t to state j in week $t + 1$ with probability p_{ij}. Using similar notation as in [9], the random variable W_{ij} denotes the number of patients that goes from queue i to queue j in a certain week and w_{ij} denotes its realization. Take $\mathbf{W}_i = (W_{i1}, W_{i2}, ..., W_{iJ})$, with J equal to the number of states patients can visit. This includes the state in which the treatment of the patients has been terminated. The following holds:

$$P\Bigg(\mathbf{W}_i = (w_{i1}, w_{i2}, ..., w_{iJ})\Bigg) = \binom{w_i^d}{w_{i1}, w_{i2}, ..., w_{iJ}} \prod_{j=1}^{J} p_{ij}^{w_{ij}},$$

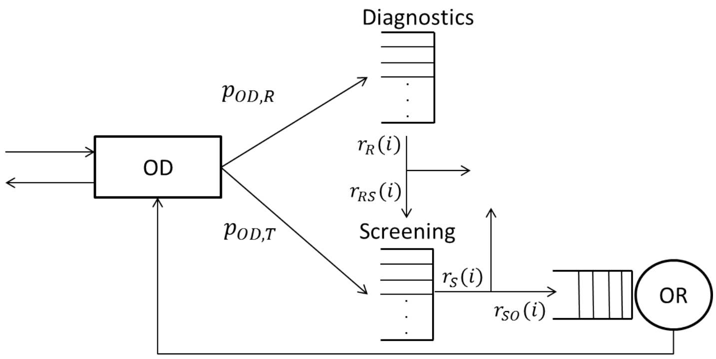

Figure 5: Adjusted patient pathway with probabilities

with w_i^d the number of patients that are in state i at the beginning of a certain week. So $\mathbf{W}_i$ follows a multinomial distribution, which means that its mean vector is equal to $\boldsymbol{\mu}_i = w_i^d \mathbf{p}_i$, with $\mathbf{p}_i = (p_{i1}, p_{i2}, ..., p_{iJ})$. The covariance matrix is equal to $\Sigma_i = \text{diag}(\mathbf{p}_i) - \mathbf{p}_i^T \mathbf{p}_i$. The diagonal matrix with the elements of $\mathbf{p}_i$ on the diagonal is denoted by $\text{diag}(\mathbf{p}_i)$. Note that Σ_i is singular.

Using the fact that the process follows a multinomial distribution and using the probability notation of Figure 5. We can now compute the conditional expected values of the relevant variables. The number of arrivals to diagnostics in week t depends on the number of patients seen at the OD in week t. So the expected number of patients that is referred to radiology in week t is equal to:

$$E[R_{t,0}|a_t] = p_{OD,R}C_t$$

The number of patients that will complete $i+1$ weeks of being in state diagnostics in week t is equal to the number of patients that completes i weeks of diagnostics in week $t-1$ and stays in that state for at least one more week:

$$E[R_{t,i+1}|\sigma_{t-1}, a_{t-1}] = (1 - r_R(i))R_{t-1,i}, \text{ with } i \in \{0, 1, .., z_R - 1\}$$

The number of patients arriving at the OR queue in week t is equal to the amount of patients that leaves state screening in week t:

$$E[A_t|\sigma_t, a_t] = \sum_{i=0}^{z_S} r_S(i) r_{SO}(i) T_{t,i}$$

The number of referrals to screening in week t is a combination of the patients who just had an OC and the patients who are done with diagnostics.

$$E[T_{t,0}|R_{t,0}, R_{t,1}, ..., R_{t,z_R}, a_t] = p_{OD,T}C_t + \sum_{i=0}^{z_R} r_R(i) r_{RS}(i) R_{t,i}$$

The number of patients that will complete $i+1$ weeks of being in state screening in week t is equal to the number of patients that completes i weeks of screening in week $t-1$ and stays in that state for at least one more week:

$$E[T_{t,i+1}|\sigma_{t-1}, a_{t-1}] = (1 - r_S(i))T_{t-1,i}, \text{ with } i \in \{0, 1, .., z_S - 1\}$$

These equations imply that T_t depends on R_t. Which is not desirable if we want to model the system as a Markov chain, because we want the probability that we are in a certain state in week t to only depend on information in week $t-1$. Change T_t and R_t into $T_t = \{T_{t,1}, T_{t,2}, ..., T_{t,z_S}\}$ and $R_t = \{R_{t,1}, R_{t,2}, ..., R_{t,z_R}\}$. Instead of the equations mentioned before, we will use the following set of equations:

$$\begin{aligned}
E[R_{t,1}|\sigma_{t-1}, a_{t-1}] =& (1 - r_R(0))p_{OD,R}C_{t-1} \\
E[R_{t,i+1}|\sigma_{t-1}, a_{t-1}] =& (1 - r_R(i))R_{t-1,i}, \text{ with } i \in \{1, 2, .., z_R - 1\} \\
E[T_{t,1}|\sigma_{t-1}, a_{t-1}] =& (1 - r_S(0))\left(p_{OD,T}C_{t-1} + r_R(0)r_{RS}(0)p_{OD,R}C_{t-1} + \sum_{i=1}^{z_R} r_R(i)r_{RS}(i)R_{t-1,i}\right) \\
E[T_{t,i+1}|\sigma_{t-1}, a_{t-1}] =& (1 - r_S(i))T_{t-1,i}, \text{ with } i \in \{1, 2, .., z_S - 1\} \\
E[A_t|\sigma_t, a_t] =& \sum_{i=1}^{z_S} r_S(i)r_{SO}(i)T_{t,i} + \sum_{i=1}^{z_R} r_R(i)r_{RS}(i)r_S(0)r_{SO}(0)R_{t,i} \\
&+ \big(p_{OD,T} + r_R(0)r_{RS}(0)p_{OD,R}\big)r_S(0)r_{SO}(0)C_t
\end{aligned}$$

The states $R_{t,0}$ and $T_{t,0}$ are now not explicitly used to construct these relations. Which is why they are not included in the state description. Now T_t depends on R_{t-1}.

Lastly, recall that it is possible for patients to switch between surgeons. After a POS consultation at the OD with the new surgeon the patient immediately enters the OR queue. This is implicitly included in the model. It is included in the probability $p_{OD,T}r_S(0)r_{SO}(0)$.

3.3.1 OR queue length

In this section we explain how the conditional expectation of the OR queue length can be computed given the state of the system and the action chosen in the previous week. Take not_plan_t to be the number of patients that is in the OR queue at the beginning of week t, but does not want to be scheduled for the OR in week t. The number of patients that wants to be scheduled is equal to $plan_t$. The queue length at the beginning of week $t+1$ is computed as follows:

$$X_{t+1} = not\_plan_t + A_t + \max\{0, plan_t - D_t\}.$$

Note that $X_t = plan_t + not\_plan_t$. The expected OR queue length at the beginning of week $t+1$ given the state of the system at the beginning of week t and the action chosen in week t is computed as follows:

$$\begin{aligned}
&E[X_{t+1}|\sigma_t, a_t] = E[not_plan_t + A_t + \max\{0, plan_t - D_t\}|\sigma_t, a_t] \\
&= E[not_plan_t|\sigma_t, a_t] + E[A_t|\sigma_t, a_t] + E[\max\{0, plan_t - D_t\}|\sigma_t, a_t],
\end{aligned}$$

Given X_t, not_plan_t and $plan_t$ follow a binomial distribution with success probability r_t and and $1-r_t$ respectively. Which is why $E[not\_plan_t|\sigma_t, a_t] = r_tX_t$. The expression for $E[A_t|\sigma_t, a_t]$ is given in the previous section.

Index	State
1	OD
2, 3, .., 1+z_R	$R_{t,1}$, $R_{t,2}$,..., R_{t,z_R}
2+z_R, 3+z_R,... ,$1 + z_R + z_S$	$T_{t,1}$, $T_{t,2}$,..., T_{t,z_S}
$2 + z_R + z_S$	OR queue
$3 + z_R + z_S$	Home

Table 2: Indices and corresponding states

If $X_t \geq D_t$:

$$\begin{aligned} E[\max\{0, plan_t - D_t\}|\sigma_t, a_t] &= \sum_{k=0}^{X_t} \max\{0, k - D_t\} P(plan_t = k|\sigma_t, a_t) \\ &= \sum_{k=D_t}^{X_t} (k - D_t) \binom{X_t}{k} (1 - r_t)^k r_t^{X_t - k}. \end{aligned}$$

If $X_t < D_t$: $E[\max\{0, plan_t - D_t\}|\sigma_t, a_t] = 0$.

3.4 One-step transition probabilities (Method 1)

The internal transition process (w^a) of the system we consider in this thesis follows a multinomial distribution. We are interested in the one step transition probabilities: $P(\sigma_{t+1}|\sigma_t, a_t)$. For readability we often refer to the realization when we mention σ_t. See Table 2 for an an explanation for which indices belong to which state. A list of the most relevant transition probabilities:

$$\begin{aligned} p_{12} &= p_{ODR}(1 - r_R(0)) \\ p_{1,2+z_R} &= p_{OD,T}(1 - r_S(0)) \\ p_{ij} &= r_R(i) r_{RS}(i)(1 - r_S(0)) && , i = 2, 3, ..., 1 + z_R; j = 2 + z_R \\ p_{i,3+z_R+z_S} &= r_R(i)(1 - r_{RS}(i)) && , i = 2, 3, ..., 1 + z_R \\ p_{ij} &= 1 - r_R(i) && , i = 2, 3, ..., z_R; j = i + 1 \\ p_{ij} &= r_R(i) r_{RS}(i) r_S(0) r_{SO}(0) && , i = 2, 3, ..., 1 + z_R; j = 2 + z_R + z_S \\ p_{ij} &= r_S(i) r_{SO}(i) && , i = 2 + z_R, ..., 1 + z_R + z_S; j = 2 + z_R + z_S \\ p_{i,3+z_R+z_S} &= r_S(i)(1 - r_{SO}(i)) && , i = 2, 3, ..., 1 + z_R \\ p_{ij} &= 1 - r_S(i) && , i = z_R + 2, ..., z_R + z_S; j = i + 1 \end{aligned}$$

The transition probabilities of the internal transition process can be computed as follows:

$$P(w^a|\sigma_t, a_t) = \sum_{w_{ij} \in \mathcal{W}} \prod_{i=1}^{1+z_R+z_S} \binom{w_i^d}{w_{i1}, w_{i2}, ..., w_{i,3+z_R+z_S}} \prod_{j=1}^{3+z_R+z_S} p_{ij}^{w_{ij}}, \quad (1)$$

with $w_i^d = \sum_j w_{ij}$ being the number of patients that leaves state i. Note that due to the structure of our model, this is equal to the number of patients that was in state i. The set $\mathcal{W}$ consists of all

w_{ij} that satisfy:

$$\begin{aligned}
C_t =& w_{1,1} + w_{1,2} + w_{1,2+z_R} + w_{1,2+z_R+z_S} + w_{1,3+z_R+z_S} && \\
R_{t,i} =& w_{i+1,i+2} + w_{i+1,2+z_R} + w_{i+1,3+z_R+z_S} && , i \in \{1, ..., z_R - 1\} \\
R_{t,z_R} =& w_{1+z_R,2+z_R} + w_{1+z_R,3+z_R+z_S} && \\
T_{t,i} =& w_{i+1+z_R,i+2+z_R} + w_{i+1+z_R,i+2+z_R} + w_{i+1+z_R,3+z_R+z_S} && , i \in \{1, ..., z_S - 1\} \\
T_{t,1+z_R+z_S} =& w_{1+z_R+z_S,i+2+z_R} + w_{1+z_R+z_S,3+z_R+z_S} && \\
R_{t+1,1} =& w_{1,2} && \\
R_{t+1,i} =& w_{i,i+1} && , i \in \{2, ..., z_R\} \\
T_{t+1,1} =& w_{1,2+z_R} + \sum_{j=1}^{z_R} w_{j+1,2+z_R} && \\
T_{t+1,i} =& w_{i+z_R,i+1+z_R} && , i \in \{2, ..., z_S\} \\
A_t =& \sum_{j=1}^{1+z_R+z_S} w_{j,2+z_R+z_S} && \\
w_{ij} =& 0 && , \text{for all other } w_{ij}
\end{aligned}$$

We now explain how $P(X_{t+1} = i|\sigma_t, a_t)$ can be computed. Denote the realization of X_t by x_t. If $x_t \geq D_t$:

$$\begin{aligned}
P(X_{t+1} = x_{t+1}|\sigma_t, a_t) =& \sum_{n=0}^{x_t} P(X_{t+1} = x_{t+1}|\sigma_t, a_t, plan_t = n)P(plan_t = n|\sigma_t, a_t) \\
=& \sum_{n=0}^{D_t-1} P(X_{t+1} = x_{t+1}|\sigma_t, a_t, plan_t = n)P(plan_t = n|\sigma_t, a_t) \\
&+ \sum_{n=D_t}^{x_t} P(X_{t+1} = x_{t+1}|\sigma_t, a_t, plan_t = n)P(plan_t = n|\sigma_t, a_t) \\
=& \sum_{n=0}^{D_t-1} B(n, x_t, 1 - r_t)P(A_t = x_{t+1} - x_t + n|\sigma_t, a_t) \\
&+ \sum_{n=D_t}^{x_t} B(n, x_t, 1 - r_t)P(A_t = x_{t+1} - x_t + D_t|\sigma_t, a_t)
\end{aligned}$$

If $x_t < D_t$

$$\begin{aligned}
P(X_{t+1} = x_{t+1}|\sigma_t, a_t) =& \sum_{n=0}^{x_{t+1}} P(X_{t+1} = x_{t+1}|\sigma_t, a_t, not_plan_t = n)P(not_plan_t = n|\sigma_t, a_t) \\
=& \sum_{n=0}^{x_{t+1}} B(n, x_t, r_t)P(A_t = x_{t+1} - n|\sigma_t, a_t)
\end{aligned}$$

So if we want to compute $P(\sigma_{t+1}|\sigma_t, a_t)$ we have to combine the transition probabilities for the OR queue length with the transition probabilities for the internal transition process.

4 Queue length boundaries and optimal allocation of the sessions

Per surgeon there is a budgeted amount of OD and OR sessions. The allocation of these sessions throughout the year affects the OR queue length. First, in section 4.1, we explain how an SDP can be used to find the optimal allocation of the OD and the OR sessions per week throughout the year. Then, based on our performance measures, we determine boundaries for the OR queue length. This is done in section 4.2. In section 4.3 we explain how uncertainties regarding a planning horizon can be incorporated in the model. Queue length boundaries help us to stick to the performance measures of the SMK and provides us with a way to reduce the state and action space of the SDP.

4.1 Stochastic dynamic programming

We want to optimally allocate the OD and OR sessions while sticking to a budgeted amount of sessions. We have modeled the system as a Markov chain. We consider discrete time periods during a finite horizon and there are transitions that depend on the action we take. SDP is a suitable method for these types of problems. As illustrated in Figure 6, we want to go from the beginning of the year where we have not used up any of the budget (A) to the end of the year where we have used up all of the budget (B). There are several possible ways to get from point A to point B and we want to find the optimal one. With the SDP, instead of stepping forwards, we step backwards in time and at each step, for each possible state the system can be in, we find the optimal action. This action minimizes the immediate costs plus the expected future costs. The output of the SDP gives us the following information: for every week and every possible state the system can be in and every possible amount of OD and OR sessions we have left to allocate, we know what the optimal action is for that week.

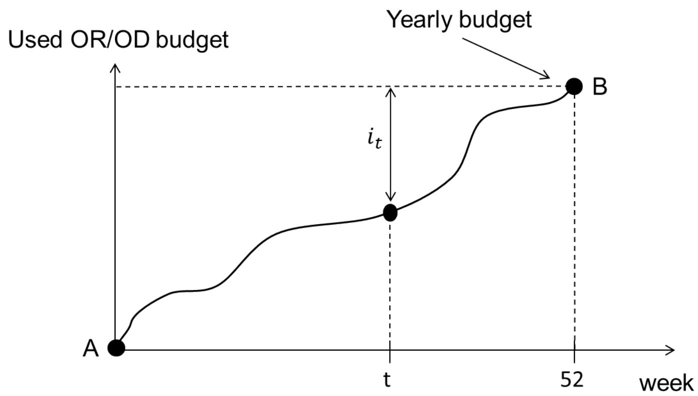

Figure 6: Used budget

The recursion:

$$f_t(\sigma_t, i_t) = \min_{a \in A_t(\sigma_t, i_t)} \{c_t(\sigma_t, a) + \sum_{j \in \mathcal{S}_{t+1}} P(j|\sigma_t, a) f_{t+1}(j, i_t - a)\}$$

$$f_{52}(\sigma_{52}, i_{52}) = c_{52}(\sigma_{52}, i_{52})$$

with budget $i_t = (i_{t,1}, i_{t,2})$

- $f_t(\sigma_t, i_t)$ is the minimal expected cost for the weeks t, $t+1$, ...,52 given that at the beginning of period t the system is in state σ_t and we have to distribute $i_{t,1}$ OD and $i_{t,2}$ OR sessions among week t until week 52. This function is also called the cost-to-go.
- $A_t(\sigma_t, i_t)$ is the set of possible actions we can take in week t given that the current budget is i_t
- $c_t(\sigma_t, a)$ is equal to the immediate costs incurred in week t if we are in state σ_t and choose action a. For our system $c_t(\sigma_t, a) = E[X_{t+1}|\sigma_t, a]$.

We want to minimize the waiting time of the patients. According to Little's law the mean waiting time of patients in the OR queue is equal to the mean OR queue length divided by the average number of patients that enters the OR queue per time unit. The average number of patients that enters the OR queue is budgeted as a result of the OD budget. So we can assume that the allocation of the OR and OD sessions does not affect the average arrival rate to the OR queue in the year of interest. This means that if we want to minimize the waiting time of the patients in the OR queue, we can minimize the OR queue length instead and take $c_t(\sigma_t, a) = E[X_{t+1}|\sigma_t, a]$.

State space week t:

$$\mathcal{S}_t = \{\sigma_t = (R_t, T_t, X_t)\colon\ 0 \le R_{t,i} \le U^R_{t,i}, 0 \le T_{t,i} \le U^S_{t,i}, s_t \le X_t \le S_t\}$$

We only include states with queue length that are within the queue length boundaries. We construct upper boundaries for diagnostics and screening by looking at the amount of workdays in the previous weeks. Take b_t to be equal to the number of days a surgeon works in week t. We use a binomial distribution with $h_{OD}\min\{3, w_{t-i}\}$ trials to compute an upper boundary for $R_{t,i}$, because $R_{t,i}$ depends on the number of patients seen at the OD in week $t-i$. The highest value with probability greater than or equal to 0.1 is set as the upper bound for $R_{t,i}$. Based on

$$C_t \xrightarrow{p_{OD,R}(1-r_R(0))} R_{t+1,1} \xrightarrow{1-r_R(1)} R_{t+2,2} \xrightarrow{(1-r_R(2)} ...,$$

the following holds:

$$U^R_{t,i} = \max_j\{B\big(j, h_{OD}\min\{3, w_{t-i}\}, p_{OD,R}(1-r_R(0))...(r_R(i-1))\big) \ge 0.1\}.$$

Similar arguments hold for $U^S_{t,i}$, but here we make use of the $U^R_{t,i}$'s as well. We construct two vectors $\mathbf{y}_1$ and $\mathbf{y}_2$. Then we the convolute these vectors and set the highest value with probability greater than or equal to 0.1 equal to $U^S_{t,i}$. The probability that j patients go from C_{t-i} to $T_{t,i}$ is equal to

$$\mathbf{y}_{1,j} = B\big(j, h_{OD}\min\{3, w_{t-i}\}, p_{OD,T}(1-r_S(0))...(r_S(i-1))\big).$$

The probability that j patients go from R_{t-i} to $T_{t,i}$ is a discrete convolution of binomial distributions with the following pairs of (number of trials, success probability): $\big(U^R_{t-i,l}, r_R(l)r_{RS}(l)(1-r_S(0))...(1-r_S(i-1))\big)$, with $l \in \{1, 2, ..., z_R\}$.

Action space: $A_t(\sigma_t, i_t)$, with $i = (i_{t,1}, i_{t,2})$

- The following must hold: # OR sessions +# OD sessions in week $t \leq b_t$
- # OD sessions in week $t \in \{0, \frac{1}{2}, 1, 1\frac{1}{2}, ..., \min\{b_t, i_{t,1}, 3\}\}$ and # OR sessions in week $t \in \{0, 1, .., \min\{b_t, i_{t,2}, 3\}\}$.
- For each t the following must hold: $\sum_{k<t} \min\{3, w_k\} \leq i_{t,1} + i_{t,2} \leq \sum_{k \geq t} \min\{3, w_k\}$. This results in a restriction on the action space and the possible values for i_t to check.

In the SDP we only include states with $s_t \leq X_t \leq S_t$, but in reality the queue length may fall outside these boundaries. For these cases we advise the following policy: if $X_t < s_t$, then $a_t = (\min\{b_t, 3\}, 0)$ and if $X_t > S_t$, then $a_t = (0, \min\{b_t, 3\})$.
Furthermore, if a valid combination of (σ_t, i_t) does not guarantee that $P(s_{t+1} \leq X_{t+1} \leq S_{t+1} | \sigma_t, a) \geq 0.8$ for any valid action a, then we delete this combination from the possible (state,budget) options we can visit in week t. This causes a reduction in possible states for week t.
Lastly, in practice it is desirable that the sum of the OD and the OR sessions in a certain week is at least equal to 2.5 and at most equal to 4. We ignore this in the thesis, because the test cases we consider are to small to add this restriction.

4.2 OR queue length boundaries

The optimal OR queue length is a trade-off between the waiting time of the patients and the amount of unused OR time. The lower the queue length, the higher the probability of too much unused OR time. Which is why we introduce a lower boundary s_t for the OR queue length in this section. Higher queue lengths result in higher waiting times. In this section we provide an upper boundary S_t for the OR queue length such that patients have the opportunity to be scheduled for the OR within W_{goal} weeks with a certain probability. These boundaries can differ per week because the number of workdays and the rescheduling probability can differ per week.

To be able to construct these boundaries, we need to have an idea of how the schedulers allocate the OD and the OR sessions. For example, suppose the schedulers allocate the sessions such that the OR queue length is always as close as possible to the lower boundary. This policy results in a higher probability of having too much unused OR time than when the sessions are allocated such that the OR queue length is close to the upper boundary. A higher lower boundary is needed for this policy. We do not know yet what the optimal allocation of the sessions is. Which is why we use an initial schedule to construct the queue length boundaries. This schedule assigns an average amount of sessions to each week relative to the amount of workdays in that week. Note that in practice it can be desirable for surgeons to have a stable schedule, i.e. doing approximately the same number of each session per week. By computing the boundaries using this initial schedule, we try to keep the optimal allocation close to that initial schedule.

Take OR_{ses} and OD_{ses} to be the budgeted amount of OD and OR sessions for the year of interest. Take b_t to be the amount of days the surgeon works in week t and take $b = \sum_{t=1}^{52} \min\{3, b_t\}$. We take the minimum of b_t and 3, because a surgeon can do at most 3 OR sessions and at most 3 OD sessions during a week. According to the initial schedule $\hat{a}$ the surgeon does $\hat{a}_{t,1} = \frac{OD_{ses} \min\{b_t,3\}}{b}$ OD sessions and $\hat{a}_{t,2} = \frac{OR_{ses} \min\{b_t,3\}}{b}$ OR sessions in week t. We can compute the number of patients seen at the OD and the OR in week t according to this initial schedule, C_t and D_t respectively. We consider patients, which is why we have to round off these fractions to integers. For the lower boundary we round up these values and for the upper boundary we round them down.

These boundaries result in a reduction in the state space and the action space of the SDP. We only allow states with an OR queue length within the boundaries. We also only allow actions in a certain week such that the probability that we are within the boundaries at the beginning of the next week is at least equal to α. So a_t must be such that:

$$P(s_{t+1} \leq X_{t+1} \leq S_{t+1} | \sigma_t, a_t) \geq \alpha.$$

4.2.1 Upper boundary

The SMK wants at least 80% of its patients to have an internal access time within the time indicated by the urgency label of the patients. The internal access time is the sum of the time spend in state Screening and the time spend in the OR queue. Because we assume a fixed surgery duration and an autonomous diagnostics and screening process, we do not distinguish between urgency labels. Inspired by this goal we construct an upper boundary for the OR queue length such that a patient has the opportunity to be scheduled for the OR within W_{goal} weeks with a probability higher than or equal to γ.
In this section we define the waiting time of a patient in the OR queue to be the time between the moment the patient enters the OR queue and the moment the patient can be planned for the OR according to the SMK. This may be an earlier date than the date on which the patient wants to be planned, e.g. think of patients who are on holiday. Let us tag the patient that arrived last to the OR queue in week $t-1$. Assume that all the patients prior to this tagged patient are allowed to reschedule, but that the tagged patient never wants to reschedule. Assume that this patient has the lowest form of priority. So he can only be scheduled for the OR if the number of patients prior to this patient that want to be scheduled for the OR in that week is less than the amount of patients that can be planned for the OR in that week. If this tagged patient can be scheduled in week $t+k$ ($k \geq 0$) then all his predecessors had the opportunity to be scheduled earlier, but some of them may have chosen not to. Recall that a patient arriving at the OR queue in week $t-1$ can receive surgery in week t the earliest. A patient arriving at the OR queue in week $t-1$ and receiving surgery in week t has a waiting time of 1 week. Take S_t such that if we have $X_t = S_t$, then the tagged patient waits at most W_{goal} weeks with a probability higher than or equal to γ. Given the initial schedule and $X_t = S_t$ we can compute the probability that the tagged patient waits at most W_{goal} weeks as follows:

$$\begin{aligned}
\gamma \leq & P(\text{waits } \leq W_{goal} \text{ weeks}) \\
= & \sum_{i=0}^{W_{goal}-1} P(\text{OR in week } t+i) \\
= & \sum_{i=0}^{W_{goal}-1} P(\text{OR in week } t+i | \text{ No OR in prior weeks}) P(\text{No OR in prior weeks})
\end{aligned}$$

If the tagged patient can enter the OR in week t, then at most $D_t - 1$ of the other patients want to be scheduled for the OR in that week. We have chosen $X_t = S_t$, so there are $S_t - 1$ patients prior to the tagged patient. We denote the probability that we have m successes out of n trials while having success probability p by $B(m, n, p)$. The number of successes has a binomial distribution.

$$B(m, n, p) = \binom{n}{m} p^m (1-p)^{n-m}$$

The tagged patient, that arrived in week $t-1$, can be scheduled for the OR in week t with probability:

$$\sum_{k=0}^{D_t-1} B(k, S_t - 1, 1 - r_t).$$

So we get the following:

$$P(\text{OR in week } t+i \mid \text{ No OR in prior weeks}) = \sum_{k=0}^{D_{t+i}-1} B(k, S_t - 1 - \sum_{j=0}^{i-1} D_{t+j}, 1 - r_{t+i})$$

$$P(\text{No OR in prior weeks}) = \prod_{j=0}^{i-1} \Big(1 - P(\text{OR in week } t+j \mid \text{ No OR in prior weeks})\Big)$$

Take the empty sum to be equal to 0 and the empty product equal to 1. Combining all this information, we can compute the probability that the tagged patient waits at most W_{goal} weeks. Take S_{t+1} to be the largest value for which this probability is larger than or equal to γ.

4.2.2 Lower boundary

The fraction of unused OR time slots in week t is equal to:

$$U_t = \frac{\max\{0, D_t - plan_t\}}{D_t}.$$

Every week we want to keep this fraction bellow a certain level called β ($0 < \beta \leq 1$)with probability larger than or equal to $1 - \epsilon$. If $D_t = 0$, then take s_t=0. Now assume that $\sigma_t = (R_t, T_t, X_t) = (\mathbf{k}_1, \mathbf{k}_2, k_3)$ and a_t is such that $D_t > 0$:

$$\begin{aligned}
\epsilon >& P(U_t \geq \beta|\sigma_t, a_t) = P(\max\{0, D_t - plan_t\} \geq \beta D_t|\sigma_t, a_t) \\
=& \sum_{k=0}^{k_3} P(\max\{0, D_t - k\} \geq \beta D_t|\sigma_t, a_t, plan_t = k) P(plan_t = k|\sigma_t, a_t) \\
=& \sum_{k=0}^{(1-\beta)D_t} B(k, k_3, 1 - r_t)
\end{aligned}$$

So take s_t to be the smallest possible k_3 for which the above equation holds. Note that $(1 - \beta)D_t$ may not be an integer, which is why we use $\lfloor(1 - \beta)D_t\rfloor$ instead.

4.3 Planning horizon

The updated production plan and minimal staffing requirements are determined two months prior to the date they are used. One plan covers a period of four consecutive weeks. After that the SMK does not want the plan to plan to be altered. So in reality, instead of deciding at the beginning of every week, like in the SDP, the amount of OD and OR sessions to do in that week, we have a planning horizon. One pro of this planning horizon is that the staff knows beforehand how many of which activity to perform in the following weeks. A downside is that we then assign a number

of sessions to a certain week based on the expected state the system is in at the beginning of that week, but the realizations may not be as expected during this planning horizon. This can be the case if more patients exit the system before entering the OR queue than expected or patients take longer to reach the OR queue than expected.
See Figure 7, suppose that in week 1 we make a schedule for week 5. Based on the expected queue length and the expected number of patients in the system is decided how many OR and OD sessions to assign to week 5. The realized OR queue length may be lower than expected, which means that we have a higher probability of too much unused OR time. If we know the probability that the OR queue length will be below the lower boundary in week 5, we can decide to increase the lower boundaries for week 5. If we do decide to increase the lower boundary for week 5, the output of the SDP will be such that we take an action in week 4 such that we are with high enough probability within the updated boundaries at the beginning of week 5. Note that the realized OR queue length can still be below the updated lower boundary, but the probability that it will be below the initial boundary is lower than it was before.

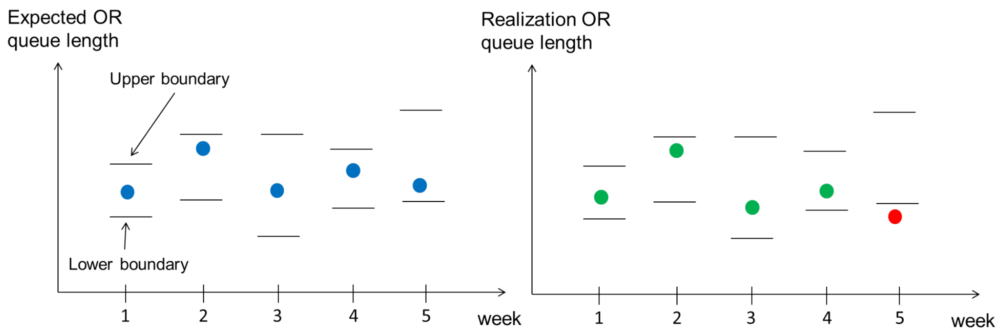

Figure 7: Planning horizon

Suppose we make a schedule in week k and we have a planning horizon of n weeks. For the case of the SMK we take $n = 12$, because the updated production plan is made 2 months (8 weeks) prior to the date it is used and one plan covers a period of 4 consecutive weeks. In this section we explain how we can compute the probability that given σ_k and the initial schedule used for the queue length boundaries, the OR queue length will drop below the lower boundaries anywhere from week $k + 1$ up to and including week $k + n$. We do not want the OR queue length to be lower than the lower boundary which is why we make all the states with a queue length lower than the lower boundary of that week absorbing states. We also do not want the queue lengths to be higher than the upper boundary, but because we are now focused on the lower boundaries, we include states with queue lengths of at most $\lceil 1.1S_t \rceil$ in the computations of this section. So we allow a 10% deviation from the upper boundary. We denote the probability that we are in an absorbing state in week $k + i$ as $P(X_{k+i} \text{ abs})$ for $i \in \{1, 2, ..., n\}$.

We do not know σ_k yet at the moment. By the time we are at the beginning of week k we do. We also do not know σ_t for $t \in \{k + 1, k + 2, ..., k + n\}$. Because we do not know the exact value of σ_k yet, we have to compute the i-step transition probability $P(X_{k+i} \text{ abs}|\sigma_k, \hat{a})$ for each possible value of σ_k. With $\hat{a}$ we denote the actions we take according to the initial schedule used for the OR queue length boundaries. Our initial idea was to use the Chapman-Kolmogorov equations to compute $P(\sigma_{k+i}|\sigma_k, \hat{a})$. These equations allow us to compute the probability that we are in state j

after n transitions given that we are now in state i. For this method we require products of transition probability matrices. These matrices can become very large very easily, specially for Model 1. In these transition probability matrices we have to include transitions from every valid σ_t to every valid σ_{t+1} with $t \in \{k, k+1, ..., k+n\}$. The number of states grows exponentially in z_R and z_S (see more in section 5). Use the programming language R for the implementation of the models. For a case the size of the SMK, we get transition probability matrices that are much larger than the maximum allowable matrix size in R. This method can be used for systems with a small state space. More details about this can be found in Appendix D. We are now going to explain a method to compute $P(X_{k+i} \text{ abs}|\sigma_k, a_k)$ that is less sensitive to z_R and z_S.

Basically, this method computes the probability that there is a certain number of patients in the OR queue in week $k+i$ given X_k and the number of patients seen at the OD prior to week $k+i$ according to the initial schedule. We now use the Chapman-Kolmogorov equations to compute $P(X_{k+i} \text{ abs}|X_k, \tilde{a})$. Here the transition probability matrices are of size $(3+z_R+z_S)\times(3+z_R+z_S)$. For readability we present the matrix P for the case in which $z_R = z_S = 2$:

$$P = \begin{array}{c} \\ C_t \\ R_{t,1} \\ R_{t,2} \\ T_{t,1} \\ T_{t,2} \\ A_{t-1} \\ \\ \end{array} \begin{array}{c} \begin{array}{cccccc} C_{t+1} & R_{t+1,1} & R_{t+1,2} & T_{t+1,1} & T_{t+1,2} & A_t \end{array} \\ \begin{pmatrix} 0 & p_{12} & 0 & p_{14} & 0 & p_{16} & 0 \\ 0 & 0 & p_{23} & \hat{p}_{2,4} & 0 & p_{26} & 0 \\ 0 & 0 & 0 & p_{34} & 0 & p_{36} & 0 \\ 0 & 0 & 0 & 0 & p_{45} & p_{46} & 0 \\ 0 & 0 & 0 & 0 & 0 & p_{56} & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 \end{pmatrix} \end{array}$$

Recall that the number of patients that is referred to the OR in week t is denoted by A_t and these patients can receive surgery in week $t+1$ the earliest. According to the Chapman-Kolmogorov equations, the probability that a patient that was at the OD in week t enters the OR queue in week $t+x$ is equal to $P^x_{1,2+z_R+z_S}$. This is equal to the element $(1, 2+z_R+z_S)$ of P^x. Then the following holds:

$$(\text{number of patients that goes from } C_{t-n} \text{ to } X_t) \sim B(P^n_{1,2+z_R+z_S}, C_{t-n}),$$

with the C_t's and D_t's following from $\hat{a}$. If we want to compute the number of patients that arrives at the OR queue in week t, we need to know how many patients were seen at the OD in the weeks t, $t-1$,...,$t-z_R-z_S$.

$$\begin{aligned} &P(A_t = i|\hat{a}) = \\ &= \sum_{n_1=0}^{i} \dots \sum_{n_{z_S+z_R}=0}^{i-n} B(n_1, C_t, \hat{P}_{1,2+z_R+z_S}) \dots B(n_{z_S+z_R}, C_{t-z_R-z_S+1}, \hat{P}^{z_R+z_S}_{1,2+z_R+z_S}) \times \\ &B(i-n-n_{z_R+z_S}, C_{t-z_R-z_S}, \hat{P}^{z_R+z_S+1}_{1,2+z_R+z_S}), \end{aligned}$$

with $n = \sum_{j=1}^{z_R+z_S-1} n_j$.

Assume that ϵ and β are such that $s_{t+1} \geq D_{t+1}$. This is a reasonable assumption because else there will be unused OR time for sure if we use the initial schedule. Take $X_k = x_k$, with $s_k \leq x \leq \lceil 1.1 S_k \rceil$.

$$\begin{aligned} &P(X_{k+1} = x_{k+1}|X_k = x_k, \hat{a}) \\ &= \sum_{j=0}^{C_t-1} B(j, x_k, 1-r_k) P(A_k = x_{k+1} - x_k + j|\hat{a}) + \sum_{j=C_t}^{x} B(j, x, 1-r_k) P(A_k = x_{k+1} - x_k + C_t|\hat{a}) \end{aligned}$$

The probability that we are in an absorbing state in week $k+1$ is equal to:

$$P(X_{k+1} \text{ abs}|\sigma_k, \hat{a}) = P(X_{k+1} < s_{k+1}|\sigma_k, \hat{a}) = P(X_{k+1} < s_{k+1}|X_k = x_k, \hat{a})$$
$$= \sum_{j=0}^{C_t-1} B(j, x_k, 1-r_k) P(A_k < s_{k+1} - x + j|\hat{a}) + \sum_{j=C_t}^{x} B(j, x_k, 1-r_k) P(A_k < s_{k+1} - x + C_t|\hat{a}).$$

The probability that we are not in an absorbing state in week $k+1$ is equal to:

$$P(X_{k+1} \text{ not abs}|\sigma_k, \hat{a}) = \sum_{i=s_{k+1}}^{\lceil 1.1 S_{k+1} \rceil \}} P(X_{k+1} = i|X_k = x_k, \hat{a}).$$

For $P(X_{k+i} \text{ abs}|\sigma_k, \hat{a})$ with $i \in \{2, 3, ..., n\}$ we use the following recursive formula:

$$P(X_{k+i} \text{ abs}|\sigma_k, \hat{a}) =$$
$$P(X_{k+i-1} \text{ abs}|\sigma_k, \hat{a}) + P(X_{k+i-1} \text{ not abs}|\sigma_k, \hat{a}) P(X_{k+i} \text{ abs}|\sigma_k, \hat{a}, X_{k+i-1} \text{ not abs})$$

Notes on the implementation of the model in R can be found in Appendix E. The method can easily be implemented in R and works much faster than using the Chapman-Kolmogorov equations to compute $P(\sigma_{k+i}|\sigma_k, \hat{a})$. If $P(X_{k+n-3} \text{ abs}|\sigma_k, \hat{a})$, ...,$P(X_{k+n} \text{ abs}|\sigma_k, \hat{a})$ are higher than desired we can choose higher lower boundaries for X_{k+n-3}, ...,X_{k+n}. This idea is a bit rough, because it does not say how high we should set the updated lower boundaries. Furthermore, it is possible that the lower boundaries for k+1,...,$k+n-1$ are already updated, so we should use the updated boundaries instead of the initial ones. Nonetheless, this method gives us a good idea of what the probability is that the queue length will drop below the lower boundary during the planning horizon, which is useful information if we want to know if and when we have to update the lower boundaries.
We still do not know how high we have to choose the lower boundaries. We do not give a clear answer to this question in this thesis. Higher lower boundaries result in higher waiting times while lower lower boundaries result in a higher probability of too much unused OR time. The rescheduling probability also affects the amount with which we have to increase the lower boundaries. The amount with which we increase the lower boundaries should also be a function of the probability that we are in an absorbing state at the beginning of the week. Furthermore, one should keep in mind that we have used the initial schedule for the computations instead of the output of the SDP. In conclusion, our advise here is to find the amounts with which we should increase the lower boundaries by trial and error and the proposed method in this section provides us with an indication of when we need to update the lower boundaries.

5 Numerical tractability

A advantage of the Markov chain described in Chapter 3 is that it is very detailed. A con is that it is very time consuming to find the optimal allocation of the sessions using this Markov chain. Due to the 'curse of dimensionality', as the size of the problem increases, solving the SDP can become numerically intractable (see also [9] and [11]). The one-step transition probabilities are very time consuming as well. Looking at equation 1, we see that we have to sum op over all feasible $w_{ij} \in \mathcal{W}$. The number of feasible actions is at most equal to 26 per week. The number of possible states for week t is equal to:

$$(S_t - s_t + 1) \prod_{i=1}^{z_R} (U_{t,i}^R + 1) \prod_{j=1}^{z_S} (U_{t,j}^S + 1).$$

This number grows exponentially in z_R and z_S. According to [3], data showed that from 2012 up to and including the first part of 2016 patients had an average diagnostics delay of more than 130 days. Data from 2016 showed that 13.61% of the patients who are referred to screening have a screening delay of 0 weeks, while 86,39% has a mean screening delay of approximately 4.57 weeks. This means that for the detailed model $z_R > 18$ and $z_S > 4$. To get an idea of how large $\mathcal{W}$ can become we present the amount of possible vectors $(w_{11}, w_{12}, ..., w_{3+z_R+z_S,3+z_R+z_S})$ for two test cases. For these test cases we have merged w_{11} and $w_{1,3+z_R+z_S}$, this causes a small reduction in the size of $\mathcal{W}$. Assuming that all the probabilities in Figure 3 are nonzero, the number of possible realizations at the beginning of week $t + 1$ is equal to:

- 825 if $C_t = 50$, $z_R = z_S = 1$, $A_t = 8$ and $\sigma_t = \sigma_{t+1} = c(20, 20, 16)$
- 1811 if $C_t = 50$, $z_R = 2$, $z_S = 1$, $A_t = 8$ and $\sigma_t = \sigma_{t+1} = c(20, 18, 20, 16)$.

The size of the state space and the structure of the transition probabilities make this method computationally expensive. In the following sections we describe our attempts in trying to make the SDP numerically tractable and speed up the computation time. First, in Section 5.1, we construct a model that is less detailed and has a smaller state space. In section 5.2 we propose approximate transition probabilities that are less time consuming. We also explain how the probability of reaching an invalid state can be computed. Then we explore the possibility of speeding up the computation time of the transition probabilities even more by using numerical approximations. This is done in Section 5.3. In the ideal case, the output of the SDP results in a practically applicable policy that can be used by the SMK. In Section 5.4 we explain whether this is possible or not.

Lastly, we want to mention that we have also considered using time aggregation to speed up the computations. Recall the time aggregation method proposed by Cao et al [4]. This method requires the state space to be divided in two complementary subsets $\mathcal{S}_1$ and $\mathcal{S}_2$. These subsets must be such that actions are only taken when the system is in a state that is in $\mathcal{S}_1$ and the transition probabilities for states in $\mathcal{S}_2$ do not depend on actions. We did not manage to construct such sets. An alternative is to not look at 52 separate weeks when we want to compute the optimal action for week 1, but aggregate future time periods. For example, consider the first 4 weeks separately and divide the remaining 48 weeks in 4 blocks of 12 weeks. The transition probabilities from one block to another can be computed by using the Chapman-Kolmogorov equations to compute the 12-steps transition probabilities $P(\sigma_{t+11}|\sigma_t, a_t)$. However, as explained in Appendix D, this can be computationally expensive due to the size of the transition matrix.

5.1 Alternative model (Model 2)

In this section we propose a model (Model 2) that causes a significant reduction in the state space compared to our previously mentioned, detailed model (Model 2). For this model we only keep track of the amount of patients that is in diagnostics and screening. We do not keep track of how long patients have already been in either of these states. Figure 8 gives a schematic depiction of the new system. Note that this is not the same as taking $z_R = z_S = 1$ in Model 1, because then we have a system in which patients are at most one week in diagnostics and one week in screening. In Model 2, every week a patient can either leave diagnostics or stay there for at least one more week. So, theoretically speaking, a patient can stay an infinite amount of time in diagnostics and in screening. The transition labeled as [1] is for patients who switch between surgeons and patients who finish diagnostics and screening fast enough to go from the OD in week t to the OR queue in week $t+1$. Transition [2] is for patients who finish screening fast enough to go from diagnostics in week t to the OR queue in week $t+1$. If we would not include the transitions [1] and [2] then patients always have a delay of at least 3 weeks between the OD and the OR.

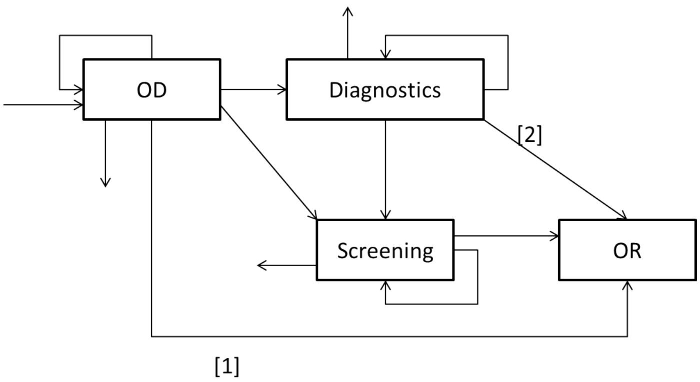

Figure 8: Simplified system

The internal transition process of Model 2 follows a multinomial distribution as well:

$$P(w^a|a_t) = \sum_{w_{ij} \in \mathcal{W}} \prod_{i=1}^{3} \binom{w_i^d}{w_{i1}, w_{i2}, w_{i3}, w_{14}, w_{15}} \prod_{j=1}^{5} p_{ij}^{w_{ij}}, \tag{2}$$

with indices: 1=OD, 2=diagnostics, 3=screening, 4=OR, 5=home. The set $\mathcal{W}$ constist of all w_{ij} that satisfy:

- $w_1^d = C_t$, $w_2^d = R_t$, $w_3^d = T_t$, $w_4^d = 0$ and $w_5^d = 0$. Recall that $w_i^d = \sum_{j=1}^{5} = w_{ij}$.
- $w_{12} + w_{22} = R_{t+1}$
- $w_{14} + w_{24} + w_{34} = A_t$
- $w_{13} + w_{23} + w_{33} = T_{t+1}$
- $w_{ij} = 0$ for all the w_{ij} that are not mentioned

The computations for $P(X_{t+1} = i|\sigma_t, a_t)$ remain the same as in Model 1. To be able to compute this transition probability, we need to compute A_t several times. Which implies that we have to compute equation 2 several times for several values of A_t and this is time consuming.

The structure of the SDP remains the same. The only difference is the computation of the upper boundaries for diagnostics and screening, U_t^R and U_t^S respectively. By repeatedly inserting the expression for $E[R_k]$ in the expression for $E[R_{k+1}]$ we get the following:

$$\begin{aligned}
E[R_{t+1}] &= E[E[R_{t+1}|C_t, R_t]] \\
&= E[E[p_{12}C_t + p_{22}R_t|C_t, R_t]] \\
&= p_{12}E[C_t] + p_{22}E[R_t] \\
&= p_{12}E[C_t] + p_{22}(p_{12}C_{t-1} + p_{22}R_{t-1}) \\
&\dots \\
&= p_{12}\sum_{i=0}^{n} p_{22}^i E[C_{t-i}] + p_{22}^{n+1}E[R_{t-n}].
\end{aligned}$$

The we compute $m = \text{argmax}_i\{p_{12}p_{22}^i \geq \zeta_R\}$ to get:

$$U_t^R = p_{12}\sum_{i=0}^{m} p_{22}^i h_{OD} \min\{3, w_{t-i}\}.$$

Depending on how many values we want to include, we choose ζ_R and ζ_S. Lower values of ζ_R and ζ_S give higher upper boundaries. If we use similar reasoning for U_t^S and repeatedly insert the expression for $E[T_k]$ in the expression for $E[T_{k+1}]$ we get the following:

$$E[T_{t+1}] = p_{13}\sum_{i=0}^{n} p_{33}^i E[C_{t-i}] + p_{23}\sum_{i=0}^{n} p_{33}^i E[R_{t-i}] + p_{33}^{n+1}E[T_{t-n}]$$

Take $m = \text{argmax}_i\{\max\{p_{13}, p_{23}\}p_{33}^i \geq \zeta_S\}$ to get:

$$U_t^S = p_{13}\sum_{i=0}^{m} p_{33}^i h_{OD} \min\{3, w_{t-i}\} + p_{23}\sum_{i=0}^{m} p_{33}^i U_{t-i}^R.$$

Note that we can use the same computations for the queue length boundaries and planning horizon for Model 2 as described in Model 1.

5.1.1 Computation time test cases Model 2 Method 1

As expected, the computation of the transition probabilities can be time consuming. The transition rate matrix and the values of ζ_R and ζ_S strongly affect the upper boundaries for screening and diagnostics, which affects the number of possible states and the size of $\mathcal{W}$.

In this section we use the transition probabilities computed in Section 6. These probabilities are partly based on data and partly are an educated guess. To speed up the computation time we merge w_{11} and w_{15} and take $w_1 = w_{11} + w_{15}$ and $p_1 = p_{11} + p_{15}$. This results in one for-loop less in the implementation of the method. Let us first look at a small test case

OD_{ses}	2	w_{goal}	3
OR_{ses}	5	h_{OD}	2
r_t ($\forall t$)	(0.1,0.2,0.1)	h_{OR}	2
b_t ($\forall t$)	(3,4,3)	Horizon	3 weeks
ζ_R	0.02	ζ_S	0.05

Table 3: Parameters test case 1

The SDP takes approximately 12 seconds to solve. The boundaries are: $s_1 = s_3 = 5$, $S_1 = S_3 = 9$, $s_2 = 6$, $S_2 = 9$ and $U_t^R = U_t^S = 1$ for $t \in \{1, 2\}$. We have taken $\zeta_R = 0.02$, because $p_{12}p_{22} \approx 0.034$. We then randomly chose a value that is a bit smaller than 0.034. A similar argument holds for ζ_s. Of course, there are also other ways one could set ζ_S and ζ_R.

As the previous SDP took just a couple of seconds, let us look at a bigger system:

OD_{ses}	4	w_{goal}	4
OR_{ses}	8	h_{OD}	20
r_t	(0.1,0.1,0.2,0.2)	h_{OR}	2
b_t	(2,3,4,5)	Horizon	4 weeks
ζ_R	0.02	ζ_S	0.05

Table 4: Parameters test case2

After approximately 9.5 hours of running the code, the SDP had found the cost-to-go for week 4 and has computed for 13 out of the 1260 possible states what the optimal action is in week 3. Note that computing the optimal action for one state in this example takes more time than computing the optimal action for one state in the previous example. This is due to the difference in size of $\mathcal{W}$. The boundaries are:

week	1	2	3	4
U_t^R	18	17	17	17
U_t^S	10	10	9	9
s_t	4	7	8	8
S_t	14	14	14	14

Table 5: Boundaries test case2

In conclusion, we have converted our detailed model into one in which the state space is described by only three properties. Due to the time consuming transition probabilities this model is not applicable in practice yet. In the next section we explain an alternative method that makes use of approximations of the transition probabilities.

5.2 One-step transition probabilities (Method 2)

The previously mentioned transition probabilities (Method 1) are time consuming. This is due to the fact that the number of patients that goes from state i to state j and the number of patients that goes from state i to state k depend on each other. The total number of patients that leaves state i

and goes to state j or state k has to be at most equal to the number of patients that was in state i. In this section we propose approximations to these probabilities. Here we assume that the number of patients that goes from state i to state k is independent of the number that goes from state j to state k. The number of patients that goes from state i to state j follows a binomial distribution with success probability p_{ij} and the number of trials being equal to the number of patients that was in state i. As a result of this we can use discrete convolutions of binomial distributions to compute the transition probabilities. These probabilities can be computed much faster than the probabilities in Method 1 (see Section 5.1.1). A downside to this method is that it is possible to enter an invalid state in which more people go from state i to either state j and state k than were in state i. An example of such a relationship is the dependency of $R_{t,1}$ and $T_{t,1}$ on C_{t-1} (see Section 3.4). A detailed description of how these transition probabilities are computed can be found in Appendix A.

Let us compare the computation time of Method 1 and Method 2 for Model 2. Recall the second, larger test case mentioned in Section 5.1.1. There, after more than 9.5 hours of computation time, the SDP only got to 13 out of the 1260 possible states of week 3. The SDP using the transition probabilities of Method 2 took approximately one hour to solve. Which means that Method 2 is much faster than Method 1.

However, when we let the SDP solve a test case using Model 2 and Method 2 with parameters of which only the horizon is unrealistic, we get high computation times. This test case considers only 5 weeks. The SDP takes approximately 29 minutes to compute the cost-to-go for week 5. After 10 hours the SDP had only computed the optimal action for 7% of the possible combinations of (σ_t, i_t) of week 4.

In conclusion, Model 2 and Method 2 result in a significant reduction of the computation time. However the computation time is still too high.

5.2.1 Probability of reaching an invalid state with Method 2

In this section we explain how, for Model 2, the probability can be computed that we reach an invalid state with Method 2. A similar idea holds if we combine Model 1 and Method 2. Suppose we are in state i, $i \in \{1,2,3,4,5\}$. Then compute $\mathbf{x}_{ij}$ for each $j \in \{1,2,3,4,5\}$, with $\mathbf{x}_{ij}(k)$ equal to the probability that k ($k \in \{0,1,...,I\}$) patients go from state i to state j.

The probability that more patients than C_t patients leave the OD and go to diagnostics, screening and the OR is equal to:

$$p_{OD} = \sum_{k=C_t+1}^{3C_t} \sum_{k_2=0}^{k} \sum_{k_3}^{k-k_2} x_{12}(k_2)x_{13}(k_3)x_{14}(k-k_2-k_3)$$

In R this can be easily computed using the built in convolution function. Note that we do not include k_{11} and k_{15} because these patient streams are less relevant for our model. We enter an invalid state if the sum of the patients who goes from the OD to screening, diagnostics or the OR queue is more than the number of patients that was originally at the OD. If this sum is less than C_t then the remaining patients stay at the OD or exit the system and go home.

The probability that more patients than R_t patients go from diagnostics in week t to diagnostics, screening and the OR queue in week $t+1$ is equal to:

$$p_R = \sum_{k=R_t+1}^{3R_t} \sum_{k_2=0}^{k} \sum_{k_3}^{k-k_2} x_{22}(k_2)x_{23}(k_3)x_{24}(k-k_2-k_3)$$

The probability that more patients than T_t patients go from screening in week t to screening and the OR queue in week $t+1$ is equal to:

$$p_T = \sum_{k=T_t+1}^{2T_t} \sum_{k_3=0}^{k} x_{33}(k_3)x_{34}(k-k_2)$$

Combining these probabilities, we get the probability that we include an invalid state in week $t+1$: $1-(1-p_{OD})(1-p_R)(1-p_T)$. The same method can be applied to compute the probability that we reach an invalid state in Model 1 if we use Method 2.

Example with realistic values
Suppose that 46 patients are seen at the OD during week t, which is a realistic number if 1 OD session is planned. Take $R_t = 10$ and $T_t = 10$ and use the transition probabilities of Table 6. Then $1-(1-p_{OD})(1-p_R)(1-p_T) \approx 0.52$, with $p_{OD} \approx 0$, $p_R \approx 0.22$ and $p_T \approx 0.38$. Even though the separate probabilities are acceptably low, the probability that the SDP will reach an invalid state is pretty high. Nonetheless, we may learn something from this method, as Method 1 only allows us to compute very small test cases.

5.3 Numerical approximation of the transition probabilities (Method 1 and Method 2)

In this section we explore several possibilities for approximations of the transition probabilities of Method 1 and Method 2. First the approximations for Method 2 are explained and then we explain our attempts at approximating the transition probabilities for Method 1.

Method 2
For Method 2 we use normal approximations to the binomial distribution: $B(n,p) \sim N(np, np(1-p)$. If n is large and p not near 0 or 1 we get good approximations. According to [13], the normal approximation will be quite good if n is such that $np(1-p) \geq 10$. If n is large and p is small the binomial distribution can be approximated by a Poisson random variable with parameter $\lambda = np$ as well. This provides us with an approximation of the probability that m $(0 \leq m \leq n)$ patients out of n patients move from state i to state j with $p_{ij} = p$. We are more interested in directly approximating $P(\sigma_t|\sigma_t, a_t)$ instead of approximating each of the transitions separately. In Appendix B we describe a multivariate normal approximation for the transition probability for diagnostics and screening together. The transition probabilities for the OR queue length can be approximated using normal distributions. A nice property about the multivariate normal distribution is that we only have to compute the mean and the variance to be able to compute the transition probabilities for screening and diagnostics.

Even though we get good approximations, they do not speed up the computation time much. This holds for both models. For Model 1, this is due to the restrictions set on the values of n for which the approximation holds. For Model 1 if C_t or any of the $R_{t,i}$'s or $T_{t,i}$'s does not contain a lot of patients or even zero patients, which is very likely to happen, then the approximation will not be very good. In that case, we have to use Method 2, which means that we often have to use Method 2 instead the approximation. If we would allow any of the $R_{t,i}$'s or $T_{t,i}$'s to be equal to 0 in Method 1, then the covariance matrix of the multivariate normal distribution is singular.
The transition probabilities are most often used when computing the expected future costs. There we loop over all the possible future states, compute the probability of ending up in that state

at the beginning of the next week and multiply it by its cost-to-go. We do not include future states that we are not likely to reach from the current state. For example, consider the possible future state $\sigma_{t+1} = (k_1, k_2, k_3)$. We first compute the probability that $R_{t+1} = k_1$, given σ_t. If this probability is lower than 0.1, we do not include this future state in the current cost-to-go. If the probability is at least equal to 0.1, we compute the probability that $T_{t+1} = k_2$, and so on. When using the numerical approximations of these probabilities, we first exclude states for which the combination $(R_{t+1}, T_{t+}) = (k_1, k_2)$ occurs with probability less than 0.01. Note that more future states are included in the latter. Which is probably why we do not see a significant reduction in the computation time when using the numerical approximations.

Method 1
We have also tried to approximate the transition probabilities of Method 1. We are interested in a good approximation of the multinomial distribution. For this we have constructed a multivariate normal approximation. If $Y \sim Multi(n, \mathbf{p})$, with $\mathbf{p} = (p_1, ...p_k)$, then we can approximate Y if n is large and $\mathbf{p}$ is not near the boundary of the following set:

$$\left\{ \mathbf{p} \in \mathbb{R}^k : 0 \leq p_i, i = 1, ..., k \text{ and } \sum_{i=1}^{k} p_i = 1 \right\}$$

as follows: $Multi(n\mathbf{p}, nM) \approx \mathcal{N}(n\mathbf{p}, nM)$, with $M = diag(\mathbf{p}) - \mathbf{p}^T\mathbf{p}$ ([6],[12]. Note that the matrix M is singular [12], because $\sum_i p_i = 1$. Which means that the regular density function does not exist [12]. We can use the following theorem from [12] to compute the probability density function of Y:

Theorem: If $Y \sim \mathcal{N}_m(\boldsymbol{\mu}, \Sigma)$ with $\text{rank}(\Sigma) = k \leq m$, then Y has probability density function

$$f(\mathbf{x}) = (2\pi)^{-\frac{1}{2}k}[\det_k\Sigma]^{-\frac{1}{2}}\exp\bigg(-\frac{1}{2}(\mathbf{x} - \boldsymbol{\mu})^T\Sigma^+(\mathbf{x} - \boldsymbol{\mu})\bigg)$$

for $\mathbf{x} \in V = \boldsymbol{\mu} + \mathcal{K}(\Sigma)^\perp$ with respect to the Lebesgue measure λ_V on V, where $\det_k\Sigma$ denotes the product of positive eigenvalues of Σ.

Here Σ^+ is the generalized inverse of Σ and $\mathcal{K}(\Sigma)^\perp$ is the column space of Σ. Generalized inverses are not unique. In R the Moore-Penrose generalized inverse is used. This inverse is also used in [12].

Surprisingly, the approximations are not good. Also this approximation is slower than the multinomial distribution. This is because we cannot approximate $P(\sigma_t|\sigma_t, a_t)$ as a whole by one distribution, we approximate each multinomial term separately. For each individual term we have to compute the generalized inverse, the eigenvalues and the value of the density function. We did consider looking at linear combinations of the multivariate normal approximations instead. According to [18], linear combinations of the elements of a multivariate normal distribution, Y_1,...,Y_m, follow a normal distribution.

$$\mathbf{aY} = \sum_{i=1}^{m} a_i Y_i \sim N(\mathbf{a}\boldsymbol{\mu}, \mathbf{a}\Sigma\mathbf{a}^T)$$

This idea does not work for our method, because for each linear combination we constructed, we did not get the desired dependency between the realizations w_{ij}.

Let us look at an example to get an idea of the difference between the multivariate normal approximation and the multinomial distribution. Take $n = 60$ and $\mathbf{p} = (0.1, 0.2, 0.7)$. We compute the following probability: $P(Y = (6, 12, 42))$, note that $\boldsymbol{\mu} = n\mathbf{p} = (6, 12, 42)$. Which means that the exponential term in the density function will be equal to 0. Using $Multi(n, \mathbf{p})$, we get $P(Y = (6, 12, 42)) = 0.022$ and using the density function for singular covariance matrices we get: $P(Y = (6, 12, 42)) = 0.013$, which is almost half of the probability we get with the multinomial distribution. We have used the mean vector $\boldsymbol{\mu}$ and the following covariance matrix that has rank 2:

$$\Sigma = nM = \begin{pmatrix} 5.4 & -1.2 & -4.2 \\ -1.2 & 9.6 & -8.4 \\ -4.2 & -8.4 & 12.6 \end{pmatrix}.$$

Even if we increase n, the approximations do not give values closer to the true value.

5.4 Monotone policy

In this section we explore the possibility of the system having a monotone policy. Recall that the output of the SDP gives the following information: for every week, every possible state and possible amount of budget left to distribute, we know the optimal action to take in that week. Even for small test cases this results in large tables with a lot of information. If this would be given to the hospitals, people will probably not use it or have no idea what to do with it. It would be better to give a clear policy to the hospital. For example, take action x_1 if the number of patients in the system is at most y_1, take action x_2 if there are more than y_1 and at most y_2 patients in the system, etc. A policy similar to the (s, S) inventory policy can easily be implemented in practice as well.
With this in mind we first tried an analytical approach, but we could not find a monotone property by looking at the equations of the models. Looking at the result of test cases, we do not see a monotone property of the kind that we discussed in the previous paragraph. We have only looked at results from Method 2, as Method 1 only allows us to compute very small test cases. It is not the case that the more patients are in the system, the more OR sessions and the fewer OD sessions have to be planned. This is also the case if we look at the number of patients in the OR queue plus the number of patients we expect to enter the OR queue from diagnostics or screening at the next time step.
What we do see for Model 1 is that for each week, states with the same set of feasible actions tend to have the same optimal amount of OR sessions. Recall that we only are only allowed to take action a in week t when in state σ_t if $P(X_{t+1}|\sigma_t, a)$.

Model 1
Looking at the results of several test cases, we see that states with the same set of valid actions tend to have the same optimal amount OR sessions. We have looked at test cases with $z_R, z_S \leq 2$. Comparing the results of the Model 1 Method 2 with the results where we assume this monotone OR policy, we see that on average the monotone assumption gives good approximations. However there are some outliers. We have tested the accuracy of the approximations by comparing the cost-to-go of both methods. The relative absolute error of the cost-to go's is considered: $|f - \tilde{f}|/f$, with $\tilde{f}$ and f being the cost-to-go resulting from the monotone policy and Model 1 respectively. The average relative error is generally below 0.07. We do have a test case with an absolute average relative error of at most 0.0098, while the maximum relative error is equal to 2.4. The approximation gives a cost-to-go in the first week for a certain state that is 2.4 times the cost-to-go of the results of Model 1 Method 2.
For the implementation of this policy we have ordered the states. First we loop over all possible values of $R_{t,1}$, then over $R_{t,2}$ and lastly over all possible values of X_t. Then for each state we compare its valid set of actions with the one of its predecessor. If they have the same set of valid

actions, choose the optimal amount of OR sessions of the predecessor as the optimal amount of OR sessions for this state. For the test cases, this policy causes approximately a 20% reduction in the computation time.

Model 2
We have looked at the output of three test cases. Two of these test cases have the exact same parameters except for the transition probabilities. The first one has the same transition probabilities as in Section 6 and the second one has different transition probabilities out of the state diagnostics. The third test case considers a system with less patients and one workday less. All the three test cases consider a horizon of 4 weeks. Looking at samples from the output of these three test cases, we see that states with the same set of feasible actions have the same optimal amount of OR sessions. For the first and second test case, these states also have the same optimal amount of OD sessions. The policy for these two test cases is as follows: if the states have the same set of valid actions, then take the action with the least amount of OD sessions and out of the remaining possible actions take the one with the highest amount of OR sessions. This policy does not hold for the third test case. The parameters for these test cases can be found in Appendix F. Due to time constraints we did not manage to extensively check this monotone OR policy for Model 2.

In conclusion, we did not manage to analytically come up with a monotone policy or groups of states that have the same optimal action. Based on the output of several test cases using Model 1, we observe that, per week, states with the same set of feasible actions tend to have the same amount of optimal OR sessions. More testing is needed for larger test cases. The average relative error appears to be acceptably low, while the largest relative error for the test cases is equal to 2.4. For Model 2, we see that, per week, states with the same set of feasible actions have the exact same amount of optimal OR sessions. However this observation is based on three test cases and needs to be tested more extensively.

5.5 Output test case Model 2 Method 2

In this section we give three possible realizations of the allocation of the OD and the OR sessions. Model 2 and Method 2 were used to determine these optimal allocations. We have used the transition probabilities described in Section 6 and we have used the the same parameters as the second, larger test case mentioned in section 5.1.1. The actual output is very detailed and tells us per week and for each valid combination of (σ_t, i_t) what the optimal action is. The amount of valid combinations of (σ_t, i_t) is equal to 19404 in week 4, 11074 in week 3, 1210 in week 2 and 746 in week 1.
Instead of presenting all this information we consider realizations with three different starting points in week 1. Case 1 considers the state with the amount of patients in the system: $\sigma_1 = (0, 1, 11)$. Case 3 considers the state with the largest amount of patients in the system: $\sigma_1 = (18, 9, 13)$ and Case 2 is randomly chosen: $\sigma_1 = (17, 7, 10)$. Figure 9, shows the used amount of OD sessions and Figure 10 shows the used amount of OR sessions. We see that Case 1 does more OD sessions in the first period and less OR session compared with the other two cases. Case 3 does the opposite and Case 2 is roughly in between the two. Note that $\sigma = (0, 0, 4)$ lies within the boundaries of the state space of the SDP and if the system is in this state, there are less present compared with Case 1. This state is eventually excluded. Probably because it is not possible to stay within the queue length boundaries with high enough probability if the system is in this state at the beginning of week 1. A similar argument holds for $\sigma_t = (18, 10, 14)$.

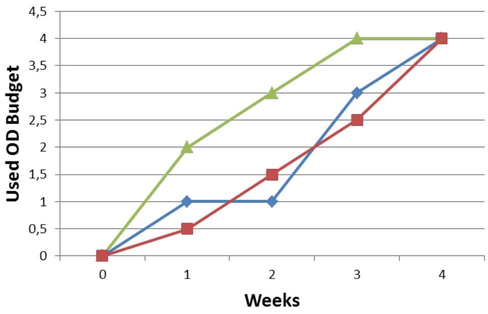

Figure 9: OD sessions

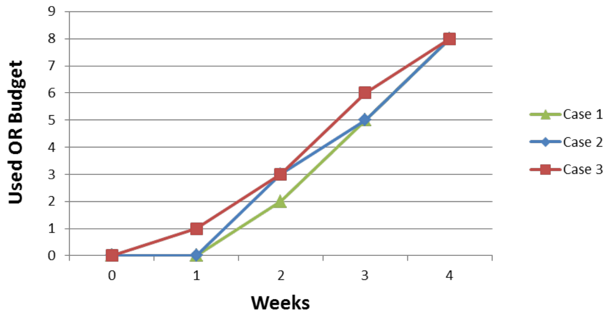

Figure 10: OR sessions

6 From data to parameters

In this section we explain what kind of data we need and how we can convert it into parameters. Unfortunately, due to time constraints we did not manage to get all the parameters from the actual data. Furthermore, we give an educated guess for the parameters for Model 2 that we do not know.

- Known per surgeon: h_{OR}, h_{OD}, OR_{ses}, OD_{ses}, d_t $(\forall t)$
- Unknown: r_t $(\forall t)$, p_{ij} $(\forall i, j)$

We need data from all the patients who had consultations with the surgeon of interest. This includes data from patients who have switched between surgeons and have had a consultation with this surgeon. From a patient's data we need to know:

- which consultations were prior to the start of the diagnostic process
- when the patient was referred to diagnostics
- when the patient had received an OR-ticket
- when the patient's screening process was finished
- the OR date
- when the patient's treatment is finished
- which OR time slots were idle prior to the patients OR date, this is needed for the rescheduling probability

If the patient receives more than one surgery during his treatment, we need all of these dates after the primary surgery as well.

Kapadia et al [7] provide us with a method to estimate the transition probabilities based on the data. The authors label patients based on the state in which the patients enter the system. We assume that patients only enter the system via the OD. The probability that a patient that is in state i goes to state j is equal to:

$$p_{ij} = \frac{\text{Number of transitions between states } i \text{ and } j}{\text{Total number of transitions out of state } i}$$

We have combined expert opinions and performance measures based on data collected from all the orthopedic surgeons together to construct an educated guess for the parameters we do not know yet. According to [3], 36% of the patients who arrives at the SMK for an initial consultation eventually gets an OR ticket. The other 64% will not receive surgery. Based on an expert opinion, 98% of the patients who receives an OR-ticket gets a screening approval. The other 2% will not receive surgery. Out of all the patients who are referred to screening, 13,61% has a screening delay of 0 weeks. The other 86,39% has an average screening delay of approximately 4,57 weeks. This is based on data from 2016. Furthermore 70% of the patients have a diagnostics delay of 0 weeks and the other 30% has an average diagnostics delay of approximately 19 weeks. These percentages are based on an expert's opinion. Based on the data we get the following one-step transition probabilities:

	OD	Diagnostics/Screening	OR	Home
OD	0.4397	0.1207	0.0208	0.4188

For this table we used the fact that the treatment of 53% of the patients who have an NC will be ended directly after this consult. We also know that 36% of the new patients will eventually get an OR ticket. So out of all the patients who enter the system 47% will be referred to the radiology department and 77% of these patients will receive an OR ticket. Combining all of the above, we get the following one-step transition probabilities:

	OD	Diagnostics	Screening	OR	Home
OD	0.4397	0.0362	0.0730	0.0323	0.4188
Diagnostics	0	0.9474	0.0348	0.0055	0.0123
Screening	0	0	0.7807	0.2149	0.0044

Table 6: Transition probabilities Model 2

The budgeted amount of OD and OR sessions is partly based on the OD hitrate of the surgeon. For example, if a surgeon has an OD hitrate of 7 and if a surgeon sees 70 patients at the OD during a certain week, then we expect that out of these 70 consultations 10 patients will enter the OR queue eventually. The hitrate of the surgeons varies between 5.3 and 15.3. The average OD hitrate is equal to 8.14. With the probabilities mentioned in Table 6 we get a slightly lower hitrate of approximately 7.6.

7 Conclusion and discussion

In this thesis we aimed at determining the optimal allocation of the outpatient department (OD) and operating room (OR) sessions per orthopedic surgeon, per year. The system is modeled as a Markov chain. This Markov chain gives a very detailed description of the system (Model 1) and its internal transition process follows a multinomial distribution (Method 1). In this thesis we have considered two main performance measures:

- Minimize the waiting time of the patients in the OR queue
- Keep the amount of unused OR time below a certain level.

The latter performance measure is used to construct a lower boundary for the OR queue length. The upper boundary guarantees that a patient has the opportunity to be planned for the OR within a certain week with a high enough probability. These boundaries result in a restriction in the state space and the action space. The values of these boundaries depend on several factors, namely the number of workdays of a surgeon per week, the maximum waiting time goal and the level of unused OR time we do not want to exceed. Several methods are proposed to compute the probability that the OR queue length drops below the lower boundary during this planning horizon. However, only the one mentioned in Section 4.3 can be used for a case the size of the SMK.
Out initial, detailed model is computationally expensive. In Section 5.1 we have proposed an alternative model (Model 2) that causes a huge reduction in the size of the state space. More research can be conducted to determine how well this alternative method mimics reality. This reduction in the size of the state space is not enough, because it takes a lot of time to compute the transition probabilities. For this we have proposed an alternative method in Section 5.2, named Method 2. The probability of reaching an invalid state with this method can become quite high, but we have used it anyway as Method 1 only allows us to compute very small test cases and maybe we can learn something from the results. Combining Model 2 and Method 2 we see a significant reduction in the computation time compared to Model 1 and Method 1. However, the computation time remains too high. In Section 5.3 we propose numerical approximations for Method 1 and Method 2. Surprisingly, we get bad approximations for Method 1, where we approximate the multinomial distribution by the multivariate normal distribution. For Method 2 we get good approximations by approximating the binomial distribution by a normal distribution. For both the models, these approximations do not result in a significant reduction of the computation time.
In the ideal case, we present the results of the SDP as a clear, practically applicable policy to the SMK. However, due to the complex structure of our system, we did not manage to find such a policy. Based on samples of the results of test cases using the transition probabilities of Method 2, we see for Method 1 that per week, states with the same set of valid actions have the same optimal amount of OR sessions. For Method 1 this is not always the case, but we see that these states do tend to have the same optimal amount of OR sessions. For Method 1 this policy gives good results on average, but there are some outliers. Keep in mind that these observations were based on the output of an SDP that approximates the transition probabilities. More testing is needed to be able to properly assess the accuracy of the policy. So we have learned that when using the approximate transition probabilities we can, per week, find groups of states that tend to have the same optimal amount of OD sessions. However this is not the practically applicable policy we were looking for.
Furthermore we would like to use parameters from real data. We did derive an educated guess based on data and expert opinions in Section 6. However these numbers are based on the performance of all the surgeons together. We aim to find the optimal allocation of the sessions per surgeon, so we prefer deriving parameters from the complete data per surgeon. Ideally, we also want to compare

our results with historical schedules to see if the allocation we come up with is very different and if it results in an improvement.

It can be argued that ignoring the state diagnostics results in an accurate model. Ignoring this state results in a reduction of the state space. After the patient's last visit to the radiology department, he is either asked to come back to the SMK another day for a follow-up consultation or he has a HNR appointment. This appointment is on the same day as his last visit to the radiology department. The appointment takes just a couple of minutes and during this appointment the patient receives the OR-ticket. We have chosen to ignore these HNR consultations in this thesis as the timing of the HNR and VC consultations depend on the duration of the diagnostic process. If one chooses to include these appointments, we can argue that we can ignore the state diagnostics as the patient always has a consultation at the OD before entering screening. Further research can be done in exploring the accuracy of this model.

For future research, an interesting approach would be to use approximate dynamic programming (ADP). ADPs are well known to overcome the 'curse of dimensionality'. However, finding the right settings of the ADP that make it a good approximation can be difficult. Due to time constraints we did not manage to use this method as well. The models proposed in this report can be used to assess the accuracy of the ADP. The paper by Hulshof et al [9] will be worth reading, as the authors consider a similar system and use an ADP to compute their optimal actions.
Another idea is to then do one iteration of policy iteration. Policy iteration can be computationally expensive, but if we use a good policy as input, we may get a better policy after doing this one iteration.

Bibliography

[1] A. Akkermans. *Constructing and evaluating a yearly capacity plan for the orthopedic department of the Sint Maartenskliniek.* Internship report, University of Twente, 2017.

[2] S Bashyam, MC Fu *Optimization of (s,S) inventory systems with random lead times and a service level constraint.* Management Science 44 (12 part 2), 243-256, 1998.

[3] K Bos, W Hasper. *Enabling measuring of the patient flow in an orthopedic clinic.* Bachelor Thesis, University of Twente, 2016.

[4] XR Cao, Z Ren, S Bhatnagar, M Fu, S Marcus. *A time aggregation approach to Markov decision processes.* Automatica 38, 929-943, 2002

[5] A Fügener, EW Hans, R Kolisch, N Kortbeek, PT Vanberkel. *Master surgery scheduling with consideration of multiple downstream units.* European Journal of Operational Research 239, 227-236, 2014.

[6] CJ Geyer. *Brand name distributions.* University of Minnesota, 2012, `http://www.stat.umn.edu/geyer/5102/notes/brand.pdf.`

[7] A Kapadia, W Chan, R Sachdeva. *Predicting duration of stay in a pediatric intensicve care unit: a Markovian approach.* European Journal of Operational Research 124 (2),353-359,2000.

[8] A Kapadia, S Vineberg, D Rossi. *Predicting course of treatment in a rehabilitation hospital: a Markovian model.* Computers & Operations Research 12, 459-469,1985

[9] PJH Hulshof, MRK Mes, RJ Boucherie, EW Hans. *Patient admission planning using Approximate Dynamic Programming.* Flexible Services and Manufacturing Journal 28 (1-2), 30-61, 2016.

[10] LGN Nunes, SV de Carvalho, R de CM Rodrigues. *Markov decision process applied to the control of hospital elective admissions.* Artificial Intelligence in Medicine 47, 159-171, 2009.

[11] WB Powell. *Approximate Dynamic Programming, Solving the Curses of Dimensionality.* 2nd ed, John Wiley & Sons, Inc., 2011.

[12] Rijksuniversiteit Groningen research database *Multivariate Statistics* `http://www.rug.nl/research/portal/files/9775153/c2.pdf.`

[13] SM Ross. *Introduction to probability models* 10th ed, Academic Press, 2010.

[14] B Sinclair *Coxian distributions.* `http://cnx.org/contents/JhSvg22o@5/Coxian-Distributions`

[15] Sint Maartenskliniek website. `https://www.maartenskliniek.nl/`

[16] B van der Staak. *Production planning in a specialized hospital in the Netherlands.* Master Thesis, Eindhoven University of Technology, 2016.

[17] HC Tijms. *Stochastic Modelling and Analysis:* A Computational Approach. John Wiley & Sons, Inc., 1988.

[18] YL Tong. *The multivariate normal distribution.* Springer-Verlag New York Inc. 1990

[19] PT Vanberkel, RJ Boucherie, EW Hans, JL Hurink, WAM van Lent and WH van Harten. *An exact approach for relating recovering surgical patient workload to the master surgical schedule.* Journal of the Operational Research Society 62, 1851-1860, 2011.

A One-step transition probabilities (Method 2)

Assume $\sigma_t = (\mathbf{j}, \mathbf{k}, l)$, with $\mathbf{k} = (k_1, k_2, ..k_{z_S})$ and $\mathbf{j} = (j_1, ..j_{z_R})$. Action a_t in week t results in C_t patients seen at the OD and D_t surgeries in week t. Recall that we denote the probability that we have m successes out of n trials while having success probability p by $B(m, n, p)$. The number of successes has a binomial distribution.

$$B(m, n, p) = \binom{n}{m} p^m (1-p)^{n-m}$$

The one-step transition probability for the OR queue length ($P(X_{t+1} = x_{t+1}|\sigma_t, a_t)$) remains the same as in Method 1.

Diagnostics:

$$P(R_{t+1} = (i_1, ..., i_{z_R})|\sigma_t, a_t) = P(R_{t+1,1} = i_1|\sigma_t, a_t) \prod_{n=2}^{z_R} P(R_{t+1,n} = i_n|\sigma_t, a_t)$$

$$= B(i_1, a_1, (1 - r_R(0))p_{OD,R}) \prod_{n=2}^{z_R} B(i_n, j_{n-1}, 1 - r_R(n-1))$$

Define the function h to be:

$$h(m, \mathbf{n}, p(x), z) = \sum_{m_1=0}^{m} \sum_{m_2=0}^{m-m_1} ... \sum_{m_{z-1}=0}^{v} \sum_{m_z=v-m_{z-1}}^{v-m_{z-1}} \prod_{x=1}^{z} B(m_x, n_x, p(x)),$$

with $v = m - \sum_{i=1}^{z-2} m_i$ and $n = (n_1, n_2, ..., n_z)$. The following equations hold:

$$P(m \text{ patients with diagnostics delay } \geq 1 \text{ to screening}|\sigma_t, a_t) = h(m, \mathbf{j}, (1-r_S(0))r_R(x)r_{RS}(x), z_R)$$

$$P(m \text{ patients with screening delay } \geq 1 \text{ to OR queue}|\sigma_t, a_t) = h(m, \mathbf{k}, r_S(x)r_{SO}(x), z_S)$$

$$P(m \text{ patients with screening delay } = 0 \text{ and diagnostics delay } \geq 1 \text{ to OR queue}|\sigma_t, a_t)$$
$$= h(m, \mathbf{j}, r_R(x)r_{RS}(x)r_S(0)r_{SO}(0), z_R)$$

Screening:

$$P(T_{t+1} = (i_1, ..., i_{z_R})|\sigma_t, a_t) = P(T_{t+1,1} = i_1|\sigma_t, a_t) \prod_{n=2}^{z_S} P(T_{t+1,n} = i_n|\sigma_t, a_t)$$

$$= P(T_{t+1,1} = i_1|\sigma_t, a_t) \Bigg(\prod_{n=2}^{z_S} B(i_n, k_{n-1}, 1 - r_S(n-1)) \Bigg),$$

with

$$P(T_{t+1,1} = i|\sigma_t, a_t)$$
$$= \sum_{n_1=0}^{i} B\big(n_1, a_1, (1 - r_S(0))p_{OD,T} + (1 - r_S(0))p_{OD,R}r_R(0)r_{RS}(0)\big) \times$$
$$h(i - n_1, \mathbf{j}, (1 - r_S(0))r_R(x)r_{RS}(x), z_R).$$

Arrivals to the OR queue

$$P(A_t = i|\sigma_t, a_t) = \sum_{n_1=0}^{i} \sum_{n_2=0}^{i-n_1} B(n_1, a_1, r_R(0) r_{RS}(0) r_S(0) r_{SO}(0) p_{OD,R}) \times$$
$$h(n_2, \mathbf{k}, r_S(x) r_{SO}(x), z_S) h(n - n_1 - n_2, \mathbf{j}, r_R(x) r_{RS}(x) r_S(0) r_{SO}(0), z_R)$$

B Approximate transition probabilities of Method 2

In this section we explain how the one-step transition probabilities of Method 2 can be approximated using a combination of the normal distribution and the multivariate normal distribution. The normal distribution will be used to approximate $P(X_{t+1}|\sigma_t, a_t)$ and the multivariate normal distribution will be used to approximate $P(R_{t+1}|\sigma_t, a_t)P(T_{t+1}|\sigma_t, a_t)$. We are first going to derive the latter. With $Y_1|Y_2, Y_3$ we denote the random variables Y_1 conditioned on the events Y_2 and Y_3.

Suppose that $\sigma_t = (\mathbf{j}, \mathbf{k}, l)$, with $\mathbf{k} = (k_1, k_2, ..k_{z_S})$ and $\mathbf{j} = (j_1, ..j_{z_R})$. And suppose that a_t results in C_t and D_t.

Diagnostics

Recall that

$$P(R_{t+1} = (i_1, ..., i_{z_R})|\sigma_t, a_t) = P(R_{t+1,1} = i_1|\sigma_t, a_t) \prod_{n=2}^{z_R} P(R_{t+1,n} = i_n|\sigma_t, a_t)$$

$$= B(i_1, a_1, (1 - r_R(0))p_{OD,R}) \prod_{n=2}^{z_R} B(i_n, j_{n-1}, 1 - r_R(n-1)).$$

Using the fact that under the right settings we can approximate a binomial distribution by a normal distribution, the following holds:

$$R_{t+1,n}|\sigma_t, a_t \sim B(j_{n-1}, 1 - r_R(n-1)) \approx N\Big(j_{n-1}\big(1 - r_R(n-1)\big), k_{n-1}\big(1 - r_R(n-1)\big)r_S(n-1)\Big)$$

for $n \in \{2, 3, ..., z_R\}$ and

$$R_{t+1,1}|\sigma_t, a_t \sim B(C_t, \hat{q}) \approx N\big(C_t\hat{q}, C_t\hat{q}(1 - \hat{q})\big),$$

with $\hat{q} = (1 - r_R(0))p_{OD,R}$. The multivariate normal approximation ([13]) for $R_{t+1}|\sigma_t, a_t$ is denoted by $\mathcal{N}(\boldsymbol{\mu}_1, \boldsymbol{\Sigma}_1)$, with

$$\boldsymbol{\mu}_1 = \Big(C_t\hat{q}, j_1\big(1 - r_R(1)\big), ..., j_{z_R-1}\big(1 - r_R(z_R - 1)\big)\Big)$$

$$\boldsymbol{\sigma}_1 = \begin{pmatrix} C_t\hat{q}(1-\hat{q}) & 0 & \cdots & 0 \\ 0 & j_1\big(1 - r_R(1)\big)r_R(1) & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & j_{z_R-1}\big(1 - r_R(z_R-1)\big)r_R(z_R-1) \end{pmatrix},$$

Screening

Using similar reasoning as for diagnostics and using the fact that the sum of independent normally distributed variables is again normally distributed, we get the following:

$$T_{t+1,n}|\sigma_t, a_t \sim B(k_{n-1}, 1 - r_S(n-1)) \approx N\Big(k_{n-1}\big(1 - r_S(n-1)\big), k_{n-1}\big(1 - r_S(n-1)\big)r_S(n-1)\Big)$$

for $n \in \{2, 3, ..., z_S\}$ and

$$T_{t+1,1}|\sigma_t, a_t \approx N\big(C_tq + \sum_{k=1}^{z_R} j_kq_k, C_tq(1-q) + \sum_{k=1}^{z_R} j_kq_k(1-q_k)\big),$$

with $q = (1 - r_S(0))p_{OD,T} + (1 - r_S(0))p_{OD,R}r_R(0)r_{RS}(0)$ and $\mathbf{q} = (q_1, ..., q_{z_R}) = \Big((1-r_S(0))r_R(1)r_{RS}(1), ..., (1-r_S(0))r_R(z_R)r_{RS}(z_R)\Big)$. The multivariate normal approximation for T_{t+1} is denoted by $\mathcal{N}(\boldsymbol{\mu}_2, \boldsymbol{\Sigma}_2)$, with

$$\boldsymbol{\mu}_1 = \left(C_t q + \sum_{k=1}^{z_R} j_k q_k, k_1\big(1 - r_S(1)\big), ..., k_{z_S-1}\big(1 - r_S(z_S - 1)\big)\right)$$

$$\boldsymbol{\Sigma}_1 = \begin{pmatrix} C_t q(1-q) + \sum_{k=1}^{z_R} j_k q_k(1-q_k) & 0 & \cdots & 0 \\ 0 & k_1\big(1 - r_S(1)\big)r_S(1) & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & k_{z_S-1}\big(1 - r_S(z_S-1)\big)r_S(z_S-1) \end{pmatrix}$$

OR queue length
Recall that we have used the following equation to derive the transition probability for the OR queue length:

$$X_{t+1} = not\_plan_t + A_t + \max\{0, plan_t - D_t\}.$$

Rewriting this equation we get:

$$X_{t+1} = X_t + A_t - \min\{plan_t, D_t\}.$$

Let us have a look at the last part of the second equation:

$$P(\min\{plan_t, D_t\} = i|\sigma_t, a_t) = \begin{cases} B(i, X_t, 1 - r_t) & \text{if } i < D_t \\ \sum_{k=D_t}^{X_t} B(k, X_t, 1 - r_t) & \text{if } i = D_t. \end{cases},$$

If $i < D_t$ we can approximate the probability with a normal distribution. When $i = D_t$ we take the sum of the tail of a binomial distribution, which can be approximated by the sum of the tail of a normal distribution. Note that we can also approximate $A_{t+1}|\sigma_t, a_t$ by a normal distribution. This goes in a similar fashion as the normal approximation of $T_{t+1,1}|\sigma_t, a_t$.

In conclusion we cannot approximate $X_{t+1}|\sigma_t, a_t$ with a normal distribution, but we can use combinations of normal distributions to approximate this probability. We do have a multivariate normal approximation for $R_{t+1}, T_{t+1}|\sigma_t, a_t$: $\mathcal{N} = (\boldsymbol{\mu}, \boldsymbol{\Sigma})$, with $\boldsymbol{\mu} = (\boldsymbol{\mu}_1, \boldsymbol{\mu}_2)$ and

$$\boldsymbol{\Sigma} = \begin{pmatrix} \boldsymbol{\Sigma}_1 & \mathbf{0} \\ \mathbf{0} & \boldsymbol{\Sigma}_2 \end{pmatrix}$$

Unfortunately, these approximations do not speed up the computations by much.

C Variables and parameters

Parameter or variable	Description
	Budget
OD_{ses}	Budgeted amount OD sessions for the year.
OR_{ses}	Budgeted amount OR sessions for the year.
i_t	Number of OD and OR sessions left to distribute from week t until the end of the year is equal to $i_{t,1}$ and $i_{t,2}$ respectively.
	Random variables
R_t	Number of patients in diagnostics, with $R_{t,i}$ the number of patients that will complete i weeks of being in state diagnostics in week t.
T_t	Number of patients in screening, with $T_{t,i}$ the number of patients that will complete i weeks of being in state screening in week t.
X_t	Number of patients in the OR queue at the beginning of week t.
A_t	Number of patients arriving to OR queue in week t.
U_t	Amount of unused OR time slots in week t.
$plan_t$	Amount of patients who wants to be scheduled for the OR in week t.
not_plan_t	Amount of patients who does not want to be scheduled for the OR in week t.
	Action and schedule
a_t	Action taken in week t. It results in D_t and C_t.
$\hat{a}$	Initial schedule that assigns to each week a average number of OD and OR sessions relative to the amount of workdays in that week.
C_t	Number of patients seen at the OD in week t.
D_t	Number of patients required to fill the OR time in week t .
	Parameters
W_{goal}	A patient has the opportunity to be scheduled for the OR within W_{goal} weeks with probability greater than or equal to 0.8.
z_R,z_S	Maximum diagnostics an screening delay respectively.
b_t	Number of workdays in week t.
r_t	Probability that a patient does not want to be scheduled for the OR in week t.
h_{OD}	Average number of time slots during an OD session.
h_{OR}	Average number of surgeries per OR session.
σ_t	State at the beginning of week t. For ease of notation we often refer to it as the realization of the state space.
w_{ij}	Realization of the number of patients that goes from state i to state j in a certain week.
w_i^d	Realization of the number of patients that leaves state i in a certain week.
	Sets and bounds .
(s_t, S_t)	Lower and upper boundary respectively for the OR queue length at the beginning of week t.
$U^R_{t,i}$	Upper boundary for $R_{t,i}$. For Model 2 we use U^R_t.
$U^S_{t,i}$	Upper boundary for $T_{t,i}$. For Model 2 we use U^S_t.
$\mathcal{S}_t$	All states that satisfy $s_t \leq X_t \leq S_t$, $0 \leq R_{t,i} \leq U^R_{t,i}$ and $0 \leq T_{t,i} \leq U^S_{t,i}$.
$\mathcal{W}$	Set of all valid realizations of the w_{ij}'s.
$A_t(\sigma_t, i_t)$	Set of valid actions when in week t we are in state σ_t and have budget i_t left.

D Planning horizon (alternative approach)

In this section we briefly explain two alternative approaches on how to compute the probability that OR queue length drops below the lower boundary during the planning horizon. This approach was our initial idea, but turned out not to work so well due to the large probability matrices needed.

Suppose we decide in week k the amount of OR and OD sessions done per week in the weeks $k+n-3$, ..., $k+n$. We are interested in the probability that the OR queue length will drop below the lower boundary anywhere between week k and week $k+n$. For this approach we use the Chapman-Kolmogorov equations to compute $P(\sigma_{k+i}|\sigma_k, a_k)$ for $i \in \{1, 2, ..., n\}$. Similar to the approach in section 4.3, we are interested in constructing a transition probability matrix, here called $\tilde{P}$. We want to construct this matrix such that $\tilde{P}^x_{ij}$ is the probability that we are in state j at the beginning of week $k+x$ given that at the beginning of week k we are in state i. Note that this probability depends on the actions taken and the rescheduling probability during this period. The actions and the rescheduling probability can differ per week, which implies that we cannot construct a matrix $\tilde{P}$ that is time invariant if we want to have an exact approach. We have two options:

Option 1
Make $\tilde{P}$ time dependent and call this new matrix $\tilde{P}_t$. With $\tilde{P}_{t,i,j}$, we denote the probability that we are in state j at the beginning of week $t+1$ given that we were in state i at the beginning of week t and we took an action in week t according to the initial schedule $\hat{a}$. This probability can be computed using the equations for the one-step transition probabilities of our model. Furthermore, the queue length boundaries can differ per week, which is another argument for making $\tilde{P}$ time dependent. We only include states with an OR queue length of at most $\lceil 1.1S_t \rceil$ and make states that have a queue length lower than s_t absorbing states. We can include this absorbing state in the $\tilde{P}_t$'s by adding an extra zero row and zero column to them with the last element of that row and column equal to 1. Suppose $\sigma_k = \mathbf{l}$ and $\sigma_{k+i} = \mathbf{j}$, then $P(\sigma_{k+i}|\sigma_k, a_k) = [\tilde{P}_k\tilde{P}_{k+1}...\tilde{P}_{k+i-1}]_{lj}$, which is the element of the product of these matrices that corresponds to a transition from $\mathbf{l}$ to $\mathbf{j}$. Due to the added absorbing state to the $\tilde{P}_t$'s, we can also derive the probability that we are in an absorbing state at the beginning of week $k+i$. The $\tilde{P}_t$'s can become very large very easily, because we have to include every possible combination of R_t, T_t and X_t and we have to construct such a matrix for each week of the year.

Option 2
For this option, instead of computing i different matrices, $\tilde{P}_k$,$\tilde{P}_{k+1}$,...,$\tilde{P}_{k+i-1}$, we use one matrix we call $\bar{P}_k$ as a basis. Recall from option 1 that the $\tilde{P}_t$'s can differ based on the rescheduling probability, absorbing state and the action taken in week t. For option 2 we assume a fixed action and rescheduling probability in the weeks k up to and including $k+n-1$. When making an updated production plan in week k we take the following rescheduling probability and action in the weeks k up to and including $k+n-1$:

$$\bar{r}_k = \sum_{j=0}^{n-1} r_{k+j} \text{ and } \bar{a}_k = \frac{1}{n}(\sum_{j=0}^{n-1} \hat{a}_{k+j,OD}, \sum_{j=0}^{n-1} \hat{a}_{k+j,OR}).$$

We still have to explain how we deal with the varying lower boundaries. In $\hat{P}_k$ we include all states with queue lengths at most equal to $\max_{t\in\{k,...,k+n\}}\lceil 1.1S_t \rceil$ and all states with queue length of at least $\min_{t\in\{k,...,k+n\}} s_t$. Then we fill this matrix by using the equations for the one-step transition probabilities of our model. When we have $\hat{P}_k$, we can for week $k+i$ adjust this matrix such that states corresponding to an absorbing state in week $k+i$ are made absorbing states by setting these rows equal to a vector $(0,...,0,1)$ of appropriate length. Just like for the $\tilde{P}_t$'s, $\hat{P}_k$ is constructed such that the last row and column always correspond to the absorbing state.

In conclusion, we have proposed two alternative methods to compute the probability that the OR queue length will drop below the lower boundary during the planning horizon. Option 2 is faster than option 1, but less exact. We advise to only use this method for small cases where the transition matrix will not be too large to handle. For a case the size of the SMK we advise to use the method explained in section 4.3.

E Implementation of the planning horizon model

In this section we explain how we have implemented the recursive formula explained in section 4.3. Suppose we make a schedule in week k and we have a planning horizon of n weeks. Define the matrix $H^{(t)}$ ($t \in \{1, 2, ..., 51\}$) as the matrix containing all the transition probabilities when going from $X_t = x$ ($s_t \leq x \leq \lceil 1.1S_t \rceil$) to $X_{t+1} = y$ ($s_{t+1} \leq y \leq \lceil 1.1S_{t+1} \rceil$) or to $X_{t+1} < s_{t+1}$. So for $s_t \leq x \leq \lceil 1.1S_t \rceil$ we get:

$$H^{(t)}_{x-s_t+1,j} = \begin{cases} P(X_{t+1} < s_{t+1} | X_t = x, \hat{a}) & \text{if } j = 1 \\ P(X_{t+1} = j - 2 + s_{t+1} | X_t = x, \hat{a}) & \text{if } 2 \leq j \leq (\lceil 1.1S_{t+1} \rceil - s_{t+1} + 2) \end{cases}$$

Notation:

- $H^{(t)}_{.,1}$ is the first column of $H^{(t)}$. It corresponds to transitions into the absorbing state of week $t + 1$
- $H^{(t)}_{.,-1}$ is $H^{(t)}$ without the first column. It corresponds to transitions into the non-absorbing states of week $t + 1$

Define the matrix L as follows, its i-th column is equal to:

$$L_{.,i} = \begin{cases} H^{(k)}_{,1} & \text{if } i = 1 \\ H^{(k)}_{,-1} ... H^{(k+i-2)}_{,-1} H^{(k+i-1)}_{,1} & \text{if } 2 \leq i < n \end{cases}$$

Here $L_{i,j}$ is the probability that $X_{k+j} < s_{k+j}$ given that $X_k = i + s_k - 1$ and $s_t \leq X_t \leq \lceil 1.1S_t \rceil$ for $t = k, k + 1, ..., k + j - 1$.If we then take the sum of each row of L, we get the probability that given $X_k = x$, for any $s_k \leq x \leq \lceil 1.1S_t \rceil$, the OR queue length will drop below the lower boundary anywhere between week $k + 1$ up to and including week $k + n$.

F Parameters test cases monotonicity Model 2

In this section we present the parameters used in the test cases described in Section 5.4 for Model 2. Test case 1 and 3 use the transition probabilities described in Section 6.

OD_{ses}	4	w_{goal}	4
OR_{ses}	8	h_{OD}	10
r_t	(0.1,0.1,0.2,0.2)	h_{OR}	2
b_t	(2,3,4,5)	Horizon	4 weeks
ζ_R	0.02	ζ_S	0.05

Table 7: Parameters test case 1 and 2

	OD	Diagnostics	Screening	OR	Home
OD	0.4397	0.0362	0.0730	0.0323	0.4188
Diagnostics	0	0.9473	0.0164	0.0026	0.0337
Screening	0	0	0.7807	0.2149	0.0044

Table 8: Transition probabilities test case 2

OD_{ses}	3	w_{goal}	3
OR_{ses}	6	h_{OD}	8
r_t	(0.1,0.1,0.2,0.1)	h_{OR}	2
b_t	(2,3,3,5)	Horizon	4 weeks
ζ_R	0.02	ζ_S	0.05

Table 9: Parameters test case 3