CONTROLLING UNCERTAINTY: OPTIMIZING CAPACITY MANAGEMENT IN AN INTENSIVE CARE UNIT

Master Thesis

Markus Lorenz Hahn, 489788

In Partial Fulfilment
of the Requirements for the Degree in
MSc Supply Chain Management

Thesis Coach: Dr Cynthia Kong
Co-reader: Dr Robert Rooderkerk

Company supervisor: Dr Lucia S.D. Jewbali

Rotterdam School of Management
15/08/2018

PREFACE

ACKNOWLEDGMENT

I would like to express my gratitude to everyone who provided support and guidance during the process of creating this thesis.

First to my thesis coach Dr Cynthia Kong: for approaching me with this highly interesting project that brought me in touch an environment I was not familiar with and subsequently enjoyed learning a lot about, as well as the regular helpful advice and guidance I received that was essential to bring the project to a fruitful completion.

To my company coach Dr Lucia S.D. Jewbali: For the very welcoming atmosphere at the Erasmus MC and the impressive amount of time, access & information I was granted. I learned a lot during every visit and was deeply impressed by the work done at the ICU every single day.

To my co-reader Dr Robert Rooderkerk: For the challenging comments on my work and sharing his expertise in academics and operations research that greatly improved the quality of these aspects in my work.

Finally to my family, which is a constant anchor of support in all situations of life.

EXECUTIVE SUMMARY

This thesis studies the intensive care unit of the Erasmus MC, an academic hospital in Rotterdam. It receives cardiology patients of two different kinds, unplanned emergency patients as well as patients that underwent elective surgery. The ICU environment is characterised by high uncertainty with regard to the two key parameters that determine its capacity usage, patient arrivals and patient length of stay. Providing life-critical care is a resource-intensive process and the uncertain environment makes efficient capacity planning difficult. Since patients arrive in a critical state, immediate care is required for any arrival. Should the ICU be at its capacity limit a patient needs to be prematurely discharged to make room, a consequence that might negatively affect the discharged patient. As part of a relocation to a new facility, the ICU unit was subjected to a restructuring that combined two separate units into a single one that now needs to care for a combined and re-defined patient mix. Faced with the challenge of adapting to its new setting, the lack of capacity management policies became evident and the question which measures can help the ICU to increase its level of patient care by minimizing patient dismissal was posed.

A review of existing academic literature on ICU capacity management set the basis for the capacity management policy development. Previous studies confirm that capacity optimization in a hospital is a relevant endeavour, especially in light of the pressure of patient growth and the need for better care (Green, 2006). Common methods to study ICU's are queuing models and simulation. However, the assumptions of a queuing model do not necessarily hold in a critical care environment (Green, 2002) and discrete event simulation is able to better represent

the operational dynamics (Günal & Pidd, 2010). For this project, the combination of solution finding and accurate modelling that can be achieved by combining optimization with discrete event simulation was determined to be the best method.

The majority of existing research improves ICU capacity management by utilizing levers not available in this study, such as allocating beds flexibly across hospital wards (Kim et al., 2000), regional collaboration (Litvak et al., 2008) and influencing available capacity (Harper & Shahani, 2002, Holm et al., 2013). All studies define their improvements by reducing the number of occurrences of a negative consequence, most commonly the refusal of patient admissions, while premature discharge was only at the centre of one prominent study by Dobson et al., 2010.

Accurately including the uncertainty that the ICU faces requires close approximation of the patient length of stay and arrival rates. Using categorical averages for these variables does not reflect reality and the approximation needs to be preceded by a patient classification that reduces within-group variability (Ridge et al., 1998). The long tail length of stay distributions can be estimated by fitting right-tailed probability distributions to the data (Harper, 2002) and arrival rates are best approximated by a time-varying Poisson arrival rate (Green, 2006).

Before the heuristic development process is started, a thorough statistical analysis of the patient population is executed, based on a data set that contains 12,170 admissions to the ICU over the span of four years. The diagnosis categories of the Erasmus MC are used as patient type classification to reduce within-group variability without compromising the interpretability of the study results. The classification is achieved by text mining the diagnosis information of each patient for keywords that are used to assign patient types.

The main insights from the statistical analysis are:

- The split between planned and unplanned patients is 37% vs 63%
- All patient types show significant length of stay outliers
- Most patients (73%) leave the ICU in less than 24 hours
- 53% of ICU capacity is used by only 14% of long-stay patients
- The largest patient classes (STEMI, Overige, CABG, Ritmestoornissen) stay only a few hours but have a very variable length of stay
- Certain patient types are characterised by a very long length of stay (ECMO, LOTX, HTX, OHCA, LVAD, Hartfalen)
- The OHCA and Hartfalen patient types use a disproportionate amount of capacity (19% and 13%) compared to their share of the patient population (5% for both)
- Patient arrivals have slightly increased over the past few years (+4.5%)
- There is a seasonal difference between summer and winter arrivals, probably caused by reduced staffing during the summer months
- There is within-week seasonality, caused by shifts in the elective patient schedule and unplanned arrivals
- There is daily seasonality with a strong peak at midday

Based on the analysis, arrival rates were modelled as time-varying Poisson distributions per patient type, considering the discovered seasonality and length of stay distributions were fitted for each patient type. It also revealed the improvement possibility of smoothing the rescheduling of elective surgeries across the week to improve capacity usage at the ICU.

The heuristic development process begins by defining a regular weekly elective surgery schedule by aggregating planned patients into five larger groups according to their length of stay characteristics. This is done since most of the planned patient categories are too small and too variable to be considered on a weekly basis. Next, a new variable that estimates average expected capacity use per patient type is defined by combining the arrival frequency with patient survivor functions which are generated by applying survival analysis to the length of stay statistics. This so-called "Loadfactor" successfully imitates the capacity usage "lag" across various days caused by the probability of extended length of stay.

An integer optimization programme allocates the newly defined planned patient Loadfactor across the week using a squared-sum objective function that evenly balances the planned patient load with the unplanned patient load. This is achieved by discouraging peaks and moving long stay patients towards the end of the week so that free capacity on the weekend is better utilized. By evening out the expected average patient load, the probability of the ICU reaching its capacity limit is reduced, which in turn reduces the rate of premature discharges.

The arrival smoothing heuristic is tested by replicating the uncertain ICU environment in a discrete event simulation. After successfully validating the accuracy of the simulation, the heuristic is tested versus the base case and successfully reduces the rate of premature discharge by 20.47%.

This newly developed heuristic succeeds in improving capacity management at the ICU and provides the opportunity for future research to iterate and improve on it.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

1. INTRODUCTION

1.1 Background and Motivation

1.1.1 *Erasmus Medical Center*

The Erasmus MC is the largest hospital in the Netherlands (CIBG, 2018) and aims to both provide excellent patient care, as well as to be a leading institution in healthcare research and education. The origins of Erasmus MC as a hospital date back to the 19th century, when Rotterdam's first hospital, the Coolsingel Hospital, was constructed. In 1965 a medical faculty was established to create an academic hospital. The Erasmus MC in its current form exists since 2002 when a collaboration agreement with Erasmus University Rotterdam was finalized. At its location in the centre of Rotterdam Erasmus MC provides the entire range of medical services and also houses its two specialized branches, Erasmus MC-Sophia, a children's hospital, and the Erasmus MC Cancer Institute (Erasmus MC, 2018).

1.1.2 *Department of Cardiology & ICU unit*

In order to provide specialized care for patients with chest disorders, Erasmus MC created the Thoraxcenter. It consists of the department of pulmonology, the department of thoracic surgery as well as the department of cardiology. Patients can receive care for any disease that affects their lungs, heart, throat or large blood vessels. Examples include heart attacks, heart failures, pulmonary hypertension as well as heart and lung transplants (Erasmus MC, 2018). The unit of analysis of this thesis is the intensive care unit within the department of cardiology, as shown in Figure 1-1.

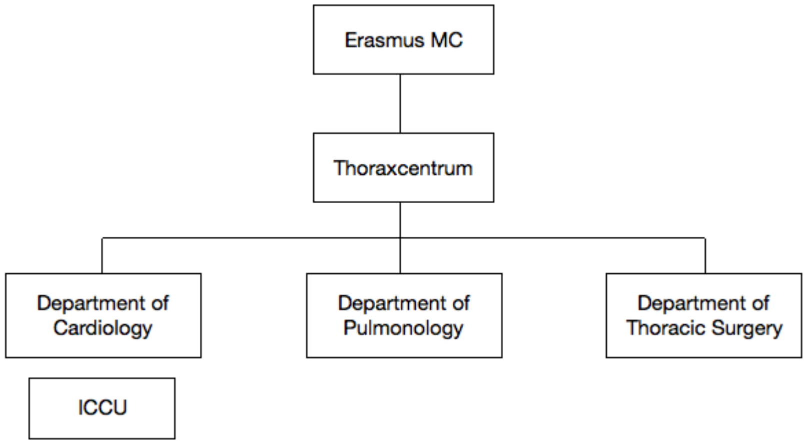

Figure 1-1: Organigram Erasmus MC

An intensive care unit (referred to as ICU from now on) is set up to care for critically ill patients by employing specialized doctors and nurses, as well as having the necessary equipment to enable constant monitoring of a patient's health status (Kelly, 2018). The ICU in question consists of two units, namely the ICCU and ICTH. The ICCU receives emergency cardiology patients that arrive via ambulance or transfer from other hospitals and currently consists of 12 beds. The ICTH, on the other hand, receives elective surgery patients from the department of thoracic surgery and consists of 7 beds. A bed is only available if a dedicated nurse is staffed. The ICU works with a shift system of three eight-hour shifts, which are from 8:00 am - 4:00 pm, 4:00 pm - 12:00 am and 12:00 am - 8:00 am.

1.1.3 *Patient types*

The ICU differentiates between patient types according to diagnosis. The appendix depicts an overview and description of these categories, which are separated into two kinds, "*unplanned*" (emergency patients, Table A-1) and "*planned*" (elective surgery patients, Table A-2) patients.

1.1.4 *Patient flows*

The ICU receives patients from various sources and is only a temporary care facility for the most critical stages of a patients healing process. It is thus important to understand the patient flows that run through the ICU, an overview of which is depicted in Figure 1-2.

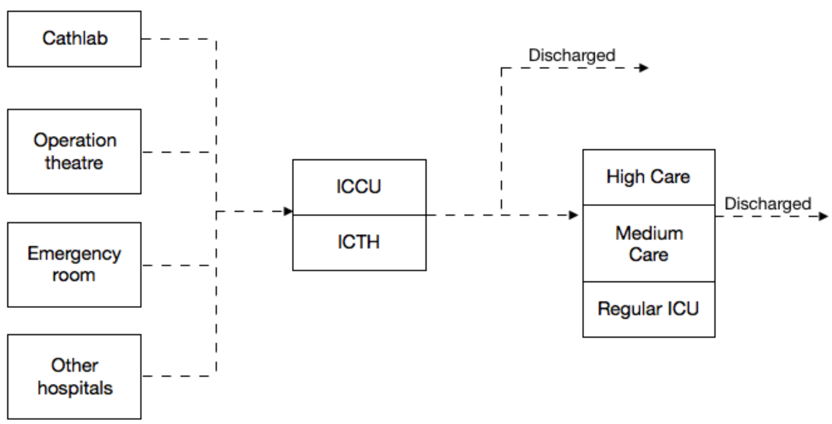

Figure 1-2: Patient flow model

Patients within the hospital arrive from the operation theatre and cath lab[1] for recovery. Newly admitted patients arrive via the emergency room or are transferred from other hospitals. After a sufficient degree of recovery, patients are discharged to lower level care facilities in the hospital (High care or medium care), a non-cardiology ICU or, rarely, discharged to their homes.

[1] A cath lab is a specialized hospital facility where cardiac catheterization is performed, a procedure that visualizes blood flow through the heart to diagnose potential heart issues (Heart.org, 2017)

1.1.5 *Bed planning at the ICU*

All elective surgeries are scheduled by the surgery department according to the operation room and surgeon availability, considering availability for emergency arrivals as well. There is limited coordination of the surgery schedule with the ICU. The ICU receives a weekly overview of elective surgery arrivals and prepares its beds accordingly. To ensure a sufficient level of attention, a bed can only be used if a dedicated nurse is available, the actual availability of beds thus fluctuates according to staffing levels. The head nurse of the department prepares a daily overview of patients currently being cared for and all expected arrivals on a whiteboard in the ICU. During their daily rounds, which happen at least twice a day, the doctors check the patients and, given a sufficient degree of recovery, decide do discharge patients to lower level care. In the case of overcrowding, a bed is made available by bumping a patient to a lower level care facility earlier than planned. It is the Thoraxcenter's policy to not cancel scheduled surgeries in such a situation.

1.1.6 *The move of the ICU*

The growth of patient numbers, as well as the wish to futureproof the Erasmus MC hospital, led to the decision to construct a new hospital building at the current site. Construction work on the west wing, which will house the ICU, was completed and in May 2018 the ICU moved to the new premises. The state-of-the-art facilities improve many aspects of the daily operations of the ICU but pose a challenge for the ICU management since it is combined with a reorganization of the ICU setup. The two separate units, the ICCU and the ICTH were integrated into one and certain patient categories are no longer being served by the ICU, resulting in a new patient mix that changes the patient load and arrival patterns at the ICU unit and requires a reassessment of ICU capacity management.

1.1.7 *Problem description*

The ICU expects patient levels to increase in the coming years and faces the challenge to provide its service with a limited number of beds. The ICU aims to accommodate and provide service to both planned and unplanned emergency patients. The uncertainty of emergency patient arrivals as well as the uncertainty of the actual length of stay (referred to as "*LOS*" from now on) of patients in the ICU, which furthermore differs greatly across patient types, results in a very difficult bed planning process. As described in the paragraph on capacity management, the bed planning efforts currently employed at the ICU are very limited and created manually on a day-by-day basis. There are no estimations of emergency patient arrivals and the actual patient LOS. Due to its policy to not cancel elective surgeries, situations of overcrowding lead to premature discharge that negatively affect patient care. Simultaneously, a high level of utilization of the ICU is important due to its high cost. The ICU wishes to improve its bed planning by introducing these estimations and creating a bed capacity management scheme that reduces the number of premature discharges by understanding the differences in arrivals and LOS of each patient type and applying this knowledge to better balance patient type composition.

1.2 **Research Objective**

Erasmus MC's ICU will be reduced in size and currently does not have a data-driven bed management scheme in place. The research objective is to assist the ICU in developing a bed capacity management scheme that is based on estimations of patient arrival rates and length of stay per patient type to improve bed planning at the ICU and ultimately its ability to provide better care to its patients.

1.2.1 *Research questions*

The main research question is:

"What bed capacity management scheme can help the Erasmus MC's intensive care unit to increase its level of patient care by minimizing premature patient dismissal?"

The following sub-questions support the main research question:

I. Are there arrival pattern differences across time and per patient type the ICU receives?
II. How can patient arrivals per type be modelled?
III. What are the LOS characteristics exhibited per patient type?
IV. How can LOS per patient type be modelled?
V. How does the combination of patient arrival rates and differing LOS influence capacity usage at the ICU?
VI. How should the intensive care unit improve the management of its limited bed capacity?

1.3 **Assignment and Deliverables**

The objective of the ICU at Erasmus MC is to professionalize its bed planning process by leveraging the availability of historical data to generate a bed capacity management scheme per patient type that minimizes the occurrences of premature dismissal.

The aim of this thesis project is to support this effort by delivering the underlying analyses and heuristic to the ICU in two stages. In the first stage, historical patient data of arrivals and diagnosis are analyzed using seasonal analysis and data mining to estimate arrival rates and length of stay per patient type.

In the second stage of the thesis project, a heuristic that aims to improve ICU performance will be developed on the basis of the statistical analysis. After optimizing it, a simulation that accurately models the ICU environment by including the factors estimated in the first stage will be used to validate the heuristic. Finally, a sound and actionable recommendation for the ICU is formulated.

1.3.1 *Managerial contribution*

This research contributes to practice by helping the ICU at the Erasmus MC to cope with the challenges described in this chapter. More specifically, it aims to alleviate the capacity management problem the ICU faces by contributing findings on two different levels. On the one hand, the statistical analysis and visualisation of the historical data will uncover facilitate the understanding of patient dynamics at the ICU as well as uncover trends and abnormalities. Secondly, the integration of the heuristic developed in this thesis into the capacity management process at the ICU will enable the ICU to better accommodate patients. This, in turn, will lead to more patients receiving better care at the ICU, increasing patient well-being.

1.3.2 *Academic contribution*

The ICU as part of the larger hospital service system represents an entity with a lot of potential for improvement through the application of operations research. Green (2005) notes that hospitals are under a lot of pressure to provide better care with limited resources, leading to a need to address operational inefficiencies and a drive to optimize operations. In their analysis of operations & supply chain management research in healthcare Dobrzykowski et. al. (2014) find that capacity planning is one of the central areas where the application of operations research has led to improvements. Harper (2002) explains that the academic challenge in modelling any hospital resource is the accurate consideration of the complexity, uncertainty, variability and resource limitations that set the dynamics of a hospital unit apart.

Academically this research will extend existing literature by developing a new capacity management scheme that aims to improve capacity management decision making under uncertainty. The heuristic development process will be novel due to its focus on aspects of ICU capacity planning that have received less attention so far, such as bumping as the only capacity alleviation option and the combination of the optimization and discrete event simulation methods during the development. A closer examination and inclusion of seasonality in arrival rates will improve the current practice of simply assuming a general Poisson distribution for patient arrivals and considering the differentiated length of stay by type instead of a simple average or overall distribution will ensure validity by addressing common criticisms of earlier models.

1.4 **Conceptual Project Design**

1.4.1 *Research strategy*

The project at hand is an example of a problem-solving design study, more specifically the research objective can be interpreted as an actionable recommendation to improve a business process. Since the aim is to deliver the ICU at Erasmus MC with an applicable solution, a detailed analysis of their situation is required. The design study will follow the steps of the problem intervention cycle:

1) Problem definition

During meetings with the director of the ICU as well as visits to the hospital the implications of the upcoming reduction in the size of the ICU due to the move were discussed and initial ideas on how to optimize the bed planning process developed.

2) Problem diagnosis

As guidance for the solution finding process, a detailed analysis of historical bed occupancy data of the ICU is necessary. To develop a fitting solution, the aim of the ICU needs to be defined and conceptualized. Furthermore, the actual distribution of length of stay and patient arrivals per type needs to be compared with the categorical assumptions currently in place.

3) Design of solution

Further discussions with the ICU staff will be held to define the design criteria of the tool and how it can fit into their bed planning process. User requirements in form of constraints of a bed capacity management scheme are of special importance. Quantitative analysis of historical data will be the basis to deliver the estimations of the previously unobserved uncertain arrival rates and length of stay per patient type. Based on inputs from the ICU management, as well as the study of previous academic literature, a heuristic that may improve capacity management is derived and optimized. A simulation model will include these inputs and will be used to test the effectiveness of the proposed heuristic.

4) Solution implementation

An implementation plan for the recommendations will be developed which defines how they can be integrated into the current bed planning process at the ICU.

5) Evaluation

The initial simulation-based evaluation of the effect of the developed recommendations will be recorded and in a future iteration of the project actual implementation can be evaluated.

1.4.2 *Research process model*

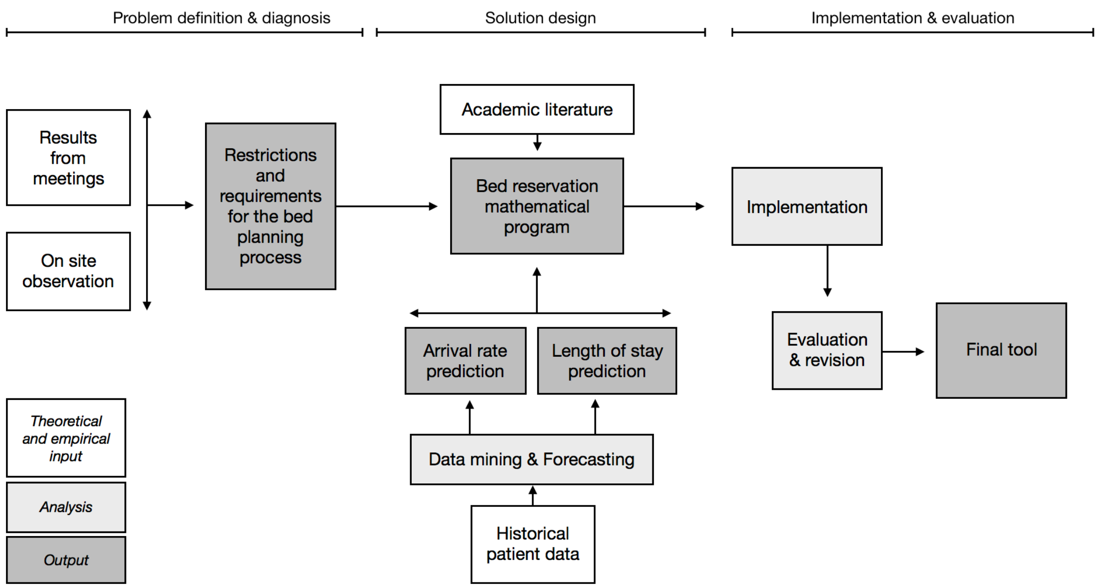

Figure 1-3: Research process model

2. LITERATURE REVIEW

2.1 Exploration of Practice

One of the main challenges faced by intensive care units is capacity constraints in face of increasing patient numbers. The OECD reports that in the next 30 years the share of the population over 65 in the Netherlands will increase from currently 18% to 28%. The ageing population is expected to increase the demand for critical care (OECD, 2017). Shortages of critical care capacity already are an issue in the Netherlands. Reports of ICU capacity shortages in the news led the Dutch health ministry to commission a thorough analysis of ICU capacities in the Netherlands. The Julius Centrum voor Huisartsgeneeskunde en Patiëntgebonden Onderzoek surveyed 113 ICU units (87% of all units in the Netherlands) over 4 months in 2001. The study found that around 10% of patients arriving during that time span were rejected and the non-availability of a bed caused 65% of these rejections. Furthermore, it found that the causes for non-availability of beds were more complex than bed capacity limits and included the efficiency of bed use (Hautvast et al., 2001).

Apart from the patient service considerations, efficient utilization of ICU beds is also important due to the high costs involved. Even though they are small in size, intensive care units account for a sizeable share of hospital costs, in certain cases up to 20% of a hospitals budget (Chalfin, 1995). An evaluation of 51 intensive care units in Germany found that these costs stem from the highly specialized equipment in use, as well as from expensive treatments that are common in critical care. The majority of the cost, however, is caused by staffing, which accounts for 56,1% of ICU cost, due to requiring full-time physicians as well as staffing nurses relative to the number of occupied beds (Moerer et al., 2007).

The issues of service degradation due to overcrowding and cost-efficient utilization of ICU resources are both highly relevant for practitioners and directly connected to the management of bed availability, resulting in a very worthwhile field of research.

2.2 **Exploration of Theory**

Modelling of intensive care unit capacity has seen increasing attention over the last few years. The before mentioned trade-off between underutilization of costly resources and the necessity to have critical care capacity available, as well as the uncertain environment, make it an interesting but challenging application of operational research methods.

2.2.1 *Capacity planning in healthcare - methods*

In general, two methods see wide application in the realm of intensive care unit capacity analysis, namely simulation and queuing analysis. Both methods allow modelling the stochastic environment of hospitals and find application when describing patient flow through the hospital as well as when studying single units, such as the ICU.

Queuing analysis is regularly applied when analyzing a hospital as a service system (Litvak et al., 2008, Kim et al., 1999). It is seen as a good fit to model healthcare environments due to its ease of use and low data requirements, it allows to quickly understand resource utilization and service performance (Green, 2006). The typical M/M/s model finds application but is usually extended to better fit the healthcare environment, for example by adding pre-emptive/non-pre-emptive priority for time-sensitive arrivals or by accounting for finite capacity (M/M/s/K model) (Green, 2006). Another variation is to model refused admissions at the ICU using the Erlang loss system (M/M/c/c model) since this phenomenon is comparable to blocking in a circuit switched telephone system (Litvak et al., 2008, Bruin et al., 2007). A queuing model can be problematic when studying an

ICU since the assumption that new patient arrivals are able to wait in a queue does not necessarily hold in a critical care environment (Green, 2002). Additionally, Green (2002) notes that the Markovian assumption of exponential service times does not hold due to the highly variable length of stay in the intensive care unit. A multi-server variation of the M/G/1 model was developed by Allen (2014) which manages to account for the increased variability of service time.

The second technique that finds application when modelling a healthcare environment is creating a simulation (Griffith et al., 2005, Costa et al., 2003, Ridge et al., 1998). Discrete event simulation is especially suited for examining "What if?" scenarios when studying resource use or patient flow in a hospital (Jacobson et al., 2006). Discrete event simulation allows modelling patient arrivals and discharges as discrete events that change the state of the unit in question to arrive at an understanding which factors influence ICU performance measures (Günal & Pidd, 2010). The simulation method also allows the incorporation of variability into the model, Harrison et al. (2005) rely on a Monte Carlo style setup that includes multiple random sampling of variables such as the patient arrival rate. Another type of simulation recognized by Günal & Pidd (2010) as relevant is compartmental modelling, which allows modelling LOS distributions by splitting them into discrete time brackets as well as how patients move between care units. Simulation is often applied to validate queuing models and examine the effect of certain assumptions (Kao & Tung, 1981, Ridge et al., 1998, Kim et al., 1999).

While these methods are well suited to test assumptions such as admission policies, they are no optimization tools by itself. Combining the ability of discrete event simulation to recreate the dynamics of uncertain environments with the optimization power of mathematical programs was done by Butler et al. (1992). Very uncommon in the literature, it promises to be a potentially effective measure and will be applied in this study.

2.2.2 *ICU capacity planning - heuristics*

The majority of research (Harper & Shahani, 2002, Ridge et al., 1998) aims at determining the optimal size of ICUs to avoid capacity constraints, assuming the possibility to increase the number of beds in the ICU. Similarly, other authors take a multi-unit view of the issue and examine how beds should be distributed across different hospital wards (Lapierre et al., 1999, Holm et al., 2013, Kao & Tung, 1981). Others, such as Kim et al. (2000) propose a somewhat flexible bed capacity that is shared with other hospital units. The dynamics between hospital wards were examined by Bruin et al. (2007), who find that economies of scale apply to hospital wards and that capacity constraints disrupt upstream care units. Litvak et al. (2008) aim to alleviate capacity constraints through better coordination with regional hospitals, reserving beds for regional overflow allows for better overall care.

Most ICU capacity studies aim to optimize patient care by reducing the occurrence of one of three negative consequences of overcrowding. The first and most popular being *patient refusal (or transfer)* in which a newly arrived patient is not admitted to the ICU in case of no bed availability and either referred to another hospital, to another ICU unit within the hospital or admitted to a lower care facility. The measure which is optimized is either the share of refusals or the probability of being refused (Litvak et al., 2008, Harper & Shahani, 2002). Another examined measure is the *cancellation of an elective surgery*, in which a pending surgery is cancelled to free a bed for an unplanned patient (Kim et al., 2000). Lastly, the *premature discharge of a current ICU patient* (referred to as "bumping") to free up a bed exists as well. This method has seen less research attention, a complete model based on bumping was developed only in 2010 by Dobson et al. Kao & Tung (1981) note in their research that a persistent difficulty in this field of study is that the occurrences of these negative consequences often are ad-hoc decisions that do not get recorded.

In their review of simulation modelling in healthcare, Günal & Pidd (2010) note that there is a lack of generalizability in the application of models since they are generally characterized by facility- and unit specific application. Nevertheless, the broad research done in this area lays the groundwork for future models to improve on earlier work and contribute a unique solution to literature.

2.2.3 *Variable approximation – length of stay*

The LOS can be interpreted as the throughput time of any hospital ward (Yoon et al., 2003). Weissman (1997) completed a thorough statistical study of LOS data on patient admissions over six years at a surgical ICU, split into separate diagnostic groups. The main finding was that the frequency distributions are strongly skewed to the right and are comprised of a "body", interpreted as "typical behaviour" as well as a long tail of "outliers". Due to the high skewness, the mode and median are most adequate to describe typical LOS. Weissman also states that beyond removing outliers caused by erroneous data entry, accurate observations of atypical LOS (possibly caused by rare complications) might influence the analysis and their removal should be explored.
Efforts to predict LOS have been made, using demographic characteristics, comorbidity[2] and physiological variables recorded at admission as explanatory variables of LOS (Chang et al., 2002, Knaus et al., 1993, Hachesu et al., 2013, Tu & Guerriere, 1993). Chang et al. (2002) and Knaus et al. (1993) use *multivariate regression* to predict LOS and achieve an

[2] Comorbidity: the simultaneous presence of more than one medical condition

R^2 of 0.32 and 0.15 respectively. Hachesu et al. (2013) use *support vector machines* and Tu & Guerriere (1993) a *neural network* to predict which patients will experience prolonged LOS by categorizing them into previously defined categories. All authors comment that a wide range of random influences and unforeseeable medical complications negatively affect the explanatory power of their models.

The methodology of considering the LOS of patients differs among studies of ICU capacity. Early studies use average LOS values in their models (Kao & Tung, 1981), however, research by Ridge et al. (1998) confirmed that using averages is mathematically not correct due to the high variation of the LOS. They suggest the application of classification trees to group patients with the aim to reduce within group LOS variation. This was put into practice by Ridley et al. (1998) as well as Costa et al. (2003). This method is problematic, however, since it is a retrospective classification and thus cannot readily be used to classify incoming patients in the day-to-day ICU environment. Another possibility of grouping patients to reduce the variability of the LOS was introduced by Griffiths et al. (2005), who differentiated by the source of the patient and found it effective in case patients from different sources have consistently different profiles.

To accurately account for LOS distributions per patient type when modelling an ICU, Harper & Shahani (2002) and Harper (2002) fit statistical distributions to the LOS of patient groups and find that the Weibull, lognormal, gamma or negative exponential distributions are suited to represent LOS frequencies. Insufficient LOS data or highly irregular LOS frequencies that cannot be described by a statistical distribution are usually incorporated by sampling from the actual historical distribution (Griffiths et al., 2005).

2.2.4 *Variable approximation – arrival rate*

Early time series analysis of patient arrivals at a US hospital done by Swartzman in 1970 found that scheduled arrivals, which are deterministic, can be described by a normal distribution around their scheduled time and unscheduled arrivals are Poisson with significant hourly differences. More recent studies model emergency arrivals as Poisson (Kim et al., 2000) as well and research by Green (2006) confirms that approximation using a Poisson arrival distribution is valid since patient arrivals behave similarly compared to other random arrivals observed in retail, call centres or emergency services. However, it is emphasized by Green (2006) that one of the defining characteristics of the Poisson process, that arrival probability is independent of time, does not hold in the hospital environment. Admissions show seasonal and daily variations and thus require shorter time intervals or a range of admission rates to be approximated as Poisson (Green, 2006). Costa et al. (2003) rely on actual arrival patterns per time of day, day and month of elective and emergency patients in their mathematical model to avoid this pitfall. Griffiths et al. (2005) incorporate arrival rates dependent on source and time, while Kao & Tung (1981) rely on forecasts based on monthly admission data to determine hospital arrivals. Harrison et. al. (2005) improve on earlier studies by incorporating a Poisson parameter that varies per weekday. Valid representation of patient arrival rates will thus require an analysis of seasonality to determine the correct interval at which the rate should be varied.

3. METHODOLOGY

3.1 Approach

In order to answer the posed research question, a heuristic needs to be developed that enables the ICU to better manage its capacity. As established in the literature review, the ICU's operations are defined by two uncertain variables, namely the patient arrival rate and the patient LOS. Accurately representing these two inputs of the model will ensure the validity of the heuristic. Key to a successful representation of these variables is a sound categorization of patients to reduce within-group variability. Thus, the categorization of patients into different types is the first step of this analysis and forms the basis for the definition and estimation of both the patient arrival rate and LOS. Next, the two input variables need to be approximated. The estimation of patient arrival rates needs to reflect the seasonality that has been observed in studies of hospital admissions. The long-tailed nature of LOS distributions needs to be represented due to the high variation and the detrimental impact of LOS outliers on bed availability. In parallel to the estimation of the variables, the historical data of the ICU is analysed to gain an understanding of the current situation, which will deliver important inputs for the development of the improvement heuristic. With the approximations and in place, the heuristic will be developed and optimized considering any restrictions that are in place in the form of constraints. Lastly, a simulation model is designed, and the effect of the heuristic on ICU performance is tested.

3.2 Data Collection and Tools

The analysis and models of this study are based on an export of all patient admissions at the ICU from the 1st of January 2012 until the 31st of December 2015 from the ICU's patient management system. Table 3-1 summarizes the fields and descriptions of the export.

Field Title	Description	Field Type
ID	A unique ID of each patient admission	String
patID	An ID unique to each patient	String
birthdate	The birthdate of the patient	Numeric (Date)
gender	The gender of the patient	String (Categorical)
admitWT	The weight of the patient	Numeric
height	The height of the patient	Numeric
careUnitLabel	The identifier of the care unit the patient is in (ICCU/ICTH/MC/HC)	String (Categorical)
beginUnitTime	The moment of arrival at the unit identified in "careUnitLabel"	Numeric (Date)
endUnitTime	The moment of discharge from the unit identified in "careUnitLabel"	Numeric (Date)
PatientAdmitTime	The moment of the admission to the hospital (can be equal to "beginUnitTime" if the ICCU/ICTH is the first unit of admission	Numeric (Date)
PatientDischargeTime	The moment of the patient's discharge from the hospital (can be equal to "endUnitTime" if the ICCU/ICTH is the unit that discharges the patient	Numeric (Date)
diagnosistime	The moment the patient was diagnosed (only applies to ICCU patients)	Numeric (Date)
diagnosis	The diagnosis of the patient's condition (only applies to ICCU patients)	String
operationtime	The moment of the patient's operation (only applies to ICTH patients)	Numeric (Date)
operation	The kind of operation executed (only applies to ICTH patients)	String

Table 3-1: Description of dataset fields

The dataset contains all admissions to both the ICTH and ICCU unit as well as incomplete data on Medium Care and High Care admissions, in total 16,442 admissions were recorded in these four years. The data was checked for completeness and consistency. It was noted that there were a large number of duplicates, caused by the fact that changing the content of the *"diagnosis"* or *"operation"* fields duplicated the entire entry. These changes were examined and only consisted of minor adaptions, with each original entry containing the most relevant information. Since a unique patient can be re-admitted and thus be present multiple times in the dataset, a new variable that identifies unique visits was created by concatenating *"patID"* and *"PatientAdmitTime".* After removing Medium Care and High Care admissions as well as filtering the dataset for unique visits, the number of relevant admissions to the ICCU and ICTH turned out to be 12,170. Apart from these adjustments, only minimal cleaning was necessary, since the export was of good quality.

To create the basis of the simulation model, a thorough analysis of the raw data was required. The initial statistical analysis, data preparation and patient classification were accomplished using the spreadsheet software Microsoft Excel 2016 (v. 16.15) as well as the R based analytical package Alteryx Designer (v. 2018.2). The seasonal analysis was done using R (v. 3.5.0), distribution fitting was done using the MATLAB package (v. R2018a).

3.3 Patient Classification

A consensus in existing literature on ICU capacity management is that modelling the ICU environment requires a classification of patients into different categories. In the academic literature after-the-fact classification in the form of classification trees is a common approach to reduce LOS variability (Ridge, 1998, Costa, 2003), while others (Weissman, 1997, Griffiths et al., 2005) use a readily observable characteristic (the diagnosis/source of the patient) as the basis for patient classification. In the spirit of arriving at an actionable recommendation, a patient classification of the easily observed second kind is preferred for this research. During discussions with the director of the Erasmus MC's ICU unit, it was established that in their daily operations they categorize patient type by diagnosis. By deciding to adopt the same categorization in this research the interpretation and implementation of the results for the ICU staff are facilitated (refer back to Table A-1 & Table A- A-2 for an overview of all patient types). This approach is in line with the aim of the categorization to reduce within-group variability since both academic literature (Ridge et al., 1998) and discussions with the ICU staff confirmed that each disease type has a certain LOS profile and differing arrival rate. The patient classification was not in place in the dataset and was derived from the string fields *"diagnosis"* (for ICCU patients) and *"operation"* (for ICTH patients) that contain a description of the condition of and the procedures executed on each patient. The data science software package "Alteryx Designer" was used to execute the text mining required to allocate all patients to types. A detailed description of the process is given in Appendix B, and visualisation of its implementation in Figure B-1.

A brief overview of the workflow is as follows: The first step was mining all fields for common phrases and keywords, this was done by building a workflow that used the "fuzzy match" technique, which accepts less than perfect matches, to compare all strings among each other and deriving the most common phrases and keywords. These were sorted by frequency of occurrence and discussed with the ICU director. Together irrelevant keywords were eliminated and all relevant keywords were connected to their corresponding patient type. After compiling all keywords into a helper sheet, a workflow was built to find and match patient type according to the information contained in the *"diagnosis"* / *"operation"* field. Finally, the occasions where more than one patient type was matched were examined and eliminated by refining the keyword list and re-running the workflow until a coherent classification was in place. The resulting classification was then checked against the yearly estimates of yearly patients per type by the Erasmus MC.

3.4 Estimation of Variables & Statistical analysis

3.4.1 *Length of stay*

The LOS of a patient represents the time a patient occupies a bed in the hospital system. Since the study is focused on the ICU unit only it is defined as:

$$LOS = endUnitTime - beginUnitTime$$

It was calculated for all ICCU/ICTH patients and the frequency distributions per patient type were compiled. For an initial understanding of the distributions, summary statistics were calculated that include the median, mode, coefficient of variance as well as the range in percentiles. As stated by Weissman (1997), the distribution of recovery time of patients tends to be distorted by extreme outliers. The statistical summary of LOS data was discussed with the ICU director and it was determined that indeed each LOS distribution contained extremely long LOS outliers, which, even though they might have been accurate recordings, were not representative of the patient type. It was thus decided to remove the top 5% observations of each patient type. With the historical LOS frequencies in place, the distribution fitting function was coded. All LOS distributions were fitted with the Lognormal, Weibull and Gamma distributions, which were found to be a good fit for LOS data by Harper (2002), as well as the Log-logistic distribution, which is an additional distribution that is well suited for fitting right-tailed distributions. Following the intuition of the maximum likelihood estimation, the distribution that maximized the log-likelihood score was chosen for each patient type.

3.4.2 *Arrival rate*

To examine overall patient arrival behaviour all arrivals were examined together at first. Counts of patient arrivals per year, month, weekday and hour were compiled using the "*beginUnitTime*" date and time field. Arrivals were then calculated by averaging the arrivals per year, month, weekday and hour across the years of available data. After studying the seasonality revealed by this initial analysis, patient arrivals were split into planned and unplanned, as well as per patient type and the calculation was repeated.

As established during the literature review, modelling unplanned arrivals as Poisson has been proven to be effective, however, it is required to observe shorter time intervals, since patient arrivals are not independent of time due to seasonality (Green, 2006). The planning horizon of the ICU is one week ahead, furthermore, since there is evidence of daily variations, the hourly level of arrivals should be considered. To make the results more practical the timing of the three nursing shifts at the ICU (8:00 am - 4:00 pm, 4:00 pm - 12:00 am and 12:00 am - 8:00 am) are used as a subdivision of each day when studying the historical arrivals, which will allow the ICU to staff according to the results more easily. The simulation will use a more detailed hourly arrival rate, thus, for each unplanned patient type, 24 Poisson arrival parameters for each weekday are defined.

Planned patients arrive according to elective surgery demand and are subject to a schedule laid out by the surgery planner, two factors which, in the current setup, are not influenced by the ICU. While the final planned patient schedule is known about a week in advance, there is an element of uncertainty with regard to the exact point in time at which an elective surgery patient arrives. The first surgery starts at 8 am in the morning and is followed by the next in line once it is completed. Each kind of surgery has a different and variable duration, a fact that introduces variability into the exact time of a planned patient arrival on a certain day. Since the creation of the surgery schedule is outside of the scope of this project and existing schedules do not go far into the future, planned arrivals need to be approximated. This is done by determining the average number of planned patients per surgery kind per day and simulating the sequence of their arrival by sampling the surgery duration from its known range, with the first surgery starting at 8 am and the following arrivals queuing until a surgery slot is available.

3.5 **Heuristic Development & Optimization**

The estimation of patient arrival rates and LOS goes hand in hand with a thorough statistical analysis of the available historical data of the ICU performance. This understanding of capacity use at the Erasmus MC as well as of its current practices is combined with approaches from past ICU studies to develop a heuristic that has the potential to improve ICU performance. The resulting heuristics will then be optimized. The basis of the optimization model is created by defining an objective function as well as parameters, decision variables and constraints. The heuristic development is detailed in chapter 5.

3.6 **Simulation Model & Validation**

Before defining the model for the simulation, restrictions that were discussed together with the ICU staff are defined and additional assumptions that are necessary for the model are detailed. Next, a detailed simulation of the Erasmus MC's ICU will be built that includes the previously described classification and estimations of the uncertain variables to realistically account for the ICU environment. As stated in the literature review, discrete event simulation presents itself as a good fit to model the dynamics of an intensive care unit and has been successfully applied in previous studies (Günal & Pidd, 2010). After testing the validity of the resulting simulation model, the effect of the bed capacity management heuristic can be tested. It will be run multiple times to improve validity and to arrive at a concise recommendation. The development of the simulation model and the validation is further discussed in chapter 6.

3.7 Recommendations & Conclusion

The findings and interpretation are summarized to arrive at a concise recommendation for the ICU. Before concluding the research, an outlook of future steps, as well as a suggestion for further research, are given.

4. STATISTICAL ANALYSIS

4.1 Patient Population

The average age of an ICU patient was 63.71 years, the gender split was 65% male, 35% female. While all patients were successfully classified into patient types, four groups of patients were excluded from further analysis: For 696 admissions the *"diagnosis"* / *"operation"* field was NULL, so these entries were eliminated. A patient type defined only as "recovery" was a catch-all for unspecified post-surgery patients and matched to 280 admissions. The patient types Electrophysiology, "Longchirurgie" and "Kinder-hartchirurgie" were matched to 130, 623 and 632 patients respectively. Discussing the result with the ICU director revealed that the ICU will not care for these four patient groups anymore after the move, so they were excluded as well, resulting in a total of 10,772 remaining matches.

Table 4-1 contains an overview of the frequency of each patient type and Table 4-2 quantifies the split between planned and unplanned patients. Studying the results of the patient classification reveals that the three largest patient types, namely STEMI, Overige and CABG make up about half of all patients (45.75%) and consequently require special attention. On the other hand, the eight smallest patient types account for only 3.87% of all patients admitted over four years and do not have a significant influence on the ICU performance. Out of the top three, CABG is the only one consisting of planned patients and is notably less frequent compared to STEMI and Overige. Consequently, the split of arrival kind in Table 4-2 confirms that the large majority of patients cared for in the ICU are unplanned, a factor that introduces significant variability at the ICU.

Patient Type	Count	Percentage	Arrival Kind
STEMI	1931	17.93%	Unplanned
Overige	1875	17.41%	Unplanned
CABG	1122	10.42%	Planned
Ritmestoornissen	904	8.39%	Unplanned
Aorta	794	7.37%	Planned
Klep	624	5.79%	Planned
NSTEMI	590	5.48%	Unplanned
OHCA	557	5.17%	Unplanned
Hartfalen	548	5.09%	Unplanned
TAVI	538	4.99%	Planned
Cardiology	316	2.93%	Planned
Pericarditis	218	2.02%	Unplanned
Overige TH	214	1.99%	Planned
Congenital	124	1.15%	Planned
Endocarditis	86	0.80%	Unplanned
Thoracotomie	69	0.64%	Planned
Re-Thoracotomie	63	0.58%	Planned
HTX	53	0.49%	Planned
LOTX	52	0.48%	Planned
Tamponade	46	0.43%	Unplanned
LVAD	30	0.28%	Planned
ECMO	18	0.17%	Unplanned

Table 4-1: Frequency table of patient types

Arrival Kind	Share
Planned	37.12%
Unplanned	62.88%

Table 4-2: Share of patient arrival types

4.2 Length of Stay

The service time of the intensive care unit is represented by a patient's length of stay. Previous research has already established that LOS is a highly variable measure, characterized by a long tail distribution due to it being influenced by unique factors, such as comorbidity, type of disease and the presence of complications. Realistic modelling of the ICU environment requires a solid understanding and incorporation of patient LOS, for which the patient classification described in the methodology section forms the basis.

4.2.1 *LOS per type*

Each patient type is characterized by a unique distribution of LOS. Table 4-3 gives a statistical overview of LOS per patient type in hours. Since there is a consensus among researchers that the simple average is not representative (Ridge et al., 1998), the listed median gives a more accurate impression of the typical LOS of each patient type. To give a perspective on the relative size of the variation, the coefficient of variation is included, defined as $CoV = \frac{SD}{Mean}$, with values less than 0,75 representing low variability and values above 1,3 representing high variability. Furthermore, the listed percentiles, as well as the listed maximum give an impression of the high outliers that can occur in any patient category. Studying the LOS characteristics of each patient types reveals a number of interesting insights, especially in connection with the number of yearly patients of certain types. Patients of both types typically tend to leave the ICU within the day, overall 73.26% of patients are discharged within 24 hours. It is notable that the three largest unplanned patient types (STEMI, Overige and Ritmestoornissen) are highly variable but are characterised by a rather short LOS. Potential blockage can come from OHCA and Hartfalen patients, whose LOS is not highly variable but is around two and three days respectively with significant outliers. These two groups make up 10.17% of all patient arrivals. The most common planned patient types (CABG, TAVI, Aorta & Klep) exhibit lower variability and are typically discharged in less than 24 hours. Patients of the transplant surgeries (HTX & LOTX) as well as LVAD patients come in small numbers but are certain to stay more than two days with some of these patients staying longer than a week. The Cardiology and Re-Thoracotomie patients show high outliers compared to their median values and consequently have a higher variability.

Patient Type	Mean	Median	SD	CoV	75^{th}P	95^{th}P	Max
ECMO	194	184	144	0.74	252	397	436
LOTX*	113	100	65	0.57	129	172	240
OHCA	96	86	68	0.71	135	196	235
HTX*	107	83	76	0.71	125	235	287
LVAD*	95	67	88	0.93	118	173	258
Hartfalen	68	44	69	1.03	94	163	222
Endocarditis	52	44	47	0.90	83	110	160
TAVI*	25	22	22	0.87	24	46	73
Re-Thoracotomie*	38	22	43	1.12	49	96	117
Aorta*	26	21	22	0.83	24	51	79
Pericarditis	29	21	28	0.96	41	69	96
CABG*	19	20	10	0.51	22	25	43
KLEP*	19	20	12	0.65	22	25	44
Tamponade	33	19	49	1.48	28	82	115
Cardiology*	31	19	39	1.27	39	83	115
Thoracotomie *	20	18	23	1.14	23	59	68
Congenital*	13	14	9	0.63	22	23	25
NSTEMI	18	10	17	0.96	25	44	53
Overige	14	7	15	1.09	21	38	48
STEMI	11	6	14	1.28	10	26	46
Overige TH*	12	5	14	1.21	20	25	37
Ritmestoornissen	10	4	14	1.37	12	24	43

Table 4-3: LOS descriptives per patient type (Planned patient types are denoted by the asterisk)

Visual examination of the histograms of LOS per patient type (Figure D-1 in the appendix) clearly shows the long right-tailed nature of the respective distributions. Additionally, a clear discharge pattern can be recognised by the peaks in the distribution, most patients get discharged within the day of admission or 24h or 48h later.

4.2.2 *ABC analysis*

How the difference in LOS among patients impacts the capacity of the ICU can be studied by applying the ABC analysis method. Usually applied to classify inventory by popularity, it turns out to be well suited to examine resource usage among patients. Patients were sorted by LOS and split into three classes. Class C contains all patients that stay less than one full day ($LOS < 20h$), class B patients that stay between one and two days ($20h <$

$LOS < 48h$) and class A all patients that stay longer than two days ($48h < LOS$). Total capacity use was calculated by summing all individual LOS and then the cumulative share of capacity usage per patient was calculated and plotted on a Pareto curve (see Figure 4-1).

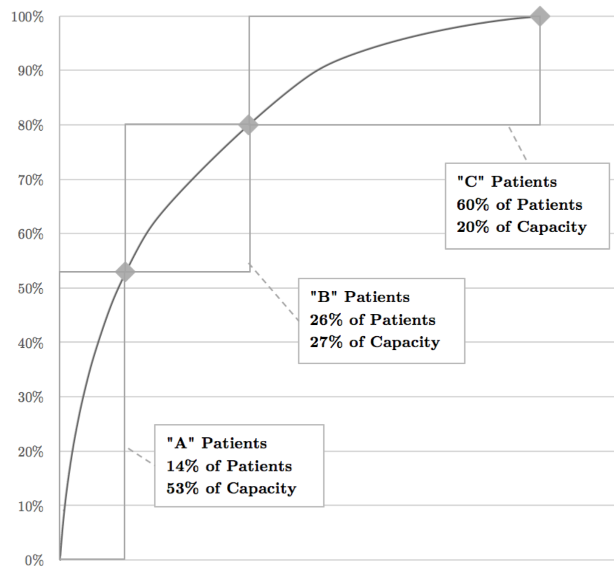

Figure 4-1: LOS Pareto curve

Even after removing outliers as described in the methodology section, the Pareto curve shows that 53% of capacity (measured in bed-days) is occupied by class A, which makes up only 14% of all patients. Class B patients make up 27% of all patient admissions and occupied 30% of used capacity. Finally, class C, the largest class, accounts for 60% of all patients and 20% of all used capacity. In order to get a clearer and a more aggregated picture of ICU capacity use, it is helpful to return to the patient type classification, which reveals a number of relevant insights. When studying the share of arrivals and capacity usage per type in more detail (Table 4-4, a visualisation is provided by Figure 4-2) one can easily recognise which patient groups tend to occupy most of the ICU capacity. As hypothesized after examining the LOS descriptives, OHCA and Hartfalen patients, two rather small classes, occupy a large share of ICU capacity, in fact, they have the highest share of ICU capacity usage and thus deserve special attention. OHCA patients account for only 5.17% of all patients but occupy 19.22% of ICU capacity, their share of capacity use is 3.71 times higher than their share of arrivals. The second largest share of capacity (13.43%) is occupied by the Hartfalen type, which also accounts for a rather small share of total patients (5.09%, resulting in a capacity use versus arrival ratio of 2.64). These two groups are likely to occupy ICU beds for a longer period of time and block capacity, especially if there is more than one admitted patient of these types. The four largest patient groups (STEMI, Overige, CABG, Ritmestoornissen) all account for a large share of patient arrivals but occupy comparatively less of the ICU's capacity (Ratios of 0.45, 0.58, 0.76 and 0.41). The ten largest patient classes account for 82.85% of the ICU capacity usage, beyond these share of capacity use of each patient type drops quickly, mostly due to the small number of patients in these groups.

Patient Type	**% of Capacity use**	**% of Patients**	**Ratio**
STEMI	7.99%	17.93%	0.45
Overige	10.04%	17.41%	0.58
CABG*	7.90%	10.42%	0.76
Ritmestoornissen	3.43%	8.39%	0.41
Aorta*	7.54%	7.37%	1.02
Klep*	4.34%	5.79%	0.75
NSTEMI	3.88%	5.48%	0.71
OHCA	19.22%	5.17%	3.72
Hartfalen	13.43%	5.09%	2.64
TAVI*	5.08%	4.99%	1.02
Cardiology*	3.36%	2.93%	1.15
Pericarditis	2.22%	2.02%	1.10
Overige TH*	0.96%	1.99%	0.48
Congenital*	0.62%	1.15%	0.54
Endocarditis	1.75%	0.80%	2.19
Thoracotomie*	0.52%	0.64%	0.81
Re-Thoracotomie*	0.87%	0.58%	1.49
HTX*	1.96%	0.49%	3.98
LOTX*	2.12%	0.48%	4.38
Tamponade	0.52%	0.43%	1.22
LVAD*	0.95%	0.28%	3.42
ECMO	1.31%	0.17%	7.84

Table 4-4: Capacity usage per patient type

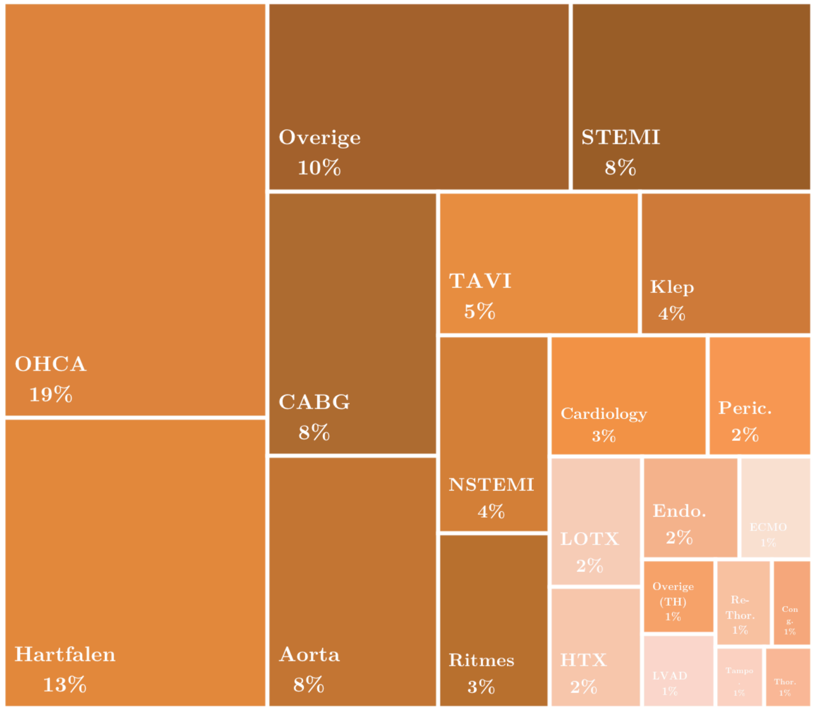

Figure 4-2: Capacity usage visualisation

4.2.3 *Distribution fitting*

Realistic representation of patient LOS in the simulation model is ensured by fitting a probability distribution to each patient type, from which individual LOS values can be drawn to generate the service time per patient in the model. As established in the methodology section all historic LOS distributions can be approximated by fitting a range of positive right tail distributions (Harper, 2002). In this study historical data was fitted with the Lognormal, Weibull, Gamma and Log-logistic distributions using MATLAB. The lowest Log-likelihood score was chosen as the criterion to determine the best fit. The results are displayed in Table 4-5 below, the histograms with the fitted distribution are displayed in Figure D-2 in the appendix.

Patient Type	Distribution	LogL Score	Parameters
STEMI	loglogistic	-7411.00	γ= 1.924 σ= 0.4624
Overige	lognormal	-7702.00	μ= 2.23 σ= 1.051
CABG*	weibull	-4174.00	α= 22.06 β= 2.079
Ritmestoornissen	lognormal	-3670.00	μ= 1.767 σ= 1.014
Aorta*	loglogistic	-3316.00	γ= 3.033 σ= 0.4019
Klep*	gamma	-2379.00	α= 2.578 β= 7.449
NSTEMI	lognormal	-2819.00	μ= 2.511 σ= 0.9334
OHCA	weibull	-3690.00	α= 105 β= 1.343
Hartfalen	weibull	-3415.00	α= 71.59 β= 0.9648
TAVI*	loglogistic	-2859.00	γ= 3.045 σ= 0.3781
Cardiology*	lognormal	-1731.00	μ= 2.787 σ= 1.192
Pericarditis	gamma	-1654.00	α= 1.665 β= 20.38
Overige TH*	lognormal	-737.20	μ= 1.929 σ= 1.102
Congenital*	weibull	-439.70	α= 15.63 β= 1.646
Endocarditis	gamma	-568.70	α= 0.8522 β= 69.89
Thoracotomie*	lognormal	-279.90	μ= 2.459 σ= 1.129
Re-Thoracotomie*	weibull	-293.40	α= 37.96 β= 0.955
HTX*	loglogistic	-298.50	γ= 4.476 σ= 0.391
LOTX*	loglogistic	-277.90	γ= 4.608 σ= 0.3161
Tamponade	lognormal	-377.70	μ= 3.166 σ= 1.378
LVAD*	gamma	-166.00	α= 1.367 β= 70.02
ECMO	uniform	-126.20	min 20.42 max 428

Table 4-5: Distribution fitting results

4.3 Patient Arrivals

To develop an understanding of patient arrivals over time at the ICU, as well as to determine trends and seasonality, arrivals per arrival type as well as per patient type are examined across different time horizons before defining arrival rates that serve as input for the model.

4.3.1 *Yearly patient arrivals*

An overview of yearly patient arrivals at the ICU is given in Table 4-6. One can observe a slight increase of about 4.5% in overall patient numbers over the years, caused by increases in both planned and unplanned patient arrivals. The steady growth observes confirms the assumption of increasing patient volume, at this level, however, it does not need to be further considered in the simulation model.

	2012	2013	2014	2015
Unplanned	1,653	1,661	1,752	1,707
Planned	970	1,003	985	1,041
Grand Total	2,623	2,664	2,737	2,748
Unplanned per day	4.53	4.55	4.80	4.68
Planned per day	2.66	2.75	2.70	2.85
Total per day	7.19	7.30	7.50	7.53

Table 4-6: Patient arrivals per kind per year

Breaking down arrivals according to patient type (Table 4-7) reveals two recent noteworthy changes, in 2015, TAVI arrivals, a scheduled surgery, jumped up by 34%, while the catch-all category Overige saw a 9% drop. Since the Overige category is a catch-all of various patients the drop could not be explained by the ICU. The rise in TAVI patients was due to increased surgeon availability, which led the Erasmus MC to offer more TAVI operations. Otherwise, admissions per patient type increase slightly over the observed four-year period.

Patient Type	2012	2013	2014	2015
STEMI	545	447	461	478
Overige	476	487	478	434
CABG*	296	297	256	273
Ritmestoornissen	146	221	283	254
Aorta*	175	176	222	221
Klep*	164	176	145	139
NSTEMI	125	138	170	157
OHCA	140	144	130	143
Hartfalen	125	130	138	155
TAVI*	103	133	127	175
Cardiology*	77	76	78	85
Pericarditis	62	55	51	50
Overige TH*	58	50	52	54
Congenital*	29	27	32	36
Endocarditis	16	22	28	20
Thoracotomie*	20	22	16	11
Re-Thoracotomie*	14	19	18	12
HTX*	12	9	18	14
LOTX*	13	14	14	11
Tamponade	14	14	9	9
LVAD*	9	4	7	10
ECMO	4	3	4	7

Table 4-7: Patient arrivals per patient type per year

4.3.2 *Monthly patient arrivals*

Drilling further down and examining aggregated average patient arrivals per month reveals a slight variation of arrivals between the summer and winter months at first sight (Figure 4-3), more noticeably for unplanned patients.

To confirm the presence of seasonality, the data is converted into a time series model and decomposed into its trend, seasonal and error components using the stl function of the integrated “stats” package in R (R Core Team, 2018). As can be seen in Figure 4-4, a yearly seasonal pattern is recognized. Peaks in January, March and October highlight an apparent difference in patient arrivals in winter compared to the summer months, with the lowest level being reached in August.

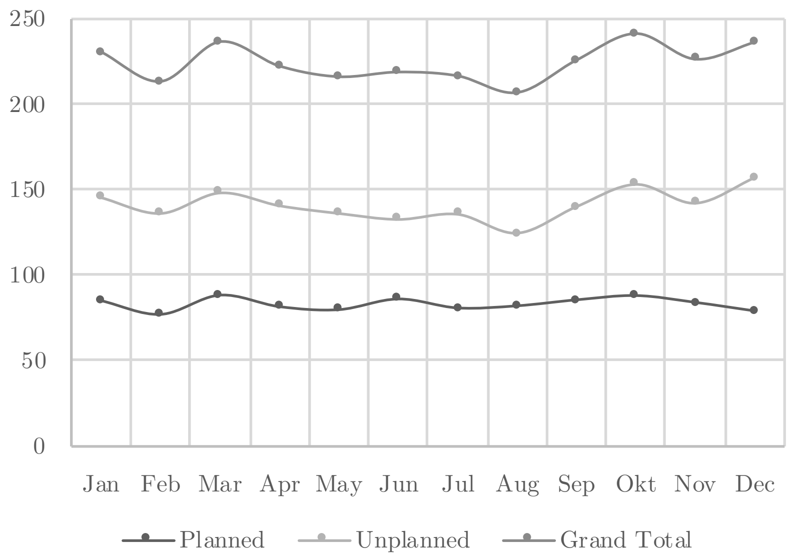

Figure 4-3: Patient arrivals per kind per month

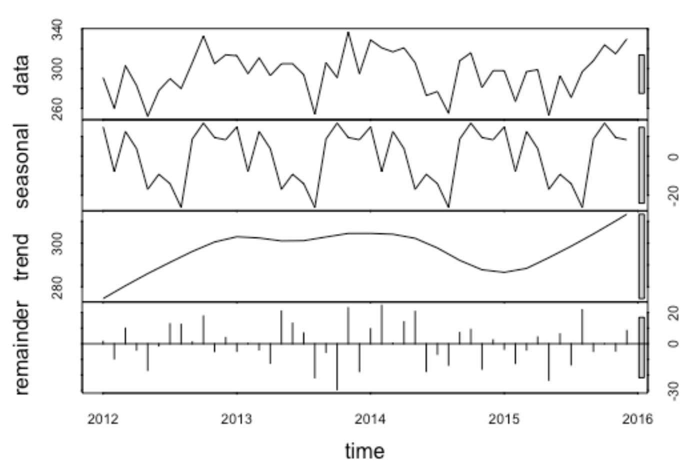

Figure 4-4: Seasonal decomposition

Discussing this finding with the ICU revealed that any seasonal differences of patient arrivals are most likely due to staffing at the ICU, the dip in the summer months caused by reduced staffing in the due to holidays. Thus, the difference in arrivals is not inherent to any patient type or characteristic and will not be further considered.

4.3.3 *Daily patient arrivals*

Inspecting patient arrivals per weekday reveals a clear pattern which is depicted in Figure 4-5. Due to no elective surgeries being scheduled on the weekend there is a clear drop in overall arrivals at the end of the week. Erasmus MC has an agreement in place with Maasstad hospital to balance the workload of emergency patients, on Mondays and Thursdays the majority of STEMI and OHCA patients are received by the Maasstad hospital, which is why a clear drop of unplanned patient arrivals can be observed.

Patient Type	**Mon**	**Tue**	**Wed**	**Thu**	**Fri**	**Sat**	**Sun**
STEMI	0.64	1.79	2.13	0.72	1.51	1.28	1.21
Overige	1.25	2.01	1.62	1.11	1.66	0.70	0.66
CABG*	1.23	1.21	0.91	0.95	1.00	0.05	0.03
Ritmestoornissen	1.18	0.35	0.92	0.40	1.04	0.22	0.23
Aorta*	0.79	0.63	0.73	0.85	0.69	0.06	0.07
Klep*	0.14	0.69	0.12	1.37	0.18	0.05	0.04
NSTEMI	0.35	0.50	0.57	0.34	0.51	0.28	0.29
OHCA	0.42	0.39	0.43	0.39	0.47	0.24	0.29
Hartfalen	0.25	0.53	0.52	0.31	0.39	0.31	0.38
TAVI*	0.59	0.64	0.56	0.56	0.61	0.02	0.01
Cardiology*	0.33	0.34	0.20	0.25	0.22	0.10	0.10
Pericarditis	0.22	0.21	0.20	0.18	0.16	0.04	0.02
Overige TH*	0.14	0.12	0.24	0.17	0.20	0.07	0.11
Congenital*	0.13	0.11	0.17	0.09	0.10	0.00	0.00
Endocarditis	0.08	0.07	0.05	0.07	0.09	0.02	0.03
Thoracotomie*	0.02	0.04	0.04	0.03	0.05	0.02	0.04
Re-Thoracotomie*	0.05	0.04	0.05	0.07	0.04	0.03	0.02
HTX*	0.05	0.08	0.06	0.07	0.05	0.01	0.01
LOTX*	0.04	0.03	0.06	0.05	0.01	0.05	0.01
Tamponade	0.03	0.01	0.06	0.02	0.01	0.00	0.00
LVAD*	0.05	0.03	0.06	0.03	0.02	0.00	0.02
ECMO	0.01	0.03	0.01	0.01	0.01	0.01	0.00

Table 4-8: Patient arrivals per patient type per weekday

Studying aggregated average arrivals per patient type per weekday more closely (Table 4-8) reveals further information about the planning currently done at the Erasmus MC. Apart from not being scheduled on the weekends, some planned surgeries have a planning pattern: CABG surgeries are scheduled more often at the beginning of the week and TAVI surgeries mostly on Tuesday and Thursday, while the others are evenly spread out over the week. The presence of seasonality was confirmed through seasonal decomposition in R. Due to the systematic nature of planned surgery variation, it needs to be taken into account when creating the simulation of the ICU environment.

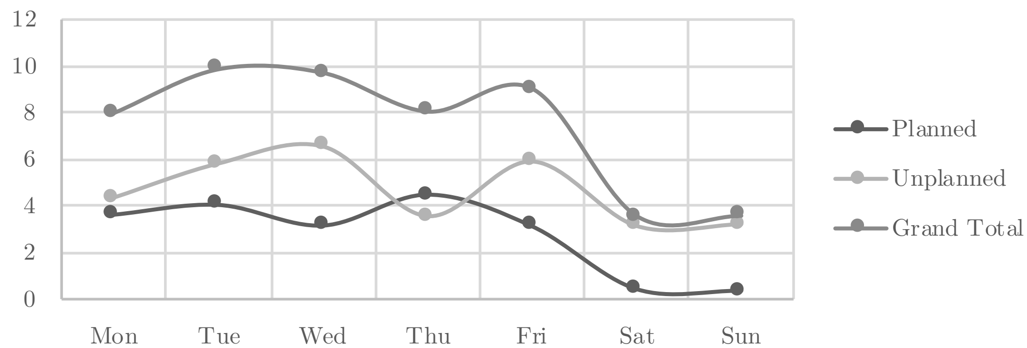

Figure 4-5: Patient admissions per kind per weekday

4.3.4 *Hourly patient arrivals*

On an hourly level, the data seasonality is present as well when aggregating average patient arrivals (Figure 4-6). After 8 am arrivals rapidly increase to reach their peak at 2 pm. After 5 pm they rapidly drop and stay relatively low during nighttime (9 pm – 8 am). Planned arrivals ramp up later, which can be intuitively explained by the duration of the surgeries that patients undergo before admission to the ICU, with the first one usually starting at 8 am.

The rapid increase in arrivals in the span of six hours that can be observed from 8 am to 2 pm demonstrates the high level of daily variation of patient arrivals. It is thus essential to account for it in the simulation model. As laid out in the methodology section, segmenting the day into the three different nursing shifts (Shift 1: 12 am – 8 am, shift 2: 8 am – 4 pm, shift 3: 4 pm – 12 pm) aides the relevancy, as well as the ease of understanding and implementation of the results for the ICU staff. The breakdown of patient arrivals aggregated into the three shifts is displayed in Table 4-9, showing especially the difference between planned and unplanned in more detail.

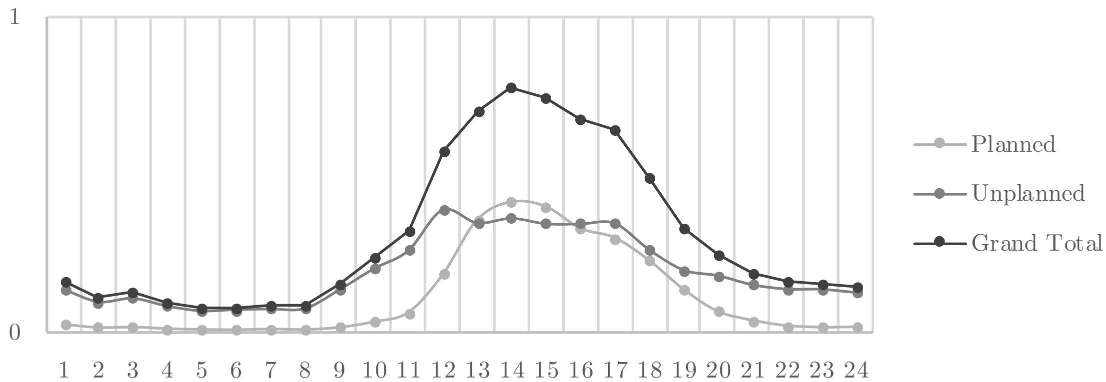

Figure 4-6: Patient arrivals per kind per hour

Patient Type	Shift 1	Shift 2	Shift 3
STEMI	0.2603	0.6137	0.4486
Overige	0.1774	0.7466	0.3603
CABG*	0.0082	0.4842	0.2760
Ritmestoornissen	0.0514	0.3925	0.1753
Aorta*	0.0199	0.3432	0.1808
Klep*	0.0048	0.3178	0.1048
NSTEMI	0.0541	0.1932	0.1568
OHCA	0.0705	0.1534	0.1575
Hartfalen	0.0753	0.1658	0.1342
TAVI*	0.0123	0.2836	0.0726
Cardiology*	0.0342	0.0986	0.0836
Pericarditis	0.0158	0.0719	0.0616
Overige TH*	0.0041	0.0959	0.0466
Congenital*	0.0000	0.0774	0.0075
Endocarditis	0.0055	0.0253	0.0281
Thoracotomie*	0.0007	0.0322	0.0144
Re-Thoracotomie*	0.0048	0.0199	0.0185
HTX*	0.0110	0.0164	0.0089
LOTX*	0.0096	0.0151	0.0110
Tamponade	0.0055	0.0158	0.0103
LVAD*	0.0007	0.0151	0.0048
ECMO	0.0014	0.0048	0.0062

Table 4-9: Patient arrivals per patient type per shift

4.3.5 *Combined seasonality*

Close examination of the seasonality patterns at different time levels revealed a number of factors that result in unique arrival patterns. At an aggregated level, patient arrivals exhibit within-week variation across weekdays and daily variation across the previously identified shifts. Figure 4-7 captures the two relevant levels of the time-varying seasonality of planned and unplanned patients in one graph, the inherent pattern is clearly visible.

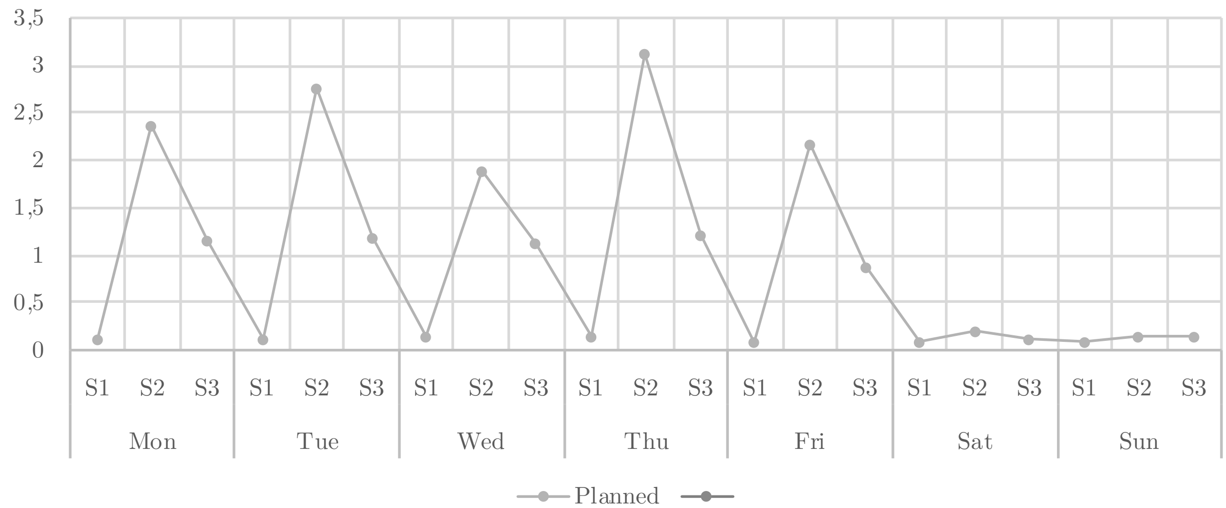

Figure 4-7: Seasonality pattern overview per patient kind

5. HEURISTIC DEVELOPMENT

Much of the academic literature focusses on improving ICU performance by determining an adequate ward size beforehand (Green, 2002) or by applying capacity related adjustments, such as flexible bed allocation (Kim et al., 2000, Kao & Tung, 1981). Since the capacity of the ICU at the Erasmus MC is fixed, the only levers available to improve the performance of the ICU are managing the uncertain variables of patient arrivals and patient LOS.

While unplanned patient arrivals can hardly be influenced, the timing of planned patient arrivals can be coordinated with the surgical planner and by moving or dispersing this flow of arrivals the overall arrival rate at the ICU can be influenced. The analysis of arrivals showed the existence of within-week seasonality for planned patients, certain kinds of elective surgery are executed on a specific day at the Erasmus MC, so there certainly is a possibility to influence their sequence. For instance, scheduling surgeries of heart or lung transplants, which are known to stay more than one day in the ICU, towards the end of the week might improve capacity utilization, since there are no new incoming planned surgeries during the weekend. On the other hand spikes of arrivals caused by unplanned patients can be smoothed out by moving the planned patient load accordingly. Therefore, the heuristic to be developed and tested in this research will be an "*Arrival smoothing heuristic*", which aims to balance ICU capacity use by distributing planned patient arrivals across the week according to their LOS characteristics as well as the levels of unplanned patient load.

5.1 **Arrival Smoothing Heuristic**

5.1.1 *Requirements*

Manipulating the influx of planned patients is one of the few levers that has a direct impact on ICU capacity usage at the disposal of the ICU. Achieving a better balance of patient load across the week will decrease the number of times the ICU fills up, which entails a reduction in the number of patients that need to be discharged prematurely. In order to consciously distribute planned patient arrivals across the days of the week, it is necessary to combine the previously analysed statistics to better quantify the planned and unplanned patient load per weekday since the patient arrival rate only is not the right proxy for ICU patient load. This is due to the implications of the second uncertain variable, LOS. The majority of patients leaves within 24 hours, while other patients (OHCA patients for example) occupy a bed for multiple days and thus have a larger impact on ICU capacity use. Since both the surgery scheduler and the ICU work with a one-week planning horizon, the end result of the heuristic will be a suggested daily surgery quota for each day of the week. As explained in the methodology section, there is no forward-looking schedule of elective surgeries. Consequently, it is necessary to return to the classification and characteristics of planned patients to define a weekly arrival pattern for planned patients before moving on to the calculation of the patient loads.

5.1.2 *Planned patient schedule development*

The detailed patient type categorization in place does not allow for a straightforward definition of a weekly elective surgery schedule. The categorization of planned patients into 12 different types results in small patient groups with low patient arrivals per type per weekday (see Table 4-8) as well as irregular patient arrivals per week (see CoV in Table 5-1). These two factors impede the definition of an integer weekly planned patient schedule per patient type, it is required to aggregate these patient types in order to arrive at a reliable weekly schedule for planned patients.

Patient Type	Average Weekly Arrivals	St. Dev.	CoV
Aorta	3.81	1.92	0.51
CABG	5.36	2.29	0.43
Cardiology	1.54	1.35	0.88
Congenital	0.59	0.88	1.48
HTX	0.27	0.54	2.02
Klep	3.00	1.73	0.58
LOTX	0.26	0.49	1.90
LVAD	0.17	0.40	2.34
Overige TH	1.04	1.12	1.08
Re-Thoracotomie	0.32	0.57	1.79
TAVI	2.60	1.63	0.63
Thoracotomie	0.34	0.64	1.87

Table 5-1: Planned patient arrivals per week

As a reminder, the categorization of patients served to reduce the within-group variation of LOS, the uncertain service time of patients in the ICU. This aim needs to guide the aggregation of patient types as well, arrivals may only be merged if the patient types exhibit the same LOS characteristics.

Patient Type	Average	Median	Max	StDev	CoV	Group
LOTX	113	100	342	64.88	0.57	1
HTX	107	83	336	75.89	0.71	
LVAD	95	67	410	88.40	0.93	
Re-Thoracotomie	38	22	228	42.99	1.12	2
Cardiology	31	19	207	39.15	1.27	
TAVI	25	22	139	21.56	0.87	3
Aorta	26	21	119	21.59	0.83	
Thoracotomie	20	18	108	23.23	1.14	4
CABG	19	20	66	9.80	0.51	
Klep	19	20	72	12.26	0.65	
Congenital	13	14	27	8.56	0.63	5
Overige TH	12	5	76	14.33	1.21	

Table 5-2: Planned patients LOS statistics

Inspecting the LOS statistics of all planned patients indicates that five subgroups seem to have comparable LOS characteristics (Table 5-2). This observation needs to be validated, the appropriate test to establish if two populations are comparable is the t-test. While the standard students' t-test assumes equal variances and an underlying normal distribution of the population, research by Fagerland & Sandvik (2009) has shown that the modified Welch-t-test, which assumes unequal variances, is suited for skewed distributions. Table 5-3 shows the results of the respective t-tests, the high p-values confirming that the patient types exhibit similar LOS characteristics and may be merged.

Welch two-sample T-test	t	df	p	Group
HTX & LOTX	-0.74	95.72	0.46	1
HTX & LVAD	0.70	49.17	0.49	
LVAD & LOTX	-1.28	44.42	0.21	
Cardiology & Re-Thoracotomie	-1.55	75.55	0.12	2
TAVI & Aorta	0.29	1071.50	0.77	3
CABG & Klep	0.42	1004.00	0.67	4
Klep & Thoracotomie	-0.61	67.97	0.55	
Thoracotomie & CABG	-0.53	65.36	0.60	
Congenital & Overige TH	1.09	317.95	0.27	5

Table 5-3: Aggregation T-Tests

Comparing the LOS histograms of each type with the LOS distribution of the respective merged group confirms the merge visually (Figure 5-1). Each group has a comparable distribution of LOS, characterized by similar peaks and maxima, which translate well into the overall group LOS histogram. The resulting groups strike a balance between reducing the weekly arrival variability through aggregation and still differentiating patients by LOS characteristics. Due to the fact that the LOS distribution shows distinctive peaks at one-day intervals, indicating regular discharges, the LOS of each group will not be fitted with a distribution, the merged historical LOS distributions will be used as empirical distributions.

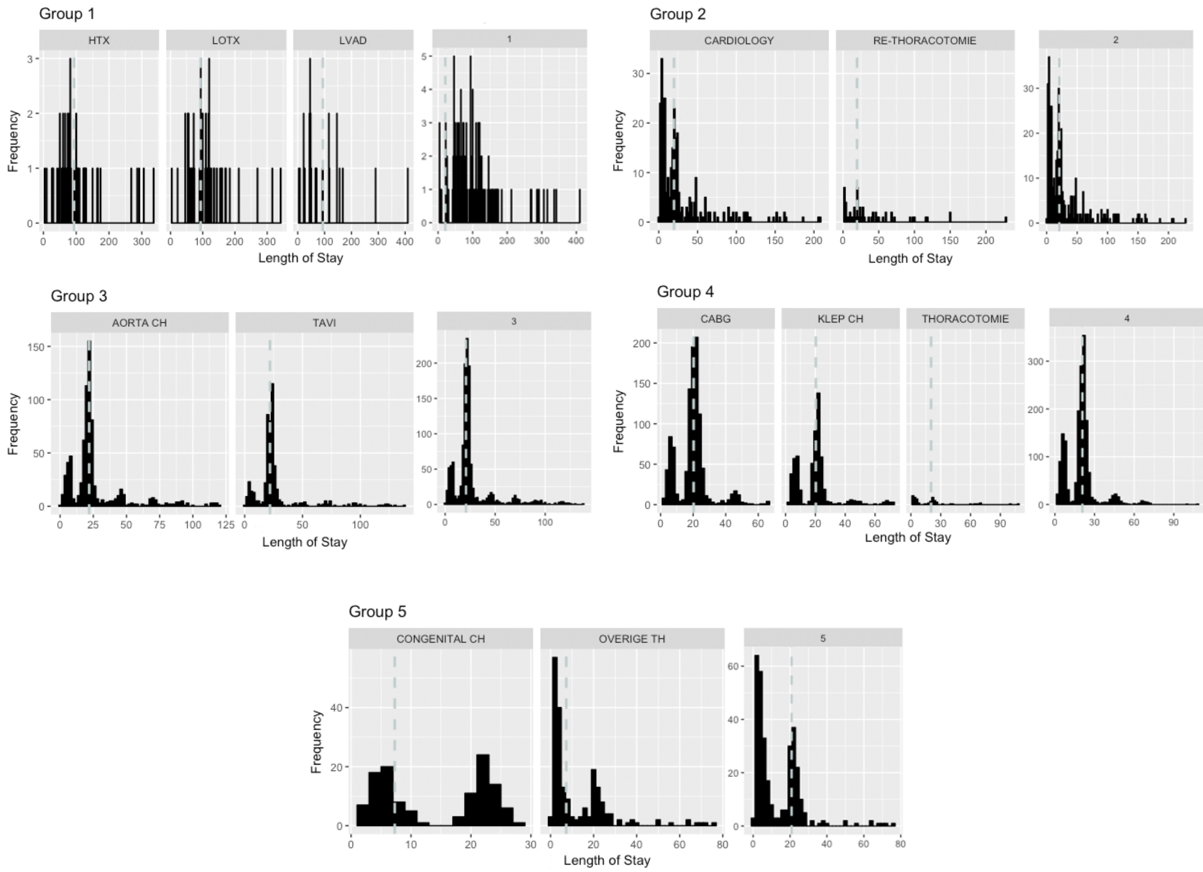

Figure 5-1: Group LOS Histograms

As expected, the weekly arrivals of the newly created groups are less variable compared to the individual patient types they are composed of (see the CoV of Table 5-4 below). The planned patient schedule resulting from the arrival smoothing heuristic should provide sufficient slots to meet surgery demand. The number of weekly surgeries required per patient group is approximated by rounding the historical weekly average to the nearest integer. The total of the rounded weekly arrivals of the aggregated groups is 20, historically the ICU recorded 19.22 planned patient arrivals per week. Considering that four of the patient groups needed to be slightly overstated, the approximation seems appropriate. The historical arrival rates per weekday are listed in Table 5-5 and are required in the next step of the heuristic development, the calculation of patient load.

Patient Group	Avg Weekly Arrivals	Rounded Arrivals	St. Dev.	CoV
1	0.65	1	0.78	1.20
2	1.82	2	1.41	0.77
3	6.38	6	2.55	0.40
4	8.75	9	2.65	0.30
5	1.62	2	1.43	0.88

Table 5-4: Aggregated weekly arrivals

Patient Group	Mon	Tue	Wed	Thu	Fri
1	0.18	0.17	0.29	0.20	0.16
2	0.47	0.47	0.33	0.40	0.33
3	0.91	1.28	0.83	2.12	0.86
4	1.96	2.03	1.61	1.66	1.74
5	0.45	0.41	0.47	0.34	0.33
Rounded Grand Total	4	4	3	5	3

Table 5-5: Daily arrival rates of planned patient groups

5.1.3 *Patient Loadfactor estimation*

The parameter that indicates how many patients are in the ICU on average each weekday will be called “Loadfactor” and needs to be calculated using the previously developed estimates of arrival rates and LOS. The first component of the Loadfactor is the average number of daily patient arrivals, here the average rates calculated as part of the statistical analysis (refer to Table 4-8 & Table 5-5) can be used and do not need to be adjusted. Next, it is necessary to quantify the average capacity the patients occupy after their arrival by incorporating their LOS. Once again, using average values of LOS is not appropriate and would not give an accurate result due to the variable nature of LOS across patients. On the basis of the statistical analysis of LOS, it is possible to make use of the so-called “*Survival Analysis*” to quantify the impact of LOS. As the name implies, this technique is usually applied to study the survival time or time until failure of subjects. However, it can also be applied to more positive “time remaining” scenarios, such as the time to recovery of the ICU patients represented by the LOS in this research.

Survival analysis is based on the properties of the probability distribution of the "time remaining" data. The "cumulative distribution function" (1) of a probability distribution can be used to evaluate the probability of an uncertain variable being lower or equal to a given value.

$$(1)\ F(t) = P(X) \leq t$$

Taking the inverse of this function results in the "survivor function" (2), which can be used to find the probability of the uncertain value being larger than the evaluated point. In other words, it will represent the share of patients who have not recovered yet at the evaluated point in time when applied to the data at hand and will be essential to calculate the Loadfactor.

$$(2)\ S(t) = 1 - F(t)$$

The survival function can be estimated from any fitted parametric distribution. Since it will be calculated for the merged planned patient groups as well, an empirical alternative based on the historical data is preferred. A common method to estimate the survival function from non-parametric empirical data is the "Kaplan-Meier Method". It calculates the survival rate from a given "time remaining" table across time and estimates the survival function. This method was applied to the LOS data per patient type using MATLAB. Plots of the survival functions can be found in the appendix (Figure D-3). Next, the resulting survival functions were used to estimate the average daily probability per patient type of remaining in the ICU was calculated. The resulting probabilities of stay in the ICU are listed in the appendix (Table D-1 & Table D-2). With the arrivals and probabilities of stay in place, the Loadfactor can be calculated.

The Loadfactor calculation takes advantage of the fact that the within-week variation is the highest level of seasonality that is being considered. It follows that the Loadfactor of a certain weekday is the sum of the average arrivals on that day and the remaining patients of previous days. The calculation is visualised in Table 5-6 below.

Day	**Loadfactor**	**Arrivals**	**+1 Day**	**+2 Days**	**+3 Days**	**+4 Days**	**+n Days**
Mon	$\sum$	λ_{Mon}	$\lambda_{Sun} * S(1)$	$\lambda_{Sat} * S(2)$	$\lambda_{Fri} * S(3)$	$\lambda_{Thu} * S(4)$	…
Tue	$\sum$	λ_{Tue}	$\lambda_{Mon} * S(1)$	$\lambda_{Sun} * S(2)$	$\lambda_{Sat} * S(3)$	$\lambda_{Fri} * S(4)$	…
Wed	$\sum$	λ_{Wed}	$\lambda_{Tue} * S(1)$	$\lambda_{Mon} * S(2)$	$\lambda_{Sun} * S(3)$	$\lambda_{Sat} * S(4)$	…
Thu	$\sum$	λ_{Thu}	$\lambda_{Wed} * S(1)$	$\lambda_{Tue} * S(2)$	$\lambda_{Mon} * S(3)$	$\lambda_{Sun} * S(4)$	…
Fri	$\sum$	λ_{Fri}	$\lambda_{Thu} * S(1)$	$\lambda_{Wed} * S(2)$	$\lambda_{Tue} * S(3)$	$\lambda_{Mon} * S(4)$	…
Sat	$\sum$	λ_{Sat}	$\lambda_{Fri} * S(1)$	$\lambda_{Thu} * S(2)$	$\lambda_{Wed} * S(3)$	$\lambda_{Tue} * S(4)$	…
Sun	$\sum$	λ_{Sun}	$\lambda_{Sat} * S(1)$	$\lambda_{Fri} * S(2)$	$\lambda_{Thu} * S(3)$	$\lambda_{Wed} * S(4)$	…

Table 5-6: Loadfactor calculation

	Mon	**Tue**	**Wed**	**Thu**	**Fri**	**Sat**	**Sun**
Unplanned Patient Load	6.20	7.73	8.94	6.27	8.11	5.70	5.28
Planned Patient Load	4.57	5.72	5.26	6.49	5.66	1.97	1.02
Total Loadfactor	10.77	13.45	14.20	12.77	13.76	7.67	6.30

Table 5-7: Loadfactor results

Studying the results (Table 5-7) reveals unsurprisingly that the Loadfactor is higher than the arrival rate across the week and that it follows the same trend as the daily arrival rate, with the patient load peaking on Wednesdays and Fridays and low levels over the weekend. This is due to the fact that the majority of patients leave within the day, leading to pronounced peaks translating to the Loadfactor as well.

A closer look at the percentage change from one day to the next of the unplanned patient arrival rate and the Loadfactor (Figure 5-2) shows a comparatively lower rate of both increases and decreases. This lag exemplifies that the Loadfactor is able to account for patients staying in the ICU and concludingly is a better representation of unplanned patient capacity usage, a necessary basis from which to determine the optimal planned patient allocation.

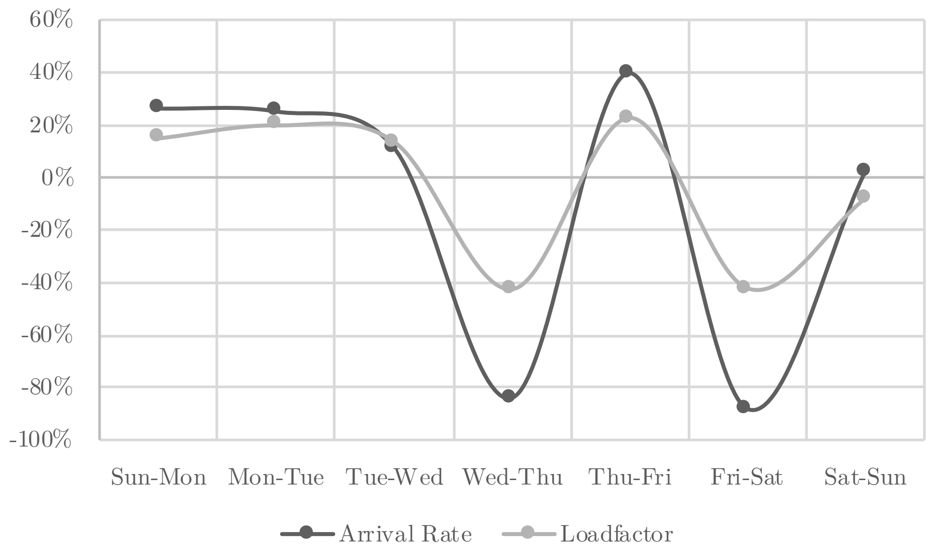

Figure 5-2: Arrival rate and Loadfactor day-to-day change comparison (unplanned patients)

For the heuristic being developed, the unplanned patient Loadfactor is an important parameter, since it is assumed to be fixed and the proposed planned patient schedule should balance it out. The total Loadfactor serves as the benchmark of the historical weekly load distribution.

5.1.4 *Planned patient schedule development*

As part of the Loadfactor calculation, the survival function of the merged planned patient groups was calculated as well (Table D-2). These listed probabilities in conjunction with the number of surgeries per weekday determined by a proposed schedule will define the Loadfactor increase that needs to be distributed across the week. an overview of the total load per patient group is given in Table 5-8.

Patient Group	Group 1	Group 2	Group 3	Group 4	Group 5
Number of required surgeries	1	2	6	9	2
Total additional load	4.85	3.66	9.03	10.87	2.30

Table 5-8: Planned patient load increase

It follows from the aim of the arrival smoothing heuristic that this increase should be balanced with the calculated unplanned patient load and minimized across each weekday to reduce the chances of the ICU reaching its capacity limit. The right method to define a planned patient schedule that meets this balance is optimization using integer programming.

5.1.5 *Optimization problem formulation*

Sets

Planned Patient Groups: $G = \{1,2,3,4,5\}$ (1)

Weekdays: $D = \{1,2,3,4,5,6,7\}$ (2)

Length of Stay: $A = \{0,1,2,3,\ldots,13\}$ (3)

Parameters

S_i = Number of weekly required surgeries for group i, $\forall \mathrm{i} \in G$ (4)

C_j = Daily surgery capacity on day j , $\forall \mathrm{j} \in D$ (5)

UL_j = Unplanned patient Loadfactor on day j , $\forall \mathrm{j} \in D$ (6)

RL_j = Remaining patient Loadfactor on day j from previous week, $\forall \mathrm{j} \in D$ (7)

M_j = Capacity limit of the ICU on day j (8)

P_{ik} = Probability of group i to remain at the ICU after k days , $\forall \mathrm{i} \in G,\ \forall \mathrm{k} \in A$ (9)

Decision Variables

X_{ij} = number of surgeries of group i scheduled on day j , $\forall \mathrm{i} \in G$, $\forall \mathrm{j} \in D$ (10)

Objective Function

$$min \sum_{j \in D} (UL_j + RL_j + \sum_{i \in G} \sum_{k=0}^{j-1} X_{ij-k} * P_{ik})^2 \tag{11}$$

Constraints

Surgery number:
$$\sum_{\mathrm{j} \in D} X_{ij} = S_i \ , \forall \mathrm{i} \in G \tag{12}$$

No surgeries on the weekend:
$$\sum_{\mathrm{i} \in G} X_{ij} = 0 \ , \forall \mathrm{i} \in G, for \ \mathrm{j} = 6 \ and \ 7 \tag{13}$$

Surgery capacity:
$$\sum_{\mathrm{i} \in \mathrm{G}} X_{ij} \leq C_j \ , \forall \mathrm{j} \in D \tag{14}$$

ICU capacity
$$\sum_{j \in D} (RL_j + UL_j + \sum_{i \in G} \sum_{k=0}^{j-1} X_{ij-k} * P_{ik}) \leq M_j, \quad \forall \mathrm{i} \in G, \ \forall \mathrm{j} \in D, \forall \mathrm{k} \in A \tag{15}$$

Definition of RL_j

$$\begin{aligned}
RL_1 = \sum_{i \in G} (&X_{i5} * P_{i3} + X_{i4} * P_{i4} + X_{i3} * P_{i5} + X_{i2} * P_{i6} + X_{i1} * P_{i7} + X_{i5} * P_{i10} \\
&+ X_{i4} * P_{i11} + X_{i3} * P_{i12} + X_{i2} * P_{i13}), \forall \mathrm{i} \in G \\
RL_2 = \sum_{i \in G} (&X_{i5} * P_{i4} + X_{i4} * P_{i5} + X_{i3} * P_{i6} + X_{i2} * P_{i7} + X_{i1} * P_{i8} + X_{i5} * P_{i11} \\
&+ X_{i4} * P_{i12} + X_{i3} * P_{i13}), \forall \mathrm{i} \in G \\
RL_3 = \sum_{i \in G} (&X_{i5} * P_{i5} + X_{i4} * P_{i6} + X_{i3} * P_{i7} + X_{i2} * P_{i8} + X_{i1} * P_{i9} + X_{i5} * P_{i12} \\
&+ X_{i4} * P_{i13}), \forall \mathrm{i} \in G
\end{aligned} \tag{16}$$

$$RL_4 = \sum_{i \in G} (X_{i5} * P_{i6} + X_{i4} * P_{i7} + X_{i3} * P_{i8} + X_{i2} * P_{i9} + X_{i1} * P_{i10} + X_{i5} * P_{i13}), \forall \mathrm{i} \in G$$

$$RL_5 = \sum_{i \in G} (X_{i5} * P_{i7} + X_{i4} * P_{i8} + X_{i3} * P_{i9} + X_{i2} * P_{i10} + X_{i1} * P_{i11}), \forall \mathrm{i} \in G$$

$$RL_6 = \sum_{i \in G} (X_{i5} * P_{i8} + X_{i4} * P_{i9} + X_{i3} * P_{i10} + X_{i2} * P_{i11} + X_{i1} * P_{i12}), \forall \mathrm{i} \in G$$

$$RL_7 = \sum_{i \in G} (X_{i5} * P_{i9} + X_{i4} * P_{i10} + X_{i3} * P_{i11} + X_{i2} * P_{i12}), \forall \mathrm{i} \in G$$

Non-negativity: $$X_{ij} \geq 0 \quad , \forall \mathrm{i} \in G, \ \forall \mathrm{j} \in D \tag{17}$$

Integer: $$X_{ij} \text{ is integer}, \forall \mathrm{i} \in G, \ \forall \mathrm{j} \in D \tag{18}$$

The decisions of this integer programme are taken across three separate sets, the planned patient groups (1), the seven days of the week (2), and the length of stay (3), which is a range of 0 until 13 days, since patients of group 1 may stay up to 13 days in the ICU. A range of parameters needs to be considered when finding the optimal solution, starting with the numbers of surgeries required per patient group (4), set by the rounded arrivals in Table 5-4. Next, the maximum number of surgeries per day needs to be considered (5), which the Erasmus MC set to five, as well as the overall capacity limit of the ICU, which is 16 (6). A fixed parameter in this model is the unplanned patient Loadfactor on each day of the week (7), listed in Table 5-7. As described in section 5.1.3, the total Loadfactor is composed of patient arrivals and the share of patients that remained in the ICU more than one day. In order to facilitate the formulation of the objective function, the optimization problem aggregates the load of all patients that arrived during a previous weekly period in the parameter (8), this parameter is influenced by the decision variable, a dynamic that will be explained later on in this section. The final parameter is the probability per patient group of reaching a certain LOS, which was derived from the survival function (9), this parameter is listed in Table D-2 of the appendix.

The decision variables that can be manipulated are the number of surgeries per patient group to be scheduled on each day of the week (10). The objective function minimizes the sum of squared Loadfactors (11). The Loadfactor sum per day consists of three different components, of which the unplanned patient Loadfactor is the only one that is constant. The second component is the total load remaining from an earlier week, defined by a constraint as explained in the next paragraph. The third and last component of the sum is the planned patient load increase which is calculated by summing the product of the surgeries scheduled on that day and of all previous days of that week with the respective probability of still remaining in the ICU. By minimizing the summed square of the daily Loadfactor it forces the distribution of the additional patient load to be as even as possible, penalizing high values and thus peak loads. This, in turn, leads to the planned patient load balancing out the unplanned patient load across the week, since the unplanned load per day is fixed. Another characteristic of the objective function formulation is that it incentivizes moving patient load onto the weekend, which has a lot of free capacity. There are no planned patient arrivals during the weekend, so its load is entirely determined by the "patients remaining" load from earlier days and the constant unplanned load. Since the sum of squares is optimized and workdays have a much higher Loadfactor (see Table 5-7), the most effective way to minimize the objective function is to decrease and even out the load during workdays as much as possible by moving load onto the less utilized weekend. The only way to do this is by scheduling long stay patients on Fridays, which pushes their "remaining patient" load onto the weekend.

A number of constraints ensure that a feasible solution is found, the first one (12) sets the number of required surgeries per patient category, the second one (13) prevents surgeries being scheduled on the weekend. Capacity limits are enforced by (14), the daily surgery limit (which on Friday is essential to keep the weekend workload increase manageable) and (15), which makes certain that the total Loadfactor on a certain day does not exceed the overall capacity limit of the ICU. This constraint is only included for good measure, it will not be binding because the Loadfactor never represents actual patient numbers at the ICU, but the expected average amount of patients at the ICU. This value does not reach the capacity limit of 16 at any point (Otherwise the utilization of the ICU would be over 100% in reality). Constraint (16) creates the link between the remaining patient loadfactor from a previous period and the schedule as determined by the optimization. Since the schedule found by solving the optimization problem will repeat itself every week, the decision variable does not change across weeks and the remaining patient load from previous periods can be defined as the share of surgery patients scheduled from Monday to Friday that stay until a new week begins. Taking Monday as an example, RL_1 includes the share of patients that arrived the Friday before that are still in the ICU after three days ($X_{i5} * P_{i3}$) up until the share of patients who arrived on a Tuesday and are reaching their thirteenth day of stay in the ICU ($X_{i2} * P_{i13}$). This inclusion guarantees that the full extent of the planned patient Loadfactor is considered and that the schedule generated for a single week is truly considered as an infinitely repeating weekly schedule. Finally, (17) and (18) are the necessary integrity constraints that define the decision variable as being non-negative and integer.

A solution for the stochastic and continuous problem of capacity allocation can only be found by restricting the solution finding to a specific period that can then be generalised across time. This is why the solution finding process was restricted to a weekly time horizon at the beginning of this chapter, based on the finding that the highest level of relevant seasonality was within week seasonality and the fact that the ICU uses a one-week planning horizon in practice. For the optimization problem described here, this entails a number of conditions that define a stable state in order to guarantee that it is able to find an optimal solution according to the KKT condition:

- ICU capacity does not change from week to week
- Surgery capacity does not change from week to week
- Unplanned patient load does not change from week to week
- The number of surgeries to be scheduled does not change from week to week (which implies an unchanged additional planned patient Loadfactor per week)

5.1.6 *The result of the optimization problem*

The optimization problem detailed in the previous section was solved using the Solver package of Microsoft Excel. The package is able to solve integer programmes, which are harder to compute than linear programmes, by utilizing the "branch and bound" algorithm, which separates the solution set into subsets ("branches") and only exhaustively tests possible solutions if a branch is within certain bounds of the current optimal solution. The default setting of this non-linear solution approach in Solver is set to allow a 1% tolerance around the integer optimal result, this parameter was reduced to 0% in order to find the true optimal solution. The optimal value of the objective function after optimization is 930.39, the more relevant and insightful results are listed in Table 5-9, 5-10 and 5-11. Inspecting the Loadfactor distribution across the week (Table 5-9) reveals that the additional planned patient load was evenly balanced across the days of the week. Wednesday and Friday remain peak days, however, the difference compared to the rest of the week was diminished. Compared to the historical distribution the largest deltas are on Monday, Tuesday and Wednesday.

Day	Mon	Tue	Wed	Thu	Fri	Sat	Sun
Historical Loadfactor	10.77	13.45	14.20	12.77	13.76	7.67	6.30
Optimized Loadfactor	12.31	12.36	12.94	12.19	13.90	8.27	6.97

Table 5-9: Optimization result 1, Loadfactor distribution

Table 5-10 depicts how the integer programme scheduled the majority of surgeries on Mondays and Thursdays, which had the lowest workday load scores originally.

Day	Mon	Tue	Wed	Thu	Fri	Sat	Sun
Unplanned Patient Loadfactor	6.20	7.73	8.94	6.27	8.11	5.70	5.28
Planned surgeries scheduled	5	3	3	5	3	0	0

Table 5-10: Optimization result 2, surgeries scheduled per weekday

Diving deeper into the proposed schedule by examining the allocation per patient group (Table 5-11) confirms the earlier hypothesis that long LOS surgeries should be scheduled towards the end of the week. The resulting optimal schedule places surgeries that result in long LOS (Groups 1 & 2) on Fridays, moving their load onto the less busy weekend. Surgeries resulting in short LOS (Group 5 & 4) are scheduled on Tuesdays, in order to reduce the load that carries over into Wednesday, which is a peak load day.

Patient Group	Mon	Tue	Wed	Thu	Fri
1	0	0	0	0	1
2	0	0	0	0	2
3	0	0	0	5	1
4	5	1	3	0	0
5	0	2	0	0	0

Table 5-11: Optimization result 3, the surgery schedule per patient group

By removing the unplanned patient Loadfactor as well as the Loadfactor increase on the first day by planned patients (which is equal to the arrival rate) from the daily total it is possible to investigate how the optimization diverted the load of all planned patients that remain at the ICU for more than one day. Looking at this value exclusively visualises the effect of the changes described in the previous paragraph. Figure 5-3 highlights the differences between the “remaining patient” load of the historical and optimized patient distribution, the shift onto the weekend and the mid-week load reduction becomes evident.

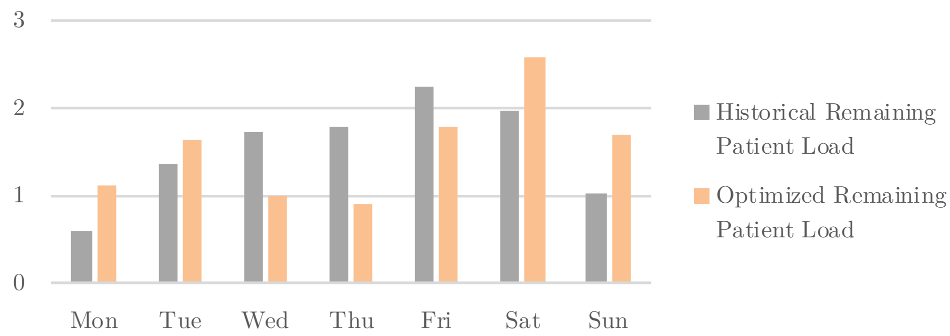

Figure 5-3: Diversion of remaining patient load

In order to execute the optimization, the variable patient load was transformed into the deterministic Loadfactor. The effect of the arrival smoothing heuristic developed here on the bumping rate needs to be tested in the variable ICU environment and compared to the base case. This is accomplished by creating a representative simulation of the ICU at the Erasmus MC.

6. VALIDATION

6.1 The Simulation

The simulation allows examining "What if?" scenarios and testing the effectiveness of the arrival smoothing heuristic (referred to as the *ASH* from now on) to improve the capacity management of the ICU. Discrete event simulation was found to be able to appropriately represent the dynamics of the ICU environment (Günal & Pidd, 2010). It is defined by events that change the state of the system, in the case of the ICU these events are the arrival and departure of patients. Thus, while the simulation time will be defined by actual days and hours, the simulation will always skip to the occurrence of the next event until the overall run is completed. The Java-based software AnyLogic (v. 8.3.2) was used to build the model, which is a simulation package well suited to model agent based (patients in this case) discrete event environments. Similarly to the integer programme, the simulation makes use of the previously estimated arrival rates and LOS statistics. Before the simulation model can be conceptualized it is necessary to describe how the results of the statistical analysis were integrated.

6.1.1 *Arrival rate definition – unplanned patients*

The analysis of unplanned patient arrival patterns revealed that there is an overall time-based seasonality of patient arrivals, as well as arrival differences exhibited by certain patient types only (e.g. caused by the deviation of patients of a certain type to another hospital). Thus, each patient type is considered separately when defining the arrival rates. As discussed, the within-week and within-day differences are being considered, leading to a total number of 7*24=168 arrival rates per patient type. Since unplanned patients are differentiated into 10 categories, the number of arrival rates that need to be defined is 1,680. The λ parameter of each Poisson distribution necessary for the simulation was calculated based on the historical average per each timeslot, the results are visually summarized in Figure C-1 of the appendix.

6.1.2 *Arrival rate definition – planned patients*

Considering planned patient arrivals in the simulation necessitates the definition of an integer allocation of planned patients per group per day for the base case. The distribution across the days of the week was done by mirroring the historical average arrivals per weekday per group as closely as possible. The rounded of the total surgeries per group allocated this way matches the total optimized by the ASH. The result can be seen in Table 6-1.

	Patient Group	Mon	Tue	Wed	Thu	Fri
Historical	1	0.18	0.17	0.29	0.20	0.16
	2	0.47	0.47	0.33	0.40	0.33
	3	0.91	1.28	0.83	2.12	0.86
	4	1.96	2.03	1.61	1.66	1.74
	5	0.45	0.41	0.47	0.34	0.33
	Rounded Grand Total	4	4	3	5	3
Approximation	1	0	0	1	0	0
	2	1	1	0	0	0
	3	1	1	0	3	1
	4	2	2	1	2	2
	5	1	0	1	0	0
	Grand Total	5	4	3	5	3

Table 6-1: Approximated weekly schedule of planned patient arrivals

As explained in the methodology section, the time of arrival on a certain weekday will be simulated by queuing arriving patients at 8 am and delaying them by the surgery duration. Surgery duration ranges were provided by the Erasmus MC, due to the lack of details a triangular distribution of the minimum, maximum and mode is assumed (Table 6-2).

Patient Group	Types	Min	Mode	Max
1	HTX, LOTX, LVAD	5.5	6	7
2	Cardiology & Re-Thoracotomie	3	5	7
3	TAVI & Aorta	4	4.5	6
4	CABG, Klep, Thoracotomie	3	4	4.5
5	Congenital & Overige	3	4	6

Table 6-2: Surgery durations

6.1.3 *Length of stay*

Length of stay can be considered by the simulation model using the fitted distributions developed in 4.2.3. For certain LOS distributions the overall fit of the fitted distributions is unsatisfactory (see Figure D-2, e.g. Pericarditis and the planned patient types), here the simulation resorts to using the historical LOS data as empirical distribution.

6.1.4 *Simulation logic*

Implementing the results of the statistical analysis assures an accurate representation of the parameters influencing the ICU. The physical restrictions, in addition to the operational dynamics of the ICU, need to be modelled closely as well in order to arrive at an accurate simulation of the ICU environment. Server capacity in the service system at hand is the number of available staffed ICU beds. The maximum possible capacity is assumed to be fixed, the theoretical maximum of 18 beds after the ICU move is reduced to 15. Two beds are not suitable due to having limited equipment and no sunlight. In order to account for the removal of the top 5% of patient records that were characterized as outliers, an additional bed is removed from the analysis. The availability of nurses, which can restrict available capacity, is not modelled due to missing information on staffing in the new ICU unit. Moving on to the patient journey, one source block per unplanned patient category is connected to a schedule that lists the Poisson arrival parameters for each hour of the week of its type, which is cycled through for all simulated weeks. According to the current simulation time, arrivals are then generated from the corresponding Poisson distribution. Planned patient arrivals are incorporated by generating the number of patients per week per group and then distributing them across the days of the week according to the quota defined by the schedule at 8 am each day. These patients enter a delay block that represents the surgery theatre, which accepts three simultaneous patients. The delay time is set by drawing the surgery duration from a triangular distribution that represents the duration range provided by the Erasmus MC (Table 5-11).

Patients that arrived at the ICU advance to the "Bed Allocation" block. Here, a function is called that evaluates the current state of the system. If there is sufficient capacity to accommodate the patient he is allocated to a bed and advances if not, the current patient population is examined and the patient with the lowest remaining length of stay (representing the "healthiest patient") is bumped and the new patient is admitted. The bumping event is recorded and marked in the simulation timeline as well.

All admitted patients advance to the Length of Stay delay block. Here the LOS (which represents service time in the ICU system) is drawn from the respective patient type probability distribution. The LOS sets the moment of discharge of the patient, once that point in time is reached, the patient leaves the system and the occupied server (the patient bed) becomes available for new patients. The following statistics are recorded in order to evaluate the validity of the model as well as the effectiveness of the ASH:

- The number of arrivals generated per patient type
- The generated LOS distribution per patient type
- The number of patients in the system
- The number of Bumping occurrences

Screenshots of the simulation model are included in the appendix, it has three different views: the routing logic (Figure E-1), the statistical backbone (Figure E-2) and the performance dashboard (Figure E-3). The logic flow of the simulation is summarized in Figure 6-1 below as well.

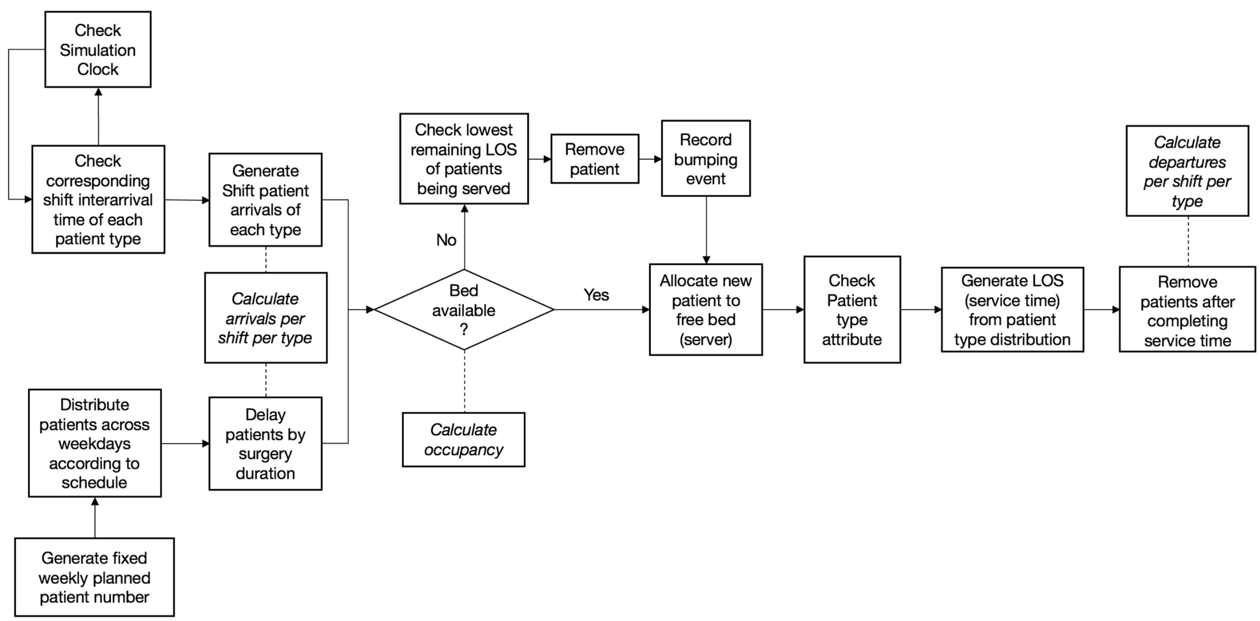

Figure 6-1: The Simulation model

6.1.5 *Simulation model assumptions*

A number of assumptions that underpin the model are mentioned here:

- Recorded arrivals at the ICU represent departures of other internal entities, such as the operating theatre and emergency room. Thus, it is assumed that the recorded historical admission rates represent the true arrival rates at the ICU.
- Since the ICU receives emergency patients with acute conditions every unplanned patient needs to be admitted & cared for.
- It is a policy of the Erasmus MC to not cancel surgeries, only delays are possible.
- The time-varying Poisson arrival assumption for unplanned patients holds.
- The approximation of planned patient arrivals holds.

- Patient arrivals per type are assumed to be independent of each other.
- Beyond the Seasonal/Weekly/Daily pattern, no other seasonality pattern influences patient arrivals.
- The fitted LOS distributions accurately model the patients' service time.
- A triangular distribution of the surgery duration ranges is representative.
- Beds become immediately available for a new patient once the previous occupant has been discharged.
- The remaining LOS of a patient is representative of his level of health.

6.1.6 *Base case validation*

Before testing the effect of the ASH, the ability of the model generate accurate arrivals and LOS is validated. The simulation is set to run for exactly one year, a timeframe that ensures that the ICU model reaches a stable setting after a run-in period. Furthermore, in order to accurately account for the stochastic environment of any statistic, the simulation was run 100 times per case and since around 2700 patients are generated in one year, it thus includes 270.000 unique draws from both the arrival and LOS distributions using different seeds.

The validation is done by comparing the number of generated arrivals per patient category as well as the generated LOS distributions per patient category to historical values. The comparison of average yearly historical arrivals and the results of the simulation can be examined in

Table 6-3. It can be concluded that the time-varying Poisson arrival rate implemented in the simulation model is able to generate an accurate number of arrivals per unplanned patient type per year.

Patient type	Historical average	Simulated arrivals
STEMI	483	480
Overige	469	475
Ritmestoornissen	226	235
NSTEMI	148	132
OHCA	139	147
Hartfalen	137	129
Pericarditis	55	61
Endocarditis	22	20
Tamponade	12	11
ECMO	5	7

Table 6-3: Validation of unplanned patient arrivals

An important characteristic of patient arrivals is the exhibited seasonality. As explained in section 4.2, the varying degrees of seasonality per hour and weekday need to be incorporated in the time-varying arrival rate. Examining the historical average of arrivals per shift in comparison to an exemplary simulation run (Figure 6-2) reveals that, while certain peaks are more pronounced (The stark difference on Mondays is due to the rounded planned patient schedule), the simulation is able to mirror the within-week variation.

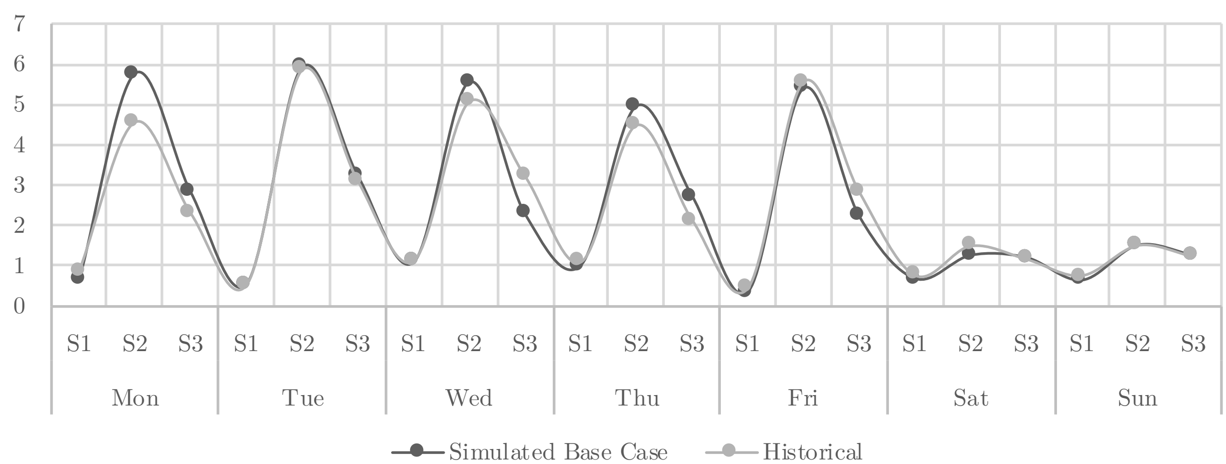

Figure 6-2: Validation of seasonality

The second uncertain variable that needs to be approximated by the simulation is the patient LOS. Table 6-4 & Table 6-5 compare three characteristics, namely the average, maximum and standard deviation of historical patient LOS per patient category with the results generated by the simulation model. The resulting values validate that the simulation is able to accurately model patient LOS.

Patient Type	Historical			Simulated		
	Average	Max	SD	Average	Max	SD
ECMO	194.28	436	143.82	189.50	475.85	126.40
OHCA	95.94	305	68.36	90.76	282.26	65.67
Hartfalen	67.54	335	69.53	69.40	332.70	72.52
Endocarditis	51.77	180	46.56	50.56	175.19	45.72
Tamponade	33.30	115	49.23	29.93	140.35	36.24
Pericarditis	27.45	135	26.39	28.39	135.61	28.46
NSTEMI	17.65	77	17.01	18.73	72.19	18.45
Overige	14.58	79	15.78	14.69	75.76	15.92
STEMI	11.02	81	14.11	11.48	85.63	13.90
Ritmestoornissen	9.90	76	13.54	10.44	75.92	13.70

Table 6-4: Validation of unplanned patient LOS distributions

Patient Group	Historical			Simulated		
	Average	Max	SD	Average	Max	SD
1	105.37	410.15	74.92	103.07	388.75	81.46
2	31.38	228.43	38.82	28.94	162.68	35.60
3	26.68	138.02	21.73	26.01	138.70	23.59
4	19.73	108.12	11.55	19.75	92.52	11.98
5	13.12	76.98	12.74	12.61	76.13	11.31

Table 6-5: Validation of planned patient LOS distributions

6.1.7 *Test of the arrival smoothing heuristic*

Following the successful validation of the simulation model, the effect of the ASH can be tested and compared to the base case. The simulation is able to generate the previously unobserved patient bumping variable in the new ICU setting, it counts the number of bumps as well as the percentage bumprate, taking into account the total number of patients received.

As described earlier, a bump occurs when a patient arrives at a full ICU. The ASH aims to reduce the chances of the ICU filling up by balancing the Loadfactor of planned and unplanned patients. Implementing the schedule found by the optimization programme it is possible to test if it improves on the base case. For that matter, the base case scenario, as well as the optimized setting were run 1000 times using random seeds (thus each run represents a unique combination of the included stochastic variables). The runs resulted in average bumprates of 2.98% and 2.37% (Table 6-6) respectively, which represents a significant reduction of 20.47% and proves that the ASH is able to improve ICU capacity management. The shift of the number of bumps between the two scenarios can be inspected on the histogram (Figure 6-3) comparing the frequency distribution of bumping occurrences.

	Average Bumprate	**Average Number of Bumps**	**SD**
Base Case	3.98%	85.07	28.44
ASH	2.37%	65.47	24.06
Change	-20.47%	-23.04%	-15.42%

Table 6-6: Arrival smoothing heuristic performance

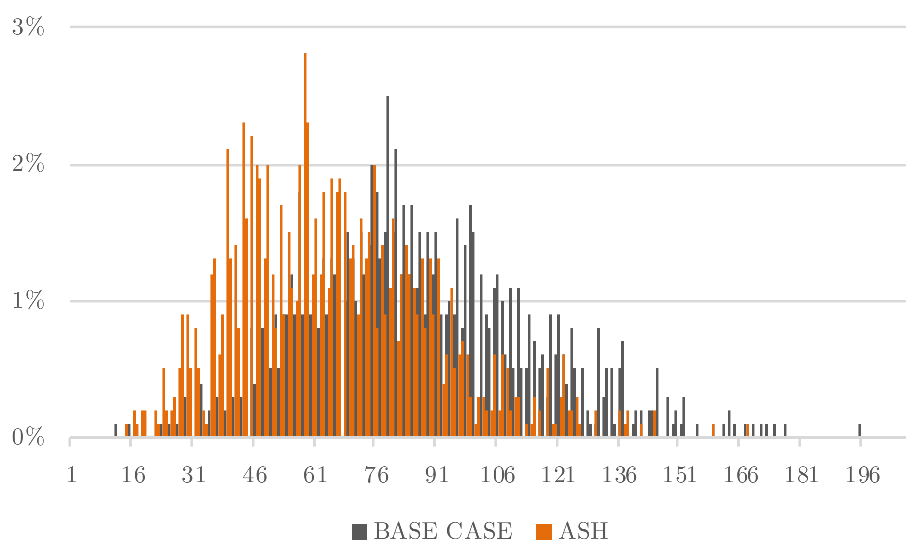

Figure 6-3: Base case vs ASH, number of bumps

The statistical significance of the reduction in the number of bumps was tested using a two-sample t-test assuming equal variances, the p-value of 0% (reported in Table 6-7) confirms that the ASH is able to reduce the occurrence of bumping.

Paired sample T-test	t	df	p
Base Case & ASH	132.46	999	0.00

Table 6-7: T-test heuristic improvement

The effectiveness of the ASH can be observed in more detail by comparing it to the base case while looking at an exemplary simulation run. Comparability is ensured by fixing the same seed of both the base case and the ASH run, thus the ICU faced identical unplanned patient arrivals and LOS characteristics over the course of the simulation run. To ensure valid average measurements of this particular scenario it was run for 10 years and thus includes 520 unique weeks. The randomly chosen seed resulted in a scenario with 2,737.4 yearly patients, of which 93.3 (3.41%) were bumped in the base case and 77 (2.81%) were bumped after applying the ASH, a bumping reduction of 17.59% in this particular scenario (average values per year). Figure 6-4 summarizes the effect of the ASH on the average weekly patient arrival rate, average weekly ICU occupancy and the number of bumps visually. All metrics are indicated per week per shift. The arrival rates are nearly identical, the only difference being that the total number of surgeries scheduled in the base case and by the ASH differs on Tuesdays and Fridays (4 vs 3 and 3 vs 4). Contrasting the trajectory of the two occupancy lines demonstrates the effect of the ASH heuristic. By moving planned patients according to their Loadfactor an evener distribution of patient load is achieved. Peak loads on Wednesdays and Fridays are reduced and days with lower capacity usage (Mondays & the weekend) are better utilized. Inspecting the number of weekly bumps per shift (the bars refer to the left axis) proves that by balancing occupancy across the week the occurrences of the ICU reaching its capacity limit can be reduced and the Bumprate is reduced as a consequence. The most significant reductions were achieved on Thursdays and Fridays, while Mondays, for example, saw an increase of Bumps.

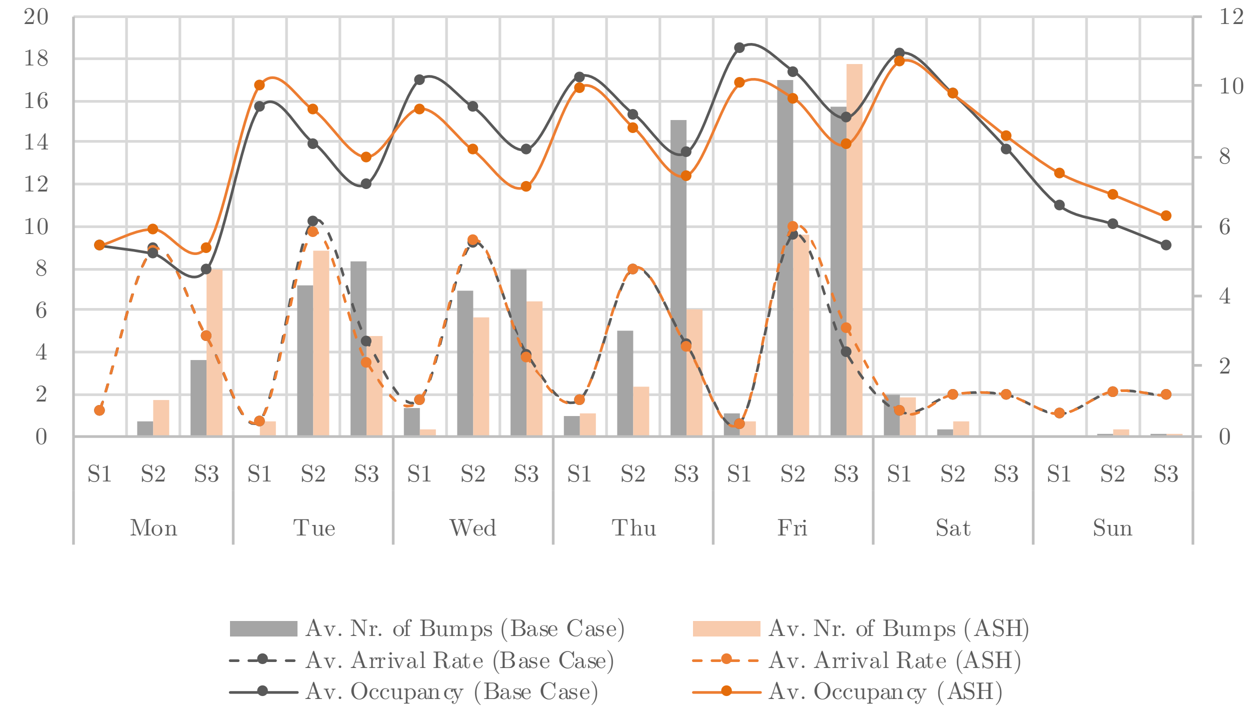

Figure 6-4: Single run comparison, base case & ASH

The ability and effectiveness of the ASH to improve capacity management at the ICU and to reduce the number of premature patient discharges is confirmed by the results of the simulation. These findings will form the core of the recommendations for the ICU's management team discussed in the next chapter.

7. CONCLUSION & RECOMMENDATIONS

7.1 Research Conclusion

This research project successfully developed a capacity management scheme that allows the ICU at the Erasmus MC to reduce the occurrences of premature discharge by calibrating the planned patient load to match the patterns of unplanned patient arrivals & LOS. It required a multi-step process that began by surveying the academic literature of previous ICU studies.

The survey first focussed on the methods used, finding that while queuing and simulation are most commonly used, the uncommon combination of optimization and discrete event simulation is the best fit for this study, for the reason that it brings solution finding together with a testing environment that is able to reflect the uncertain characteristics of the ICU environment. In terms of heuristics that improve how the ICU deals with its uncertain environment, the review found that most studies aim to reduce negative consequences in form of surgery cancellation or patient refusals, while the bumping consequence at the centre of this study has seen less attention. Most ICU studies were set in an environment in which capacity could still be influenced, a parameter that could not be changed in this study and thus required the development of an entirely new heuristic. The uncertainty an ICU faces is the defining characteristic of the operational challenges of the ICU. How to capture and account for the variability of arrival rates and LOS was studied, finding that avoiding categorical averages by using a time-varying Poisson arrival process combined with right-tailed distribution fitting for patient LOS was best suited. The estimation of these variables needed to be based on a coherent patient classification that segregates patients that exhibit differing arrival and LOS characteristics. By adopting the diagnosis categories of the Erasmus MC as patient type classification the interpretability of the study results is greatly improved while also achieving the necessary reduction of within-group variability.

Underpinning the heuristic development was a thorough statistical analysis and discussion of the ICU's historical data, which also provided the basis for the modelling of the previously unobserved arrival and LOS characteristics of each patient. Apart from providing the management of the ICU with a better understanding of these factors, it also revealed the improvement possibility of rescheduling elective surgeries on which the ASH heuristic is based. By combining arrival frequency and survivor functions expected patient capacity usage was quantified. This parameter successfully accounts for the capacity usage "lag" across various days caused by the probability of extended patient LOS. A regular weekly elective surgery schedule was defined by studying planned patient arrivals and aggregating them into five larger groups according to their LOS characteristics. With these two definitions in place, the expected patient load was allocated across the week using a squared-sum integer optimization programme that evenly balanced the planned patient load with the unplanned patient load.

The ASH heuristic was then tested by mirroring the ICU environment in a discrete event simulation, which proved that the ASH successfully reduces the number of premature dismissals at the ICU by reducing the number of times the ICU reaches its capacity limit.

While the ASH developed here is specific to the ICU of the Erasmus MC, the underlying methodology developed and explained in this research can be replicated at any ICU that receives a mix of unplanned and planned patients.

7.2 Recommendations to the Erasmus MC

Presented with the question of how to improve the capacity management at the new combined ICU unit, this research arrived at a newly developed heuristic that achieves a reduction in patient bumping. The ASH defines a weekly schedule that ideally balances planned and unplanned patient arrivals considering their arrival rates and LOS. The schedule is summarized in Table 7-1. The main recommendation that follows is to suggest the resulting allocation of elective surgeries to the surgery planning team and determine if an implementation is feasible for both sides. Barring any scheduling conflicts from the surgery planners, implementing this schedule is relatively straightforward. The amount of surgeries scheduled reflects the most common elective surgery load the ICU receives per week and the merged groups introduce some flexibility into which surgery exactly can be scheduled on each day. Implementation of the ASH will not only reduce the occurrence of premature dismissals, but the even patient load across the week also carries other benefits, such as the possibility to schedule staff more evenly.

Weekday	Elective surgery kind	Amount
Monday	CABG, Klep, Thoracotomie	5
Tuesday	CABG, Klep, Thoracotomie	1
	Congenital, Overige TH	2
Wednesday	CABG, Klep, Thoracotomie	3
Thursday	TAVI, Aorta	5
Friday	HTX, LOTX, LVAD	1
	Cardiology, Re-Thoracotomie	2
	TAVI, Aorta	1

Table 7-1: Detailed surgery schedule

Beyond the actionable recommendation of implementing the ASH, this research results in a range of other findings relevant to how the ICU at the Erasmus MC operates. First of all, it informs about the trends and characteristics of patient arrivals and LOS. This information can be used to revise assumptions about these factors at the ICU. The ICU management should continue to collect data to keep this information up to date and keep observing its development. Data collection at the ICU, in general, should be expanded and include records on ad-hoc decisions, such as patient bumping. This will allow additional future research into the dynamics of capacity management at the ICU.

The statistical analysis of this research uncovered a disproportionate capacity usage by Hartfalen & OHCA patients, which tend to stay in the ICU for long periods of time. The ICU's management should consider if the medical care requirements of these patients allow for a shortening of their stay at the ICU or if it is possible to expand the regional collaboration with Maasstad hospital to further share the patient load of these types. These efforts would make more ICU capacity available.

Another important finding was highlighting the seasonal nature of patient arrivals on a weekly and hourly basis, insights that should be considered when creating the nurse schedule of the ICU.

Lastly, the ICU should introduce LOS estimation at admission using medical judgement of the patient's condition and record it. The resulting discharge date estimate would improve short-term planning and could be integrated into a variation of the ASH that expands its way of working to include an adaptive element based on the current state of the ICU.

7.3 Limitations

The main limitation of this research is that its analysis is focussed on a single entity, the Erasmus MC, which is also the sole source of data. While this infers that the final recommended patient schedule is specific to this ICU, the approach used here is suited to study other hospital wards as well. Furthermore, the concept and methodology behind the proposed ASH, that defines patient load as the combination of arrival and stay probability, can be replicated to describe the expected capacity usage and improve on it in other hospital units as well.

Another limitation of this study is the fact that it describes a new environment (the ICU after the move), characterised by a new patient mix, based on historical data of two separate units. Capacity usage and the premature discharges resulting from reaching the capacity limit at the new unit had to be approximated. This limitation was countered by mirroring the observed dynamics of the ICU very closely and validating the individual estimations used by the simulation individually with historical data. The successful validation gives confidence in the resulting parameters.

A third limitation relates to the LOS estimation as well as the usage of survival functions as part of the average Loadfactor estimation. Both rely on probability distributions and share the common fallacy that they are based only on the data observed so far. In reality, the ICU might receive a patient of any type who's stay exceeds the highest LOS observed for this type so far and thus extends the survival function at any point in time. This limitation is more relevant for small patient types that have little historical data. Another limitation is the possibility that the described Poisson arrival rates and LOS distributions are overfitted, reducing the ability to generalise the resulting findings. This was countered by using the maximum amount of data available to estimate all parameters. Aggregating the underlying measures over the four years enables the estimation to pick up the signal within the data more easily.

7.4 Recommendations for Future Research

The ASH developed in this study is a novel heuristic that is able to improve capacity management in a fixed capacity environment in its current form. Nevertheless, there is a clear opportunity to iterate and improve on the ASH. In its current form, it takes on a deterministic form to arrive at a solution that provides a balanced patient load most of the times. Two improvements present themselves for future research. The first being an expansion of the heuristic to incorporate the ICU's current state and transforming it into a short-term state-dependent arrival smoothing heuristic. The most relevant predictor of future ICU capacity is its current state, which was not observable in this study. Another beneficial adaptation would be to make the elective surgery component that's being allocated adaptively by allowing for changing daily limits or changing numbers of required surgeries. These two improvements could also be combined to form a more flexible and even more powerful heuristic that is able to balance patient load with even higher precision, especially useful for short-term decision making.

An additional evolution of the ASH that could be a topic for future research would be to combine it with an elective surgery scheduling problem that aims to build an effective surgery schedule while keeping ICU capacity in mind. While examining this, planned patient allocation could be considered on a shift- or hourly level. Drilling this far down could bring improved performance by considering hourly differences in patient levels, which do exist as shown in this research.

Great potential lies in further exploring the estimation of LOS as well. Joint research of medical professionals and operations researchers that combines medical expertise with prediction model building would be a valuable addition to the available literature and surely contribute to improve heuristics such as the one presented in this study as well as find direct application in practice.

7.5 Next Steps in Collaboration with the Erasmus MC

The research commenced here and the collaboration with the Erasmus MC will continue after the submission of this thesis as part of my double degree programme. My MSc in Supply Chain Management, of which this thesis is the final component, will be complemented by the MSc in Business Analytics at the ESADE Business School. The double degree programme entails the creation of a combined thesis, this being the first part. As part of my second degree, I will iterate on this project together with the Erasmus MC with the aim to both expand and implement the findings in a more practical way. Additional data collection, that includes more detail of ad-hoc events such as bumping occurrences and staffing changes, was agreed on with the Erasmus MC. The ICU management plans to integrate discharge estimates into its workflow, thus the project iteration will aim to add a state-dependent ASH implementation as well as explore how LOS estimation, done by medical professionals, can be integrated. The aim of both parts combined is to facilitate integrating a refined ASH policy into the Erasmus MC's patient management system.

APPENDIX A. PATIENT TYPES

Patient Type	Diagnosis
STEMI	ST-Elevation Myocardial Infarction: a heart attack that is characterized by the blockage of one of the heart's major arteries and requires immediate attention
NSTEMI	A Non-STEMI heart attack that may or may not require immediate attention
OHCA	Out of Hospital Cardiac Arrest: the loss of heart function, caused by sudden irregularities of the heart's rhythm, requires immediate attention
ECMO	Extra Corporeal Membrane Oxygenation: a life support procedure in which the functions of the heart and lung are performed by a machine
Ritmestoornissen	Heart arrhythmia that requires cardiological treatments
Hartfalen	Heart failure
Pericarditis	An infection of the pericardium (the tissue around the heart)
Endocarditis	An infection of the inner heart valves
Tamponade	The accumulation of fluid between the pericardium (the tissue around the heart) and the heart, which can restrict the ability of the heart to function
Overige	An umbrella term to cover a variety of other cardiology diseases

Table A-1: Description of unplanned patient types

Patient Type	Diagnosis
CABG	Coronary Artery Bypass Graft: commonly known as "bypass" surgery, it restores blood flow by bypassing a coronary artery blockage
Aorta	Surgery of the Aorta (the body's main artery), often necessary to treat aneurysms, abnormalities of blood vessels
TAVI	Transcatheter Aortic Valve Implantation: a newly developed procedure to replace aortic valves that show signs of failure
EFO	Electrophysiology: a procedure to examine and determine the treatment of abnormal heart rhythms
Klep	Open heart surgery to treat heart valve irregularities
Cardiology	An umbrella term covering surgeries that need cardiology care afterwards
Congenital	Surgical treatment of heart conditions/defects existing since birth
Overige TH	An umbrella term to cover a variety of other thorax surgeries
Kinderhart CH	Children heart surgery
Long CH	Lung surgery
Thoracotomie /Re-Thoracotomie	A procedure to open the chest to access the lungs - often used for lung cancer patients. Re-Thoracotomie refers to a repetition of the same surgery
HTX / LOTX	Heart / Lung transplant surgery
LVAD	Left Ventricular Assist Device: a surgery that implants a mechanical, battery-powered blood pump, necessary for patients with advanced heart failures

Table A-2: Description of planned patient types

APPENDIX B. PATIENT TYPE CLASSIFICATION

The patient type matching process:

1. The ICU dataset is split per unit (ICCU & ICTH)
2. The “diagnosis” (ICCU) and “operations”(ICTH) fields are converted to all caps
3. ICTH only: split patient population by age (Children $<$ 18 years $<$ Rest), all children are allocated the patient type “Kinderhartchirurgie”
4. Import the ICCU/ICTH patient type keyword allocation (In its initial state it only included the name of each patient type)
5. A nested IF function checks the “diagnosis” or “operations” fields for the presence of each keyword on the list
6. All positive matches are listed in a new column “TypeMatch”
7. An IF function checks if the entries have been matched, the dataset is split into “Matched” and “No Match”, all “No Match” entries enter a separate keyword generation workflow, explained below
8. ICTH only: the “Matched” entries are joined with the “Kinderhartchirurgie”-set separated earlier
9. The “TypeMatch” column may include more than one match, possibly duplicates, it is thus required to find the best fitting match. This was already considered during the keyword list generation. Since the software checks for keywords from top to bottom the less specific “catch-all” category “Overige” was placed at the very end, since it is expected that many patients will match its keywords. Furthermore, NSTEMI is a different kind of STEMI patient that needs to be considered separately. Since all NSTEMI’s will match to the keyword “STEMI”, but not the other way around, NSTEMI was listed before STEMI in the keyword list. It was found that patient entries had up to three matches, the first step is splitting these into separate columns, using “space” as a delimiter
10. A new column “FinalMatch” is created, if a patient had only one match, it takes on that value
11. An IF function checks if the multiple matches are duplicates. If yes, it takes on the first match, if not, it checks if either match is “Overige”, if yes it takes on the first match, if not it needs to be inspected
12. Finally, the ICCU and ICTH datasets are merged again and the final matched dataset is exported

Keyword generation process, using unmatched patient entries:

1. For all patients that were not matched the "diagnosis" or "operations" fields are grouped to aggregate patients in case there is identical "diagnosis" or "operations" information, the occurrence of each grouped "diagnosis" or "operations" information is counted
2. The grouped "diagnosis" or "operations" information that appeared more than 50 times (representing common diagnoses that are easy to allocate right away) in the data is separated, the resulting list of "diagnosis" or "operations" information is provided to the ICU management which extracts keywords and allocates these to patient types
3. The grouped "diagnosis" or "operations" information that appeared less than 50 times represents a large part of data that cannot easily be summarized. To derive a manageable list of keywords from this information it goes through a "fuzzy match" process, which works as follows:
   i. Punctuation and conjunctions (such as "&", "of", "the") and single letter words are stripped from all fields, empty fields are ignored
   ii. Common Dutch words that have no relevance are stripped (e.g. "na", "bij", "rechts", "nieuwe", "het",…) this also includes medical terms that are not relevant or too general for the patient type classification (e.g. Fever: "Koorts", "Shock", lung-related: "pulmonale",…)
   iii. All remaining words in each field are used as Alphanumeric Keywords (MatchKeys) and matched with each other using a threshold of 80% (this accounts for differences in spelling or abbreviations)
   iv. The MatchKeys and Match Score is generated for all entries
4. The dataset is grouped by the generated MatchKeys, the occurrence of each Key is counted
5. All keys that occurred more than 10 times are compiled, provided to the ICU management which extracts keywords and allocates these to patient types
6. The two returned lists of keywords are used to update the ICCU/ICTH patient type keyword allocation as well as the list of words of no relevance
7. The main workflow is run again
8. The keyword generation process is repeated until all patients have been matched

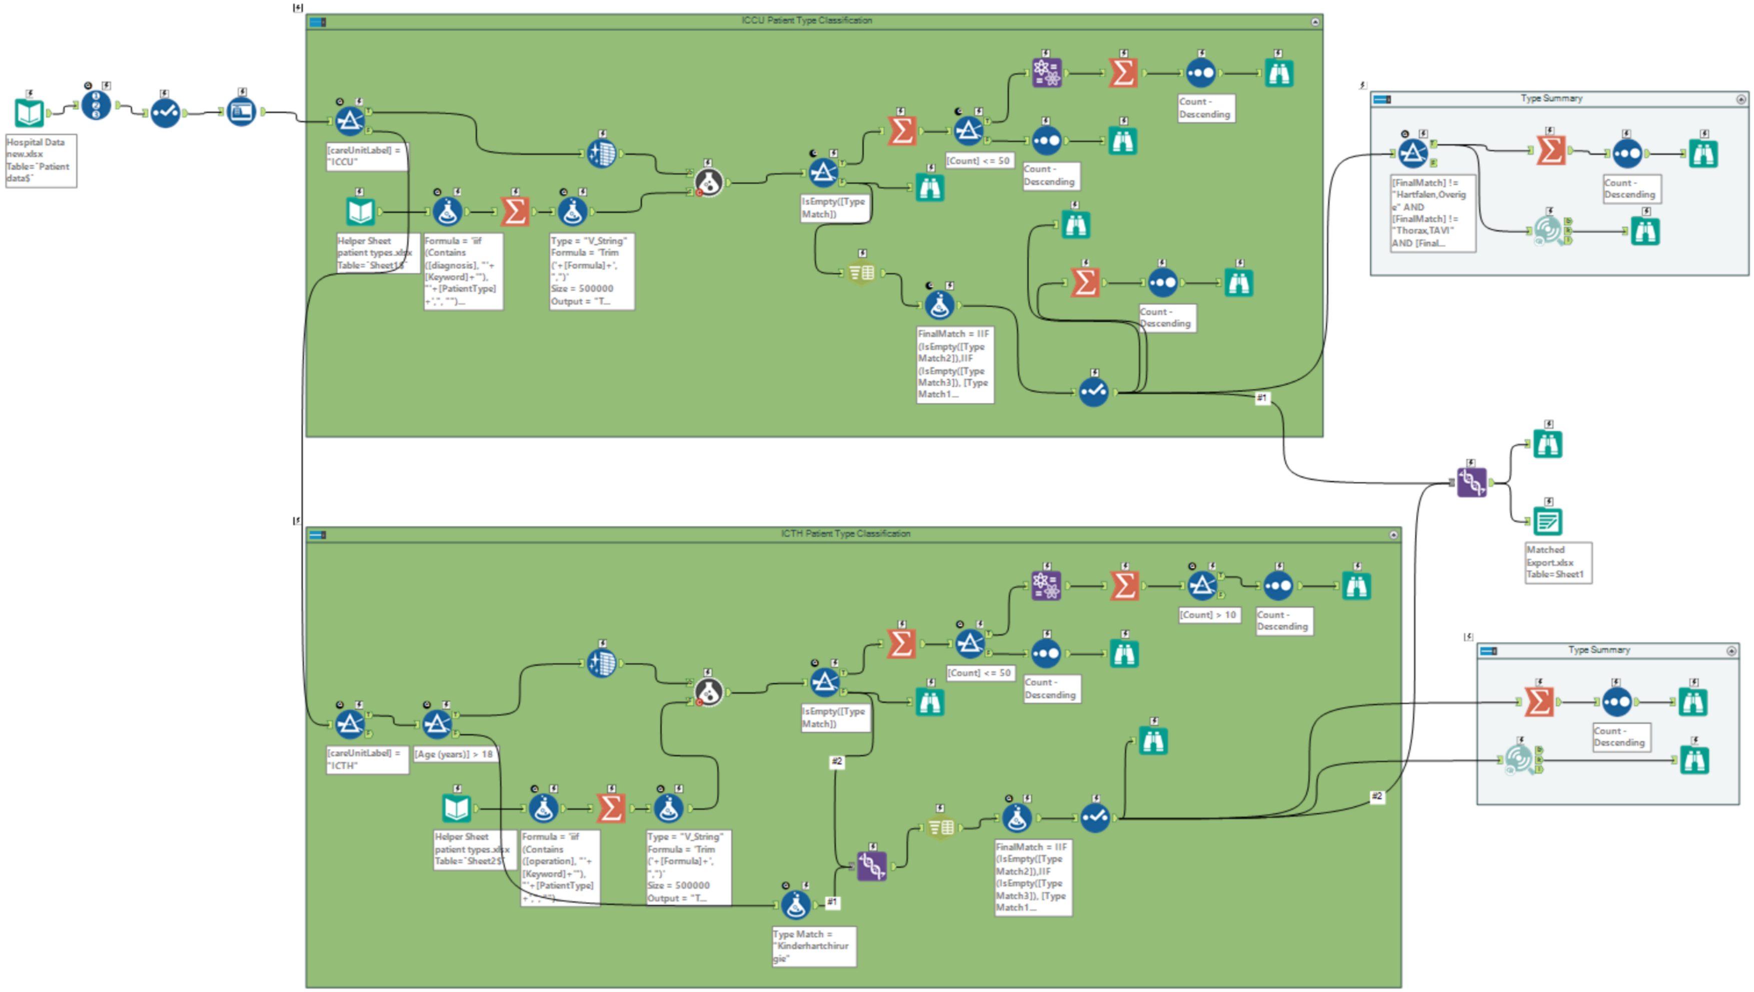

Figure B-1: Patient classification workflow in the Alteryx software

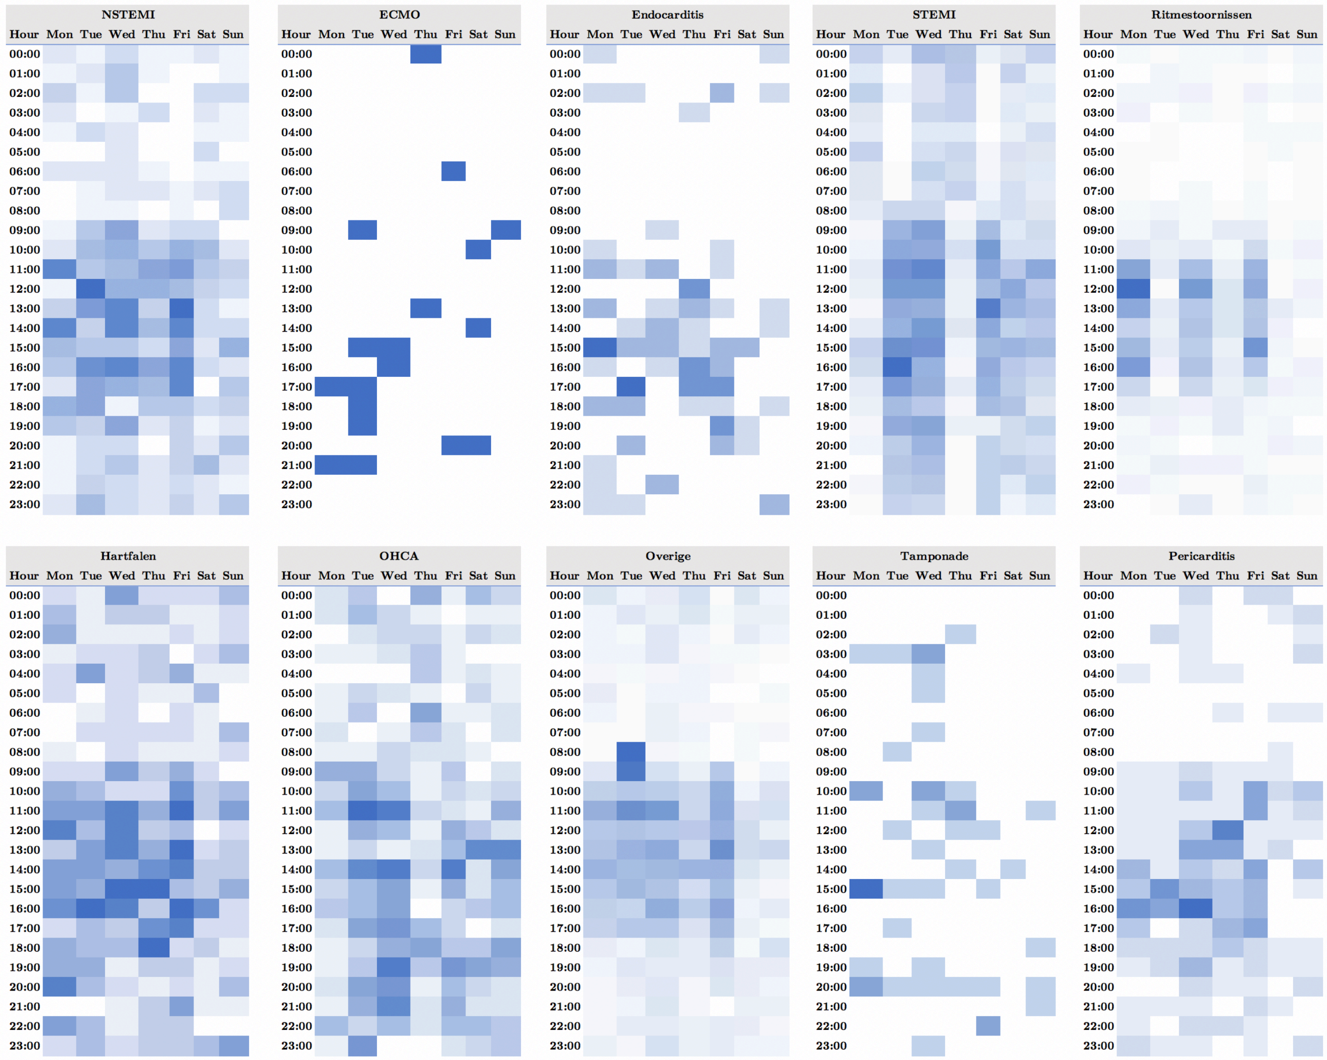

Figure C-1: Unplanned patient arrival rates

APPENDIX D. LENGTH OF STAY

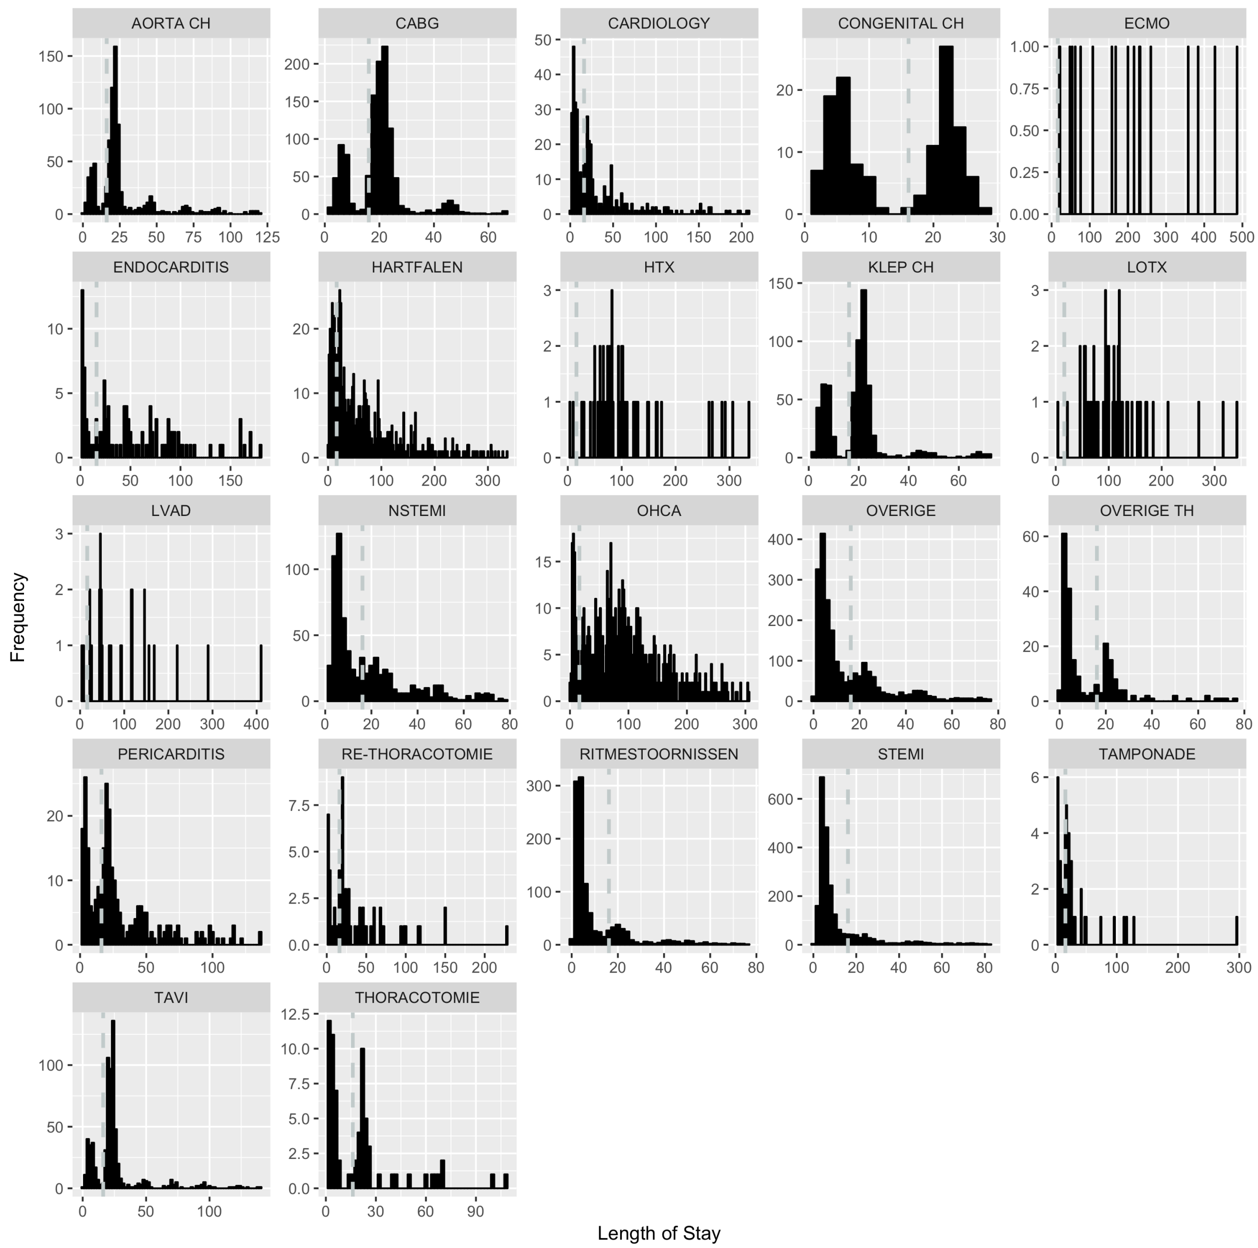

Figure D-1: LOS Histograms per patient type

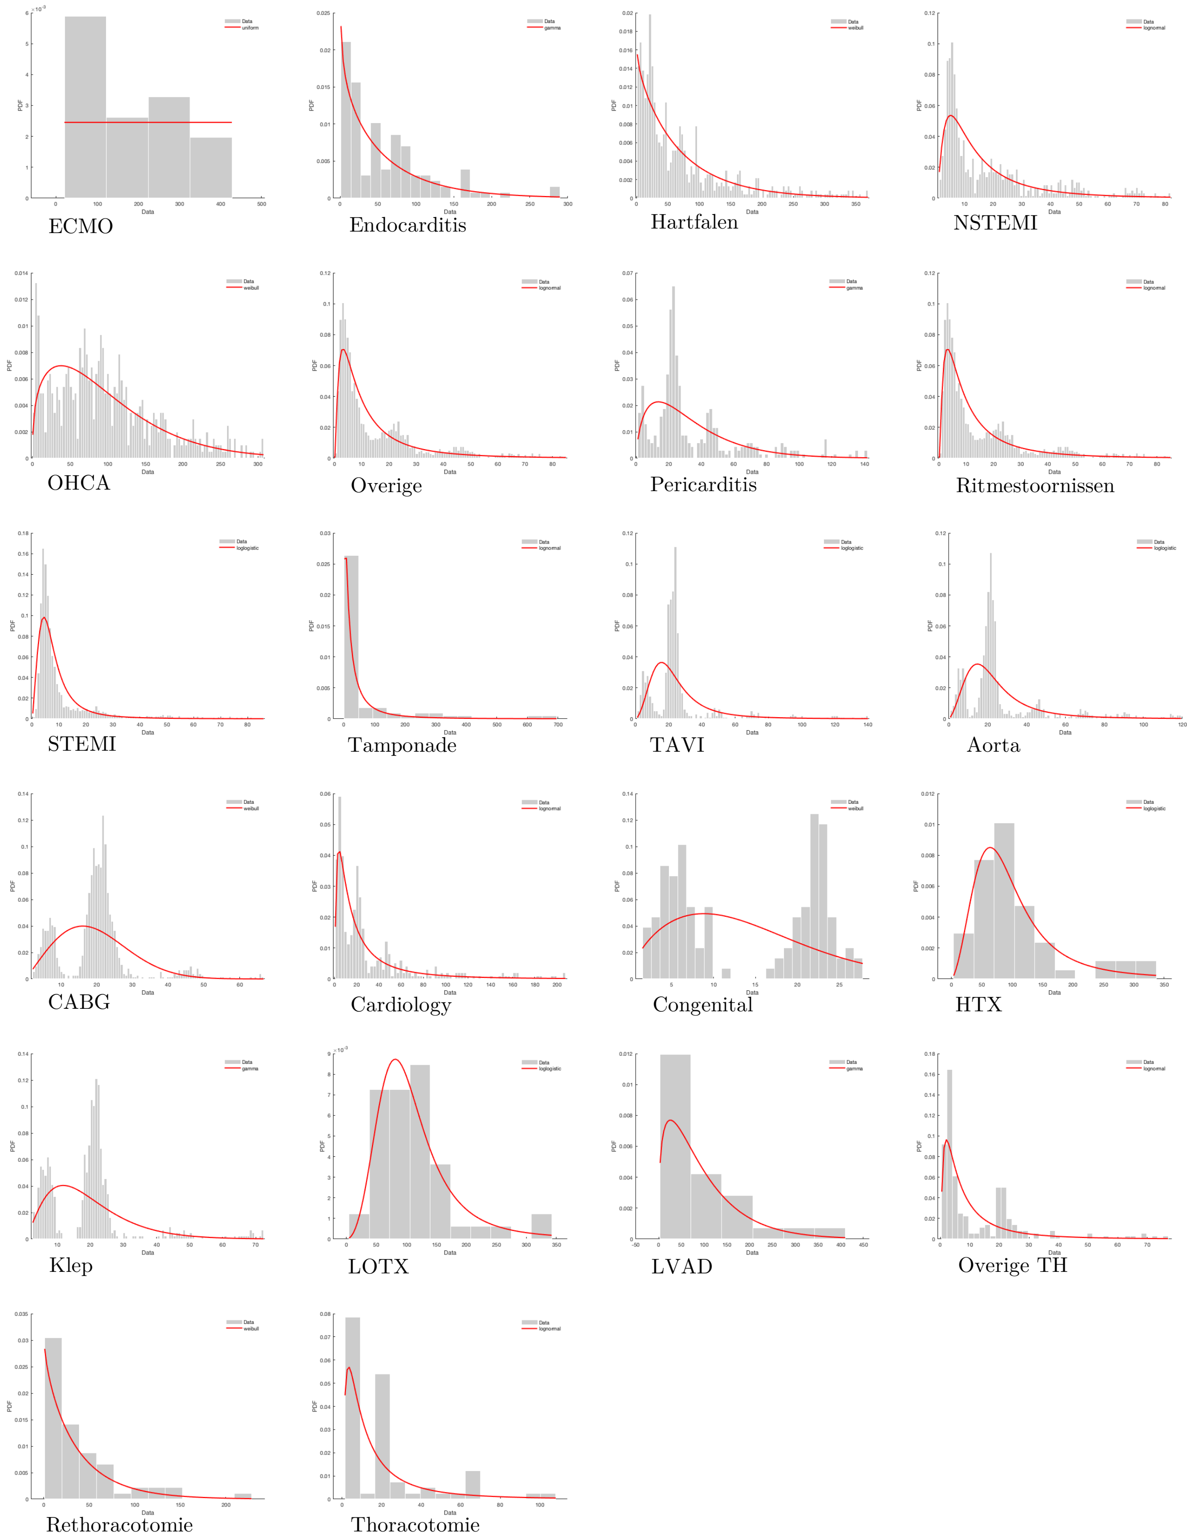

Figure D-2: Histograms with fitted distributions

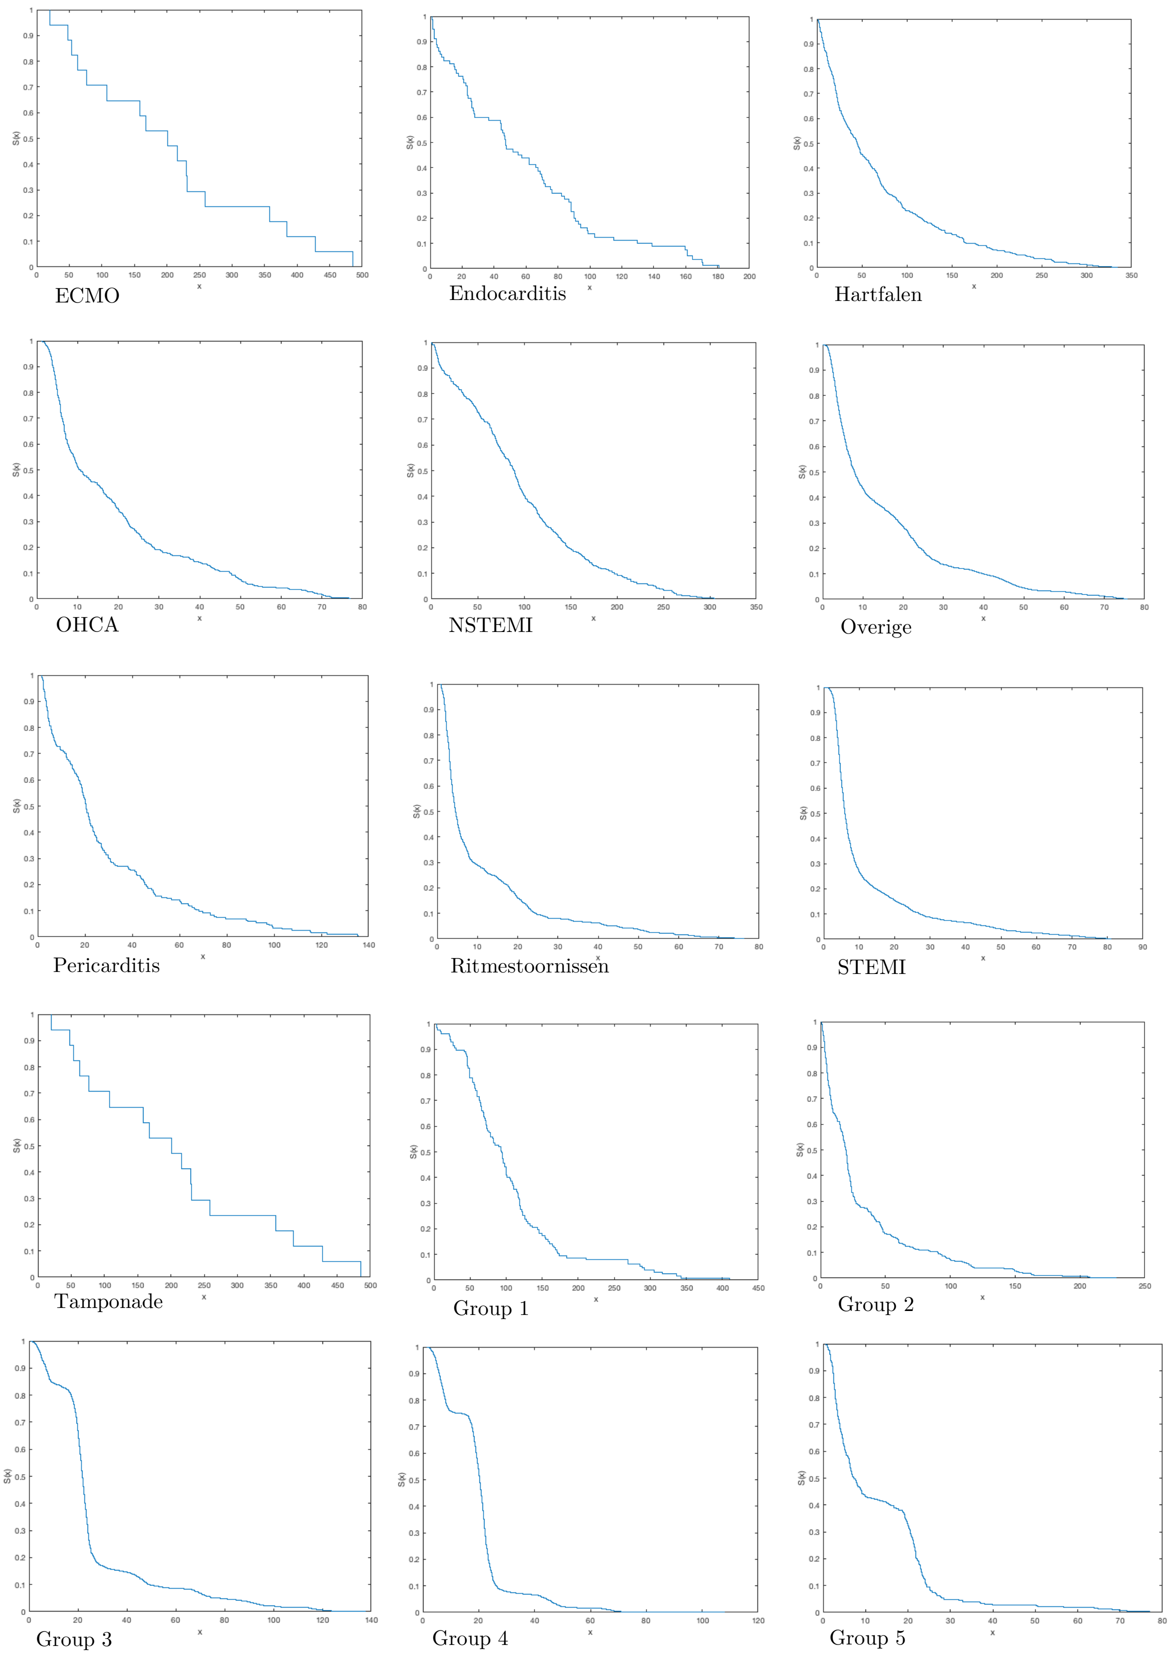

Figure D-3: Survivor functions

Days	STEMI	OHCA	Overige	Ritme-stoornisse	NSTEMI	Peri-carditis	Endo-carditis	Tampo-nade	Hartfalen	ECMO
0	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%
1	12.14%	84.03%	20.35%	10.47%	26.65%	39.51%	67.50%	30.23%	65.38%	94.12%
2	4.63%	74.33%	5.52%	4.19%	9.30%	18.54%	47.50%	16.28%	47.00%	88.24%
3	1.20%	60.08%	0.90%	0.58%	0.72%	9.27%	32.50%	10.08%	33.08%	73.53%
4	0.00%	42.97%	0.00%	0.00%	0.00%	5.37%	16.25%	7.75%	23.98%	67.65%
5	-	30.61%	-	-	-	1.46%	11.25%	4.65%	18.57%	63.73%
6	-	21.86%	-	-	-	0.00%	8.75%	0.00%	13.93%	60.78%
7	-	15.21%	-	-	-	-	3.75%	-	9.86%	52.94%
8	-	11.03%	-	-	-	-	0.00%	-	7.93%	49.41%
9	-	6.84%	-	-	-	-	-	-	6.00%	41.18%
10	-	5.13%	-	-	-	-	-	-	4.26%	27.94%
11	-	1.90%	-	-	-	-	-	-	2.51%	23.53%
12	-	0.76%	-	-	-	-	-	-	1.55%	22.06%
13	-	0.00%	-	-	-	-	-	-	0.77%	20.59%
14	-	-	-	-	-	-	-	-	0.00%	19.12%
15	-	-	-	-	-	-	-	-	-	17.65%
16	-	-	-	-	-	-	-	-	-	11.76%
17	-	-	-	-	-	-	-	-	-	8.24%
18	-	-	-	-	-	-	-	-	-	0.00%

Table D-1: Planned patient group probabilities of staying

Days	Group 1	Group 2	Group 3	Group 4	Group 5
0	100.00%	100.00%	100.00%	100.00%	100.00%
1	92.91%	34.54%	30.82%	17.37%	12.19%
2	82.68%	18.94%	10.56%	3.14%	2.50%
3	62.20%	11.42%	5.96%	0.17%	0.31%
4	46.46%	8.64%	2.46%	0.08%	0.00%
5	29.13%	3.90%	0.64%	0.00%	-
6	19.69%	3.62%	0.00%	-	-
7	12.60%	1.11%	-	-	-
8	8.66%	0.84%	-	-	-
9	7.87%	0.00%	-	-	-
10	7.20%	-	-	-	-
11	6.52%	-	-	-	-
12	5.51%	-	-	-	-
13	3.15%	-	-	-	-
14	0.00%	-	-	-	-

Table D-2: Unplanned patient probabilities of staying

APPENDIX E. SIMULATION

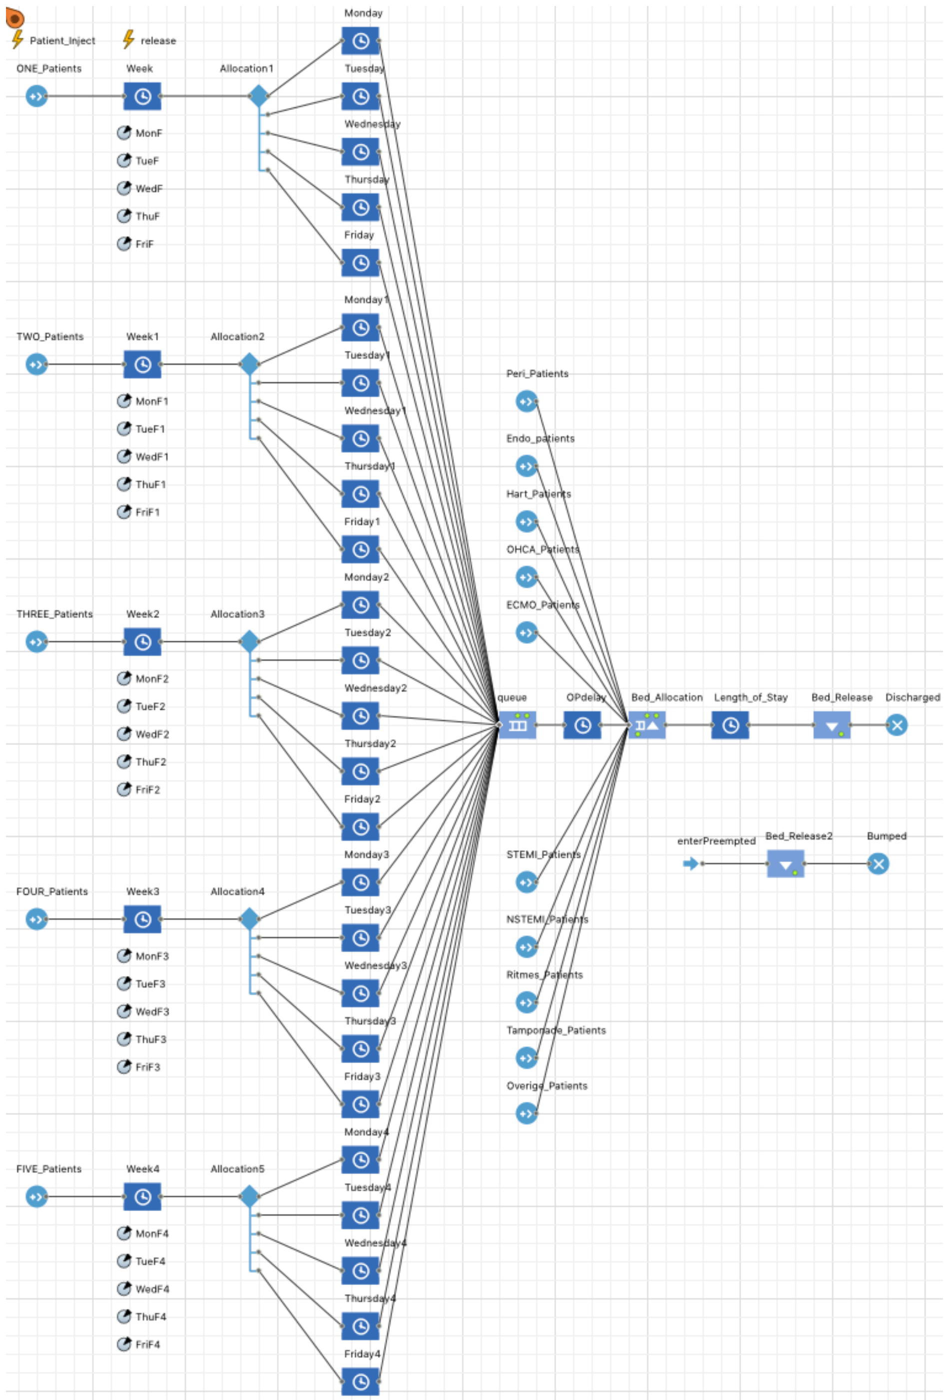

Figure E-1: Simulation routing logic

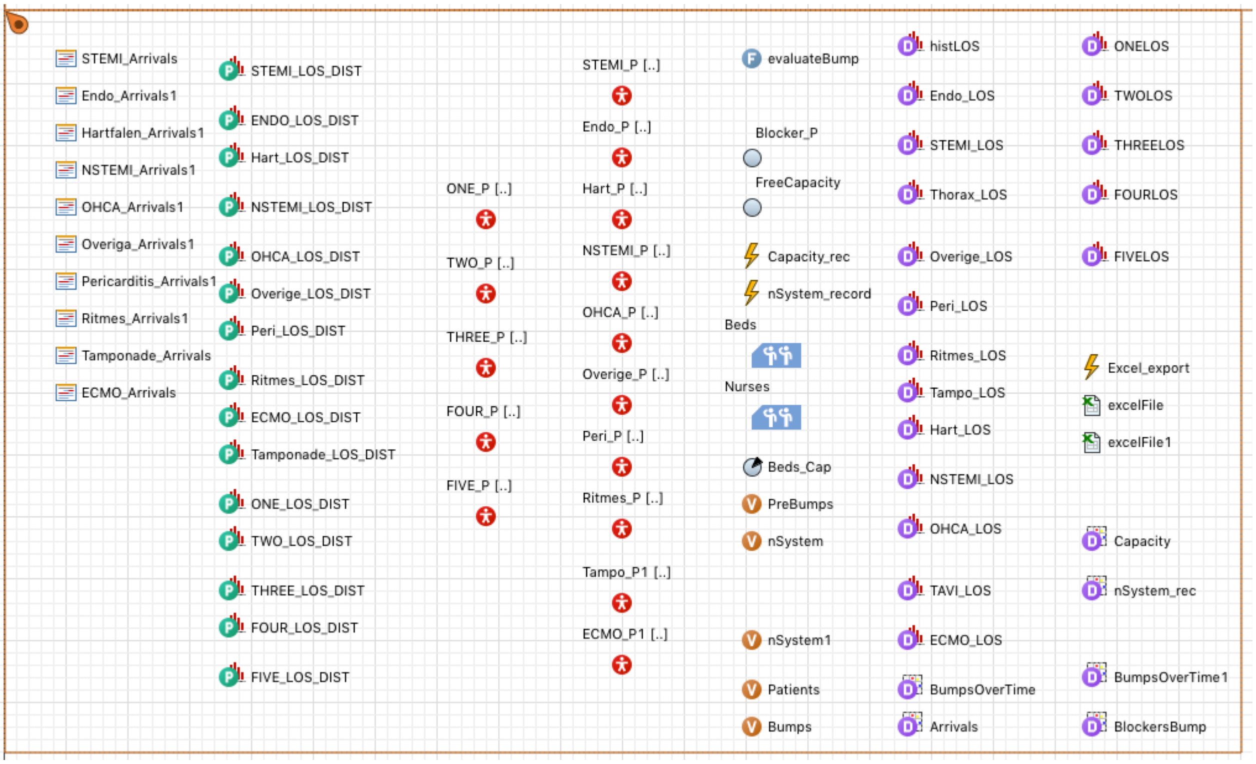

Figure E-2: Statistical backbone of the simulation

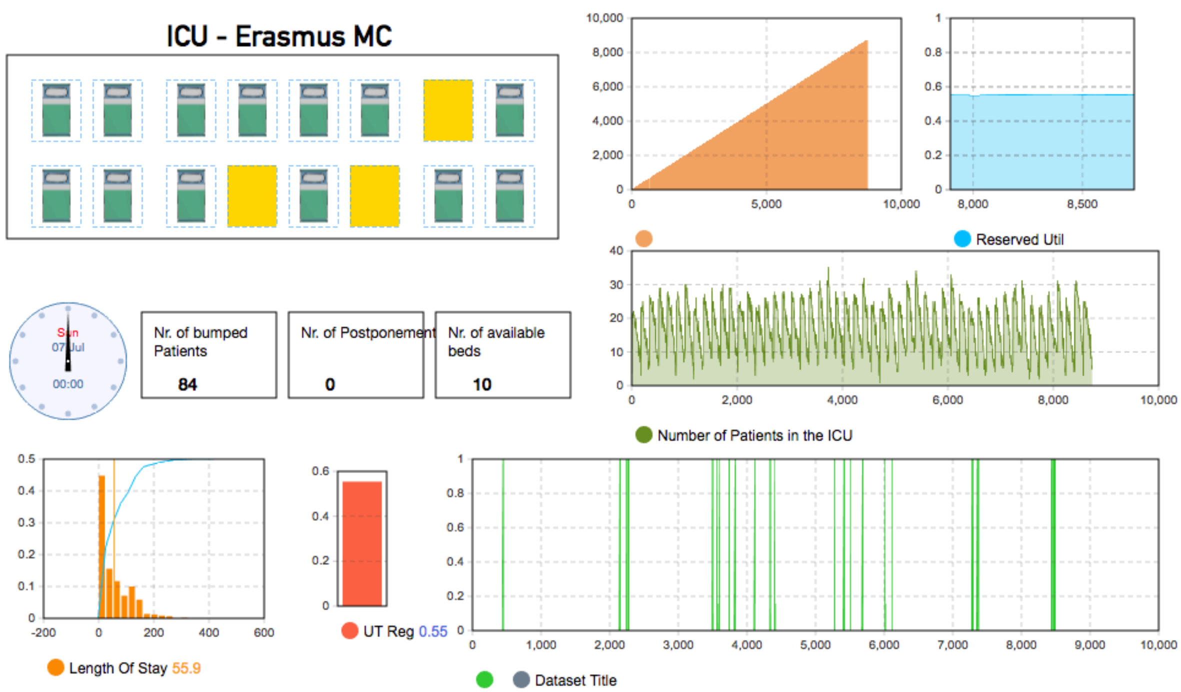

Figure E-3: Main dashboard of the simulation

REFERENCES

Allen, A. O. (2014). *Probability, statistics, and queueing theory.* Academic Press.

De Bruin, A. M., Van Rossum, A. C., Visser, M. C., & Koole, G. M. (2007). Modeling the emergency cardiac in-patient flow: an application of queuing theory. *Health Care Management Science*, *10*(2), 125-137.

Butler, T.W., Karwan, K.R., Sweigart, J.R., & Reeves, G.R., 1992, An integrative model- based approach to hospital layout, *HE Transactions,* 24(2): 144-152.

Chalfin, D. B. (1995). Cost-effectiveness analysis in health care. *Hospital cost management and accounting*, *7*(4), 1.

CIBG (2018). VWS | Jaarverslagenzorg. *Ministerie van Volksgezondheid, Welzijn en Sport.* Retrieved from https://www.desan.nl/net/DoSearch/Search.aspx?action=laatste&sub=2016

Costa, A. X., Ridley, S. A., Shahani, A. K., Harper, P. R., De Senna, V., & Nielsen, M. S. (2003). Mathematical modelling and simulation for planning critical care capacity. *Anaesthesia*, *58*(4), 320-327.

Dobrzykowski, D., Deilami, V. S., Hong, P., & Kim, S. C. (2014). A structured analysis of operations and supply chain management research in healthcare (1982–2011). *International Journal of Production Economics*, 147, 514-530.

Dobson, G., Lee, H. H., & Pinker, E. (2010). A model of ICU bumping. *Operations research*, *58*(6), 1564-1576.

Erasmus MC, (2018). Erasmus MC Website. Retrieved from https://www.erasmusmc.nl/?lang=en

Fagerland, M. W., & Sandvik, L. (2009). Performance of five two-sample location tests for skewed distributions with unequal variances. *Contemporary clinical trials, 30*(5), 490-496.

Günal, M. M., & Pidd, M. (2010). Discrete event simulation for performance modelling in health care: a review of the literature. *Journal of Simulation*, *4*(1), 42-51.

Green, L. V. (2002). How many hospital beds?. *INQUIRY: The Journal of Health Care Organization, Provision, and Financing*, *39*(4), 400-412.

Green, L. V. (2005). Capacity planning and management in hospitals. In *Operations research and health care* (pp. 15-41). Springer, Boston, MA.

Green, L. V. (2006). Queueing analysis in healthcare. In *Patient flow: reducing delay in healthcare delivery* (pp. 281-307). Springer, Boston, MA.

Griffiths, J. D., Price-Lloyd, N., Smithies, M., & Williams, J. E. (2005). Modelling the requirement for supplementary nurses in an intensive care unit. *Journal of the Operational Research Society*, *56*(2), 126-133.

Hachesu, P. R., Ahmadi, M., Alizadeh, S., & Sadoughi, F. (2013). Use of data mining techniques to determine and predict length of stay of cardiac patients. *Healthcare informatics research*, *19*(2), 121-129.

Harrison, G. W., Shafer, A., & Mackay, M. (2005). Modelling variability in hospital bed occupancy. *Health Care Management Science*, *8*(4), 325-334.

Harper, P. R. (2002). A framework for operational modelling of hospital resources. *Health care management science*, *5*(3), 165-173.

Harper, P. R., & Shahani, A. K. (2002). Modelling for the planning and management of bed capacities in hospitals. *Journal of the Operational Research Society*, *53*(1), 11-18.

Hautvast, J. L. A., Bakker, J., Boekema-Bakker, N., Faber, J. A. J., Grobbee, D. E., & Schrijvers, A. J. P. (2001). Plaats in de herberg, een studie naar determinanten van opname-en ontslagproblemen in IC-afdelingen in Nederland. *Julius Centrum voor Huisartsgeneeskunde en Patientgebonden Onderzoek, Utrecht (in Dutch).*

Heart.org. (2017). Cardiac Catheterization. Retrieved from: http://www.heart.org/HEARTORG/Conditions/HeartAttack/DiagnosingaHeartAttack/Cardiac-Catheterization_UCM_451486_Article.jsp#.WqmcOGaZNR4

Jacobson, S. H., Hall, S. N., & Swisher, J. R. (2006). Discrete-event simulation of health care systems. In *Patient flow: Reducing delay in healthcare delivery* (pp. 211-252). Springer US.

Kao, E. P., & Tung, G. G. (1981). Bed allocation in a public health care delivery system. *Management Science*, *27*(5), 507-520.

Kelly, F. (2018). Intensive care medicine is 60 years old: the history and future of the intensive care unit. Retrieved from http://www.clinmed.rcpjournal.org/content/14/4/376.long

Kim, S. C., Horowitz, I., Young, K. K., & Buckley, T. A. (1999). Analysis of capacity management of the intensive care unit in a hospital. *European Journal of Operational Research*, *115*(1), 36-46.

Kim, S. C., Horowitz, I., Young, K. K., & Buckley, T. A. (2000). Flexible bed allocation and performance in the intensive care unit. *Journal of Operations Management*, *18*(4), 427-443.

Knaus, W. A., Wagner, D. P., Zimmerman, J. E., & Draper, E. A. (1993). Variations in mortality and length of stay in intensive care units. *Annals of Internal Medicine*, *118*(10), 753-761.

Lapierre, S. D., Goldsman, D., Cochran, R., & DuBow, J. (1999). Bed allocation techniques based on census data. *Socio-Economic Planning Sciences*, *33*(1), 25-38.

Litvak, N., Van Rijsbergen, M., Boucherie, R. J., & van Houdenhoven, M. (2008). Managing the overflow of intensive care patients. *European journal of operational research*, *185*(3), 998-1010.

Moerer, O., Plock, E., Mgbor, U., Schmid, A., Schneider, H., Wischnewsky, M. B., & Burchardi, H. (2007). A German national prevalence study on the cost of intensive care: an evaluation from 51 intensive care units. *Critical care*, *11*(3), R69.

OECD (2017), "Demographic trends", in *Health at a Glance 2017: OECD Indicators*, OECD Publishing, Paris.

R Core Team (2018), R: A language and environment for statistical computing. R Foundation for Statistical Computing, Vienna, Austria. URL https://www.R-project.org/.

Ridge, J. C., Jones, S. K., Nielsen, M. S., & Shahani, A. K. (1998). Capacity planning for intensive care units. *European journal of operational research*, *105*(2), 346-355.

Ridley, S., Jones, S., Shahani, A., Brampton, W., Nielsen, M., & Rowan, K. (1998). Classification treesA possible method for iso-resource grouping in intensive care. *Anaesthesia*, *53*(9), 833-840.

Swartzman, G. (1970). The patient arrival process in hospitals: statistical analysis. *Health services research*, *5*(4), 320.

Tu, J. V., & Guerriere, M. R. (1993). Use of a neural network as a predictive instrument for length of stay in the intensive care unit following cardiac surgery. *Computers and biomedical research*, *26*(3), 220-229.

Walczak, S., Pofahl, W. E., & Scorpio, R. J. (1998). Predicting Hospital Length of Stay with Neural Networks. In *FLAIRS Conference* (pp. 333-337).

Weissman, C. (1997). Analyzing intensive care unit length of stay data: problems and possible solutions. *Critical care medicine*, *25*(9), 1594-1600.

Yoon, P., Steiner, I., & Reinhardt, G. (2003). Analysis of factors influencing length of stay in the emergency department. *Canadian Journal of Emergency Medicine*, *5*(3), 155-161.