

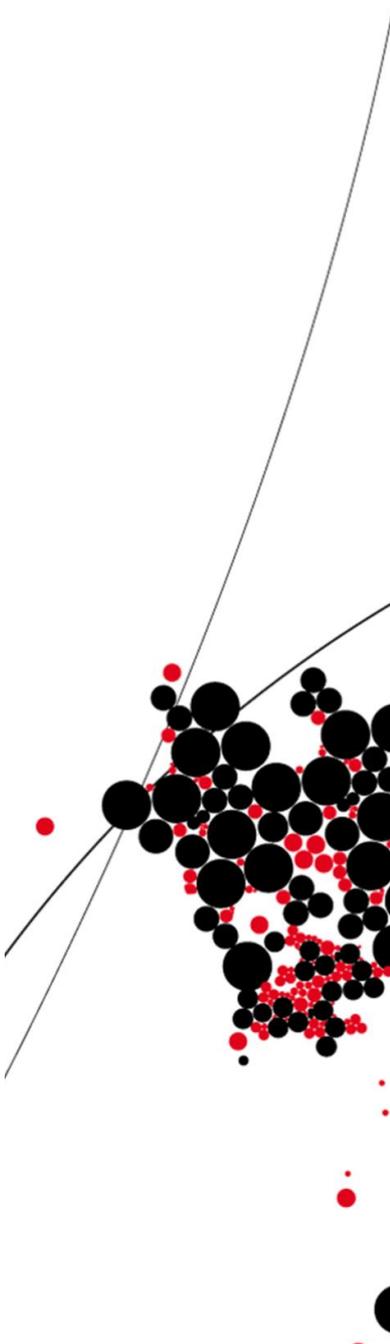


UNIVERSITY OF TWENTE.

Faculty of Science and Technology

Staffing won't be worse with one fewer nurse

**Improving the staffing of nurses in the
Intensive Care Department of the
Medisch Spectrum Twente**



**Aina Goday Verdaguer
M.Sc. Thesis
June 2017**

Supervisors:

University of Twente
Dr. P.C. Schuur
Dr. K.G.M. Groothuis-Oudschroon

Medisch Spectrum Twente
Dr. R.J Trof

Management summary

In the framework of completing the master thesis of Health Sciences, I performed research at the Intensive Care Department (ICD) of the Medisch Spectrum Twente (MST) into studying how to improve the staffing of nurses.

In January 2016 the MST moved into a new building. With this change, the ICD was expanded and the capacity increased from 28 beds -in the past situation- to a maximum capacity of 42 beds. The units also changed. In the past situation there was one Thorax and one General unit, whereas in the new situation there is one Thorax unit, two General intensive care units and one Medium care unit.

With the change of situation, the medical manager from the ICD thought about the possibility of adjusting staffing to the activity, with the purpose of increasing the efficiency, because he had the impression that there were more nurses than needed. He also had the feeling that when there are more nurses than beds to cover the motivation of the nurses decreases. Until now, the methodology used for staffing is based on the bed capacity. Thus, he wanted to study the possibility of switching to staffing based on demand, instead of capacity.

The objective of this research is to study how the current methodology employed for nurse staffing purposes can be improved. Not only do we consider staffing based on demand, but we also study other staffing approaches. This leads to the following central research question:

“ In what way can the current capacity based staffing be improved in the Intensive Care Department of the MST?”

In order to find an appropriate approach for answering the research question, we have to take into account that the ICD underwent an important expansion. There was not enough data from the new situation to extract conclusions about possible changes. Moreover, the data was not reliable. Together with the manager, the following assumption was agreed upon: if an approach would have been effective in the old ICD, it will be effective in the new ICD as well. Accordingly, this is a retrospective study using data from 2012 to 2015.

Before studying the staffing approaches, we analyzed the past situation of the ICD with data from 2012 to 2015. The most relevant information found was that the most common bed occupancy rates were between 71% and 90%.

In this project we consider three staffing approaches: staffing based on demand, based on capacity and based on a hybrid approach. Within each approach, we consider different scenarios. These approaches and scenarios are:

- (i) *Demand approach*: the demand approach consists in predicting the daily average number of beds occupied per week in 2015. The scenarios considered are: perfect prediction, naive forecast, moving averages and seasonal indexes. We also forecasted the demand using ARIMA models. Even though we found mathematically correct models, the prediction was close to the mean of the data and did not show variations. Since we were not satisfied with the results, and the application of ARIMA models is time-demanding and requires forecasting skills, we decided to develop heuristics for demand prediction. The heuristics developed are the aforementioned demand scenarios (except for the perfect prediction).
- (ii) *Capacity approach*: the capacity scenarios calculate the number of beds occupied based on a percentage of the maximum available capacity. The different scenarios considered are: 100%, 85%, 80%, 75% and 70% of capacity. For example, 85% of capacity means that we assume that 85% of the beds are occupied, even though 100% of the bed capacity is available.
- (iii) *Hybrid approach*: it is a combination of capacity and demand staffing. First, staffing is done based on a percentage of the maximum available capacity. Then, a reinforcement of this staffing is done based on demand forecasting (using seasonal indexes). This means that the hybrid approach takes the maximum value between the capacity and the demand approach. The scenarios are: 85%, 80%, 75% and 70% hybrid.

The presented scenarios were discussed from three perspectives: waste, financial and practical point of view.

From a waste viewpoint, we were looking for the scenario with the best allocation of nurses, i.e., the least amount of extra hours and unnecessary hours worked. This resulted to be the 75% capacity scenario, with a pool of nurses of 61.68 FTE. This implies 20.6 FTE less than using the 100% capacity scenario.

From a financial point of view, the 70% capacity scenario is the best one due to the reduced costs. Taking into account the pool of nurses and extra hours, the 70% capacity scenario adds up to a total amount of 1,844,296 €, whereas staffing based on 100% capacity costs 2,617,600 €. This supposes a cost reduction of 29.5%. Nevertheless, due to the lack of nurses that this scenario implies, and that the resulting nurse-patient ratios differ significantly from the optimal-defined ones, the implementation of this scenario is not wise.

So far, we have seen that the waste and financial viewpoints are not aligned. Moreover, we also have to take into account a practical perspective. In this project, reality has been simplified, thus, the scenarios mentioned so far might be too adjusted and not appropriate to be implemented. Therefore, taking into consideration the fact that changes are difficult and that reality has been simplified, we consider that a scenario closer to 100% capacity might be a good option - such as staffing based on 85% of capacity.

Answering the research question:

- (i) *Demand approach*: we consider it too risky due to the amount of extra and unnecessary hours of work that it has. Moreover, a considerable amount of reliable past data is needed, which is often difficult to obtain.

- (ii) *Capacity approach*: it is possible to improve the current staffing methodology using this approach, by reducing the percentage of capacity based on which the staffing is done. Yet, keeping all the beds open. Almost all the capacity scenarios are prepared to handle 100% of the demand, even though they do not staff based on 100% capacity. With the capacity scenarios, a cost and waste reduction can be achieved.
- (iii) *Hybrid approach*: this approach presents more waste than the capacity approach, but the risk of having less nurses than needed is reduced. Its disadvantage is the need of past reliable data.

In conclusion, we advise to consider the capacity approach and, from this, the 85% capacity scenario. We consider that further research using simulation is recommended in order to simulate a more complex environment and evaluate the performance of the scenario. In case of willing to implement the new proposed methodology, we recommend to make a real life test, in which during a month the amount of nurses working is reduced. It is important to measure the motivation of the nurses before and after the test using questionnaires, as well as to measure the test performance. This will result in valuable information in decision making regarding the change of methodology.

Even though the analysis is done based on the old ICD, taking into account the aforementioned assumption, it is possible to change the current staffing methodology in the new ICD to improve the allocation of resources.

Preface

After completing the Health Sciences course “Quantitative methods for Operations Management”, I realized that healthcare logistics and management was my field and that I wanted to devote my future to it. I approached the lecturer of this course, Peter Schuur, to find a master thesis in this field. With unconditional support, Peter helped me finding a nice topic for the thesis and, without doubting one second, we wanted to work together in this research.

After starting this project in March 2016, I decided that I wanted to enrich my knowledge in logistics and management. Hence, I decided to start the master in Industrial Engineering and Management (IEM). Unfortunately, I had not finished the thesis before starting the second master in September 2016. Therefore, some balancing with the new IEM courses and the Health Sciences thesis was needed. It turned out to be difficult to combine them, hence, the thesis took longer than expected. Nevertheless, with the support and empathy of my supervisors this was not a problem. After one year and three months, I am very glad to be able to present this thesis.

I wish to express my sincere thanks to all the people that have made this project possible. First of all, I would like to express my gratitude to Peter Schuur for his unconditional support and believe in me from the very first moment. He has not only guided me in this project, but he has enriched me as a person with all the things he has taught me and all the recommendations he has given. I would also like to thank Ronald Trof for trusting me the development of this project and his willingness to carry out research in order to improve. Furthermore, I would like to thank Marc Eijnsink for his enthusiasm, his eternal help answering all the questions I had and providing me all the data I needed. My thanks also go to Bert Beishuizen who was my contact person with the MST and the person who gave me the opportunity to carry out a project in the hospital. I also want to express my gratitude to Karin Groothuis-Oudshoorn who showed enthusiasm when I presented my research proposal, helped me with statistics and was willing to be the second supervisor of this project. Last but not least, I really appreciate the unconditional support and encouragement of my friends and, specially, of my parents. Special thanks go to Juana Borbolla, who has always been there for me motivating and helping.

Aina Goday Verdaguer
Enschede, June 2017

Contents

Management summary	i
Preface	v
List of figures	xii
List of tables	xiii
1 Introduction	1
1.1 Background information	1
1.1.1 Intensive Care Department distribution	1
1.1.2 Pool of nurses	3
1.1.3 Working shifts	4
1.1.4 Patient-nurse ratios	4
1.1.5 Nurse staffing approach	4
1.2 Problem definition	5
1.3 Stakeholder analysis	5
1.4 Research goal	6
1.5 Scope	6
1.6 Research question and sub-questions	7
1.7 Research method	8
1.8 Verification and Validation	12
1.9 Deliverables	12
2 Past situation	15
2.1 Patient admissions	15
2.2 Length of stay	17
2.3 Bed occupancy rates	18
2.4 Conclusions	20
3 Literature review	21
3.1 State of the art	21
3.2 Conclusions	24
4 Theory and concepts	25
4.1 Time series	25
4.2 ARIMA models	26
4.3 The Box-Jenkins technique	27

4.4	Forecast accuracy	27
4.5	Conclusions	28
5	Feasibility of demand prediction	29
5.1	Statistical prediction	29
5.1.1	Thorax unit	30
5.1.2	General unit	37
5.2	Heuristics prediction	40
5.2.1	Description of the approaches	41
5.2.2	Comparison with ARIMA models	43
5.2.3	Comparison among the heuristics	43
5.3	Conclusions	44
6	Desired magnitude of the pool of nurses	45
6.1	Overview of the approaches, scenarios and steps	45
6.2	Staffing based on demand	47
6.3	Staffing based on capacity	48
6.4	Staffing based on demand and capacity	49
6.5	Results and discussion	50
6.5.1	Definition of the viewpoints	50
6.5.2	Thorax unit	51
6.5.3	General unit	53
6.5.4	Intensive Care Department	54
6.6	Conclusions	55
7	How to achieve cost savings	57
7.1	Fees and budget	57
7.2	Costs results	58
7.3	Comparison with the ICD budget	59
7.4	Conclusions	59
8	Discussion	61
8.1	General discussion	61
8.1.1	Comparison of our results to the literature	61
8.1.2	Consideration of other staffing approaches	62
8.1.3	Waste, financial and practical viewpoints	62
8.1.4	Comparison of the pool of nurses needed	63
8.1.5	Forecasting an unknown year	64
8.1.6	Implementation	64
8.2	Limitations	66
8.3	Further research	66
8.4	Conclusions	66
9	Conclusions and recommendations	69

Appendix	75
Appendix A	75
Appendix B	78
Appendix C	82
Appendix D	85

List of Figures

1.1	ICD distribution before and after the change of building and variation of the number of open beds. Black beds: maximum capacity. White beds: operational capacity.	2
1.2	Stakeholder typology (Mitchell et al., 1997)	6
2.1	Total patient admissions per unit and year	15
2.2	Percentage of patient admissions per unit depending on the origin (coming from the operating room or not)	16
2.3	Percentage of patient admissions per unit and day of the week	16
2.4	Length of Stay histogram for the Thorax unit	18
2.5	Length of stay histogram for the General unit	18
2.6	Daily average occupancy rates in the Thorax unit for each year from 2012 to 2015	19
2.7	Daily average occupancy rates in the General unit for each year from 2012 to 2015	20
5.1	Daily average of beds occupied per week in the Thorax unit	30
5.2	Time series decomposition for the Thorax unit	31
5.3	Correlation plot (ACF) (left) and partial correlation plot (PACF) (right) of the weekly data in the Thorax unit	32
5.4	Correlation plot (ACF) (left) and partial correlation plot (PACF) (right) of the differenced weekly data in the Thorax unit	33
5.5	Forecast from ARIMA (1,1,2) model for the Thorax unit	37
5.6	Forecast from ARIMA (1,1,1)(1,0,0) ₅₂ model for the Thorax unit	37
5.7	Forecast for the third and fourth quarters of 2015 for the Thorax unit. Black line: original data. Red line: ARIMA (1,1,2). Green line: ARIMA(1,1,1)(1,0,0) ₅₂ . Dashed lines: 80% confidence intervals.	37
5.8	Daily average per week of the number of beds occupied in the General ICU	38
5.9	Time series decomposition for the General unit	38
5.10	Forecast from ARIMA (1,0,0) model for the General unit	39
5.11	Forecast from ARIMA (1,0,0)(1,0,0) ₅₂ ^c model for the General unit	40
5.12	Forecast for the second quarter of 2015 for the General unit. Black line: original data. Red line: ARIMA(1,0,0) ^c . Green line: ARIMA(1,0,0)(1,0,0) ₅₂ ^c . Dashed lines: 85% confidence intervals.	40
5.13	Example of the naive forecast for next week	41
5.14	Example of the naive forecast in five weeks ahead	42
6.1	Example of a demand scenario (seasonal indexes) for the Thorax unit	48

6.2	Example of a capacity scenario (85% of capacity) for the Thorax unit	49
6.3	Example of a hybrid scenario (85% of capacity and 15% seasonal indexes) for the Thorax unit	49
6.4	Summary of the nurse hours, pool of nurses, extra and unnecessary hours and the waste obtained with the different scenarios for the Thorax unit in 2015. OVC= Overcapacity. Inc= incapacity to cover demand	51
6.5	Summary of the nurse hours, pool of nurses, extra and unnecessary hours obtained with the different scenarios for the General unit in 2015. OVC= Overcapacity. Inc= incapacity to cover demand	53
6.6	Summary of the nurse hours, pool of nurses, extra and unnecessary hours obtained with the different scenarios for the ICD. OVC= Overcapacity. Inc= incapacity to cover demand	54
7.1	Summary of the costs for each scenario in the ICD in 2015	58
9.1	Thorax unit occupancy rates depending on the shift (sum of all the years) . . .	78
9.2	General unit occupancy rates depending on the shift (sum of all the years) . . .	78
9.3	Number of beds occupied in the Thorax unit in 2012 depending on the shift . .	79
9.4	Number of beds occupied in the Thorax unit in 2013 depending on the shift . .	79
9.5	Number of beds occupied in the Thorax unit in 2014 depending on the shift . .	79
9.6	Number of beds occupied in the Thorax unit in 2015 depending on the shift . .	80
9.7	Number of beds occupied in the General unit in 2012 depending on the shift . .	80
9.8	Number of beds occupied in the General unit in 2013 depending on the shift . .	80
9.9	Number of beds occupied in the General unit in 2014 depending on the shift . .	81
9.10	Number of beds occupied in the General unit in 2015 depending on the shift . .	81
9.11	Correlation plot (ACF) (left) and partial correlation plot (PACF) (right) of the weekly data in the General ICU	82
9.12	Number of FTE needed per week per unit in case of implementing the 85% capacity scenario in August	85

List of Tables

1.1	Capacity based staffing in the new ICD (based on an operational capacity of 36 beds)	5
1.2	Example of how the program chooses the maximum, mode, or mean according to the criteria explained	10
2.1	Length Of Stay analysis	17
5.1	Non-seasonal ARIMA models for the Thorax unit. The smaller the AIC_c and the $RMSE$, the better. A high p-value is desired because it indicates that the residuals behave like white noise.	34
5.2	Seasonal ARIMA models for the Thorax unit. The smaller the AIC_c and the $RMSE$, the better. A high p-value is desired because it indicates that the residuals behave like white noise.	35
5.3	RMSE values for the different methodologies used to forecast. Night shift from 27 th to 52 nd week of 2015	43
5.4	RMSE values for the naive, moving averages and seasonal indexes forecasts. Period from 1 st January to 31 st December 2015. Av.=Average.	43
6.1	Nurse-patient ratios' thresholds	47
6.2	Ratios nurse-patient in the Thorax unit for the situation when there is 100% occupancy and the nurses are calculated with different scenarios	52
6.3	Ratios nurse-patient in the Thorax unit for the situation when there is 100% occupancy and the nurses are calculated with different scenarios	54
7.1	Average fees for normal and extra hours	57
9.1	Non-seasonal ARIMA models for the General unit. The smaller the AIC_c and the $RMSE$, the better. A high p-value is desired because it indicates that the residuals behave like white noise.	83
9.2	Seasonal ARIMA models for the General unit. The smaller the AIC_c and the $RMSE$, the better. A high p-value is desired because it indicates that the residuals behave like white noise.	84

Chapter 1

Introduction

In the framework of completing the master thesis of Health Sciences, I performed research at the Intensive Care Department (ICD) of the Medisch Spectrum Twente (MST) into studying if demand nurse staffing is feasible and which other ways of staffing could be used in order to make the ICD more efficient.

In this chapter we make an overview of the organization and way of working in the Intensive Care Department (Section 1.1), we explain the problem we want to tackle (Section 1.2) and the stakeholders involved (Section 1.3). This is followed by the goal of the research (Section 1.4), the scope (Section 1.5) and the research question and sub-questions (Section 1.6) that we set out to carry out the project. Afterwards we explain the methods used for the development of the study (Section 1.7), the verification and validation (Section 1.8) and the documents to deliver as an output of the study (Section 1.9).

1.1 Background information

The MST hospital in Enschede moved to a new building the past January 2016. It was not only a change of place, but it also implied various changes in the organization, including the extension of the ICD.

1.1.1 Intensive Care Department distribution

With the new situation, the ICD has experienced an increase of the capacity. The old ICD had a total capacity of 28 beds whereas the new one has a maximum of 42 beds. Figure 1.1 shows the distribution of units and beds of the ICD before and after the change of building. No beds were reserved for emergency patients, but no emergency patients were refused because the best patient in the ICD was transferred to a ward.

The old ICD consisted of two units, the Thorax and General units. In the Thorax unit mainly patients coming from thorax surgeries were treated, whereas in the General unit patients came from all over the specialties. The Thorax unit had a maximum capacity of 10 beds and the General unit of 18 beds. The number of beds opened varied depending on the shift, day of the week and the month of the year (Figure 1.1). For the Thorax unit, there was a capacity of 10 beds during the weekdays (from Monday day shift to Saturday day shift) and a capacity of 8 beds during the weekends (from Saturday evening shift to Monday night shift) and all July and August. The capacity in the General unit only changed in summer, where in July

and August the capacity became 17 beds instead of 18 like during the rest of the year.

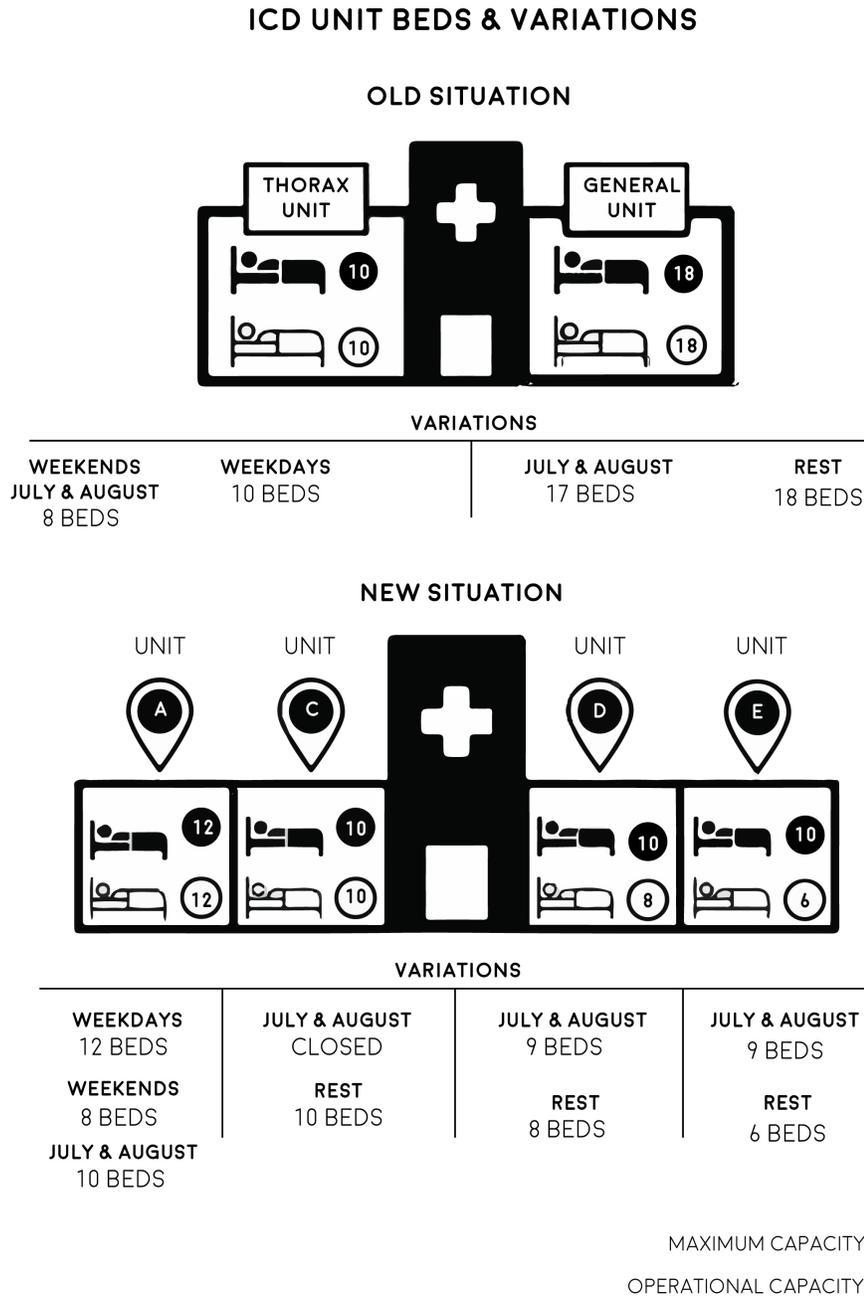


Figure 1.1: ICD distribution before and after the change of building and variation of the number of open beds. Black beds: maximum capacity. White beds: operational capacity.

In the new situation, we can find 4 units (Figure 1.1): Thorax Intensive Care Unit A, General Intensive Care Unit D, General Intensive Care Unit E and Medium Care Unit C. The thorax unit is equivalent to the old one with the difference that now it has more beds. The General units are also the same as the old one but the total capacity is higher. Apart from the Intensive Care Units (A, D and E), a fourth unit has been added which treats less severe patients: the

Medium Care Unit C. Although the new ICD has a maximum of 42 beds (12 in unit A, 10 in unit C, 10 in unit D and 10 in unit E), based on the total amount of employable nurses (in 2016) the operational capacity is limited to 36 beds (12 in unit A, 10 in unit C, 8 in unit D and 6 in unit E). In the new situation, the number of beds open also depends on the shift, day of the week and month of the year. Unit A, during the weekdays, has up to 10 beds reserved for intensive care thorax patients and 2 for medium care thorax patients. During the weekends, only 8 beds are open and during summer holidays, 10 beds are open. Unit C is closed during summer (July and August), whereas units D and E have each 9 beds open.

1.1.2 Pool of nurses

Types of nurses

Nurses in the ICD not only give direct care to patients, but some of them also have other duties. For example, there are 3 leader nurses that devote 50% of their time to management tasks and 50% to direct care. There are other nurses that have innovation tasks, in which they participate in new projects. Also, some nurses are research nurses and, thus, devote a lot of their time into medical research. Other nurses are more dedicated to teaching. For instance, there are ventilation practitioners that, apart from giving direct patient care, give instructions to nurses about ventilation (how to use the machines, how to apply the techniques...). Other teaching nurses take care of showing specific techniques such as how to perform an infusion, how to place a catheter, etc. All this implies that in order to calculate the number of nurses needed, not only the direct patient care hours have to be taken into account, but also the hours devoted to these other duties. However, for simplicity, in this project we only consider full-time employees devoted to direct care.

Organization of the nurses

Before 2012, the Thorax unit had its own dedicated pool of nurses. The work was very specific, so not a lot of nurses wanted to be dedicated only to thorax patients. In order to have nurses working in the Thorax unit, the managers decided to increase the thorax nurses' salary. In 2012 they came up with the idea of introducing a rotating system in order to overcome the problem of having to motivate the nurses to work in the thorax unit for more money. With this system all the nurses worked either in the Thorax and the General unit. The goal was to have a strong pool of nurses motivated to work in both units without an increase of salary.

The rotating system was the same in the old and in the new situation. The only difference was an increase of nurses for the new situation. Due to some issues, it is not possible to know exactly the organization of the pool of nurses in 2015. Therefore, we explain the working system in the new situation.

In the rotating system, the IC nurses are divided in 12 groups, which each of them adds up an amount of approximately 300 hours, and they rotate once in a quarter. This means that each unit has 4 groups assigned and after a quarter, two groups of each unit move to another unit. The groups that move are always the ones that have been assigned to the unit the longest. The MC unit always has its specialized nurses, who have received a specific medium care training.

1.1.3 Working shifts

Direct patient care is given in three different shifts during the day. These are:

- Night: 22:45h - 7:45h (9 hours)
- Day: 7:30h - 15:30h (8 hours)
- Evening: 15:00h - 23:00h (8 hours)

During the overlapping times, nurses from the different shifts share information to keep each other updated about the patient's situation. Moreover, every morning at 11:00 they have a meeting with the doctors to inform about how the patient went through the night and to be aware of the new decisions that doctors make.

1.1.4 Patient-nurse ratios

According to the guidelines, there are specific ratios about the number of patients that one nurse has to take care of. This should be fulfilled in order to have a good quality of care. The ratios vary in function of the shift and type of care:

- Intensive Care Unit
 - Night: 2 patients per 1 nurse
 - Day: 1.5 patients per 1 nurse
 - Evening: 1.75 patients per 1 nurse
- Medium Care Unit: 3 patients per 1 nurse (regardless of the shift)

1.1.5 Nurse staffing approach

The nurse staffing approach that has always been used in the ICD is based on the bed capacity. Even though in the new situation it is still based on capacity, they have slightly modified the approach.

In the old situation they calculated the nurses needed (for direct care) based on the ratios patient-nurse and the beds they had. This means that they calculated the nurses needed for each day of the week for each shift, based on the patient-nurse ratios, thinking that all the beds were going to be occupied. Then, they added up all the numbers and this resulted to be the number of nurses needed for one week. Then, the calculation is done for how many FTE are needed in one year.

In the new situation, the principle is the same but the methodology is slightly changed. Now, they divide the beds according the patient levels. There are three patient levels: level 1, level 2 and level 3, where level 3 are the most acute patients. They reserve a specific amount of beds for each patient level. Also, they specify how many nurses are needed per bed depending on each level. Multiplying the number of beds of each type by the nurses needed in each level and adding the results, they obtain the nurses needed. Table 1.1 shows how they make the calculation based on an operational capacity of 36 beds. They need a total of 126.5 FTE for direct care. It is important to remark that we have noticed that the sum of the number of beds reserved for each level does not add up to 36 but 35. However, this was the data given.

<i>Level</i>	<i>Beds</i>	<i>Nurses per bed</i>	<i>Total</i>
Level 3	16	4.2	67.5
Level 2	10	3.5	35
Level 1	9	2.7	24.3
Total			126.5

Table 1.1: Capacity based staffing in the new ICD (based on an operational capacity of 36 beds)

1.2 Problem definition

With the change of situation, the medical manager from the ICD thought about the possibility of adjusting staffing to the activity, with the purpose of increasing the efficiency, because he had the impression that there were more nurses than needed. He also had the feeling that when there are more nurses than beds to cover the motivation decreases, which obviously is not good.

Until now, the methodology used for staffing procedures - i.e., determining the appropriate number of nurses to employ- and scheduling was based on the maximum bed capacity. By means of this approach they are subjected to the possibility of being overstaffed in the sense of having more nurses than beds to cover because demands fluctuates over time. Therefore, the medical manager wanted to study the possibility of switching to staffing and scheduling based on demand, instead of capacity. If such a transition proved to be possible, he would expect it to minimize the number of nurses needed as well as costs, without neglecting the quality of care, in order to optimize the performance of the ICD.

Staffing and scheduling based on demand involves the use of forecasting methods to predict demand. Demand prediction in health care services is valuable to improve the allocation of human and physical resources and strategic planning and, thus, to match staffing to activity (Champion et al. (2007); Wargon et al. (2009)). At a micro level it can help staff scheduling and at a macro level it is useful for the financial and strategic planning of the hospital (Champion et al., 2007).

1.3 Stakeholder analysis

Of course, the problem involved is intrinsically complex, since it involves many stakeholders. Mitchell et al. (1997) defines three attributes in order to make a stakeholder classification. These attributes are:

1. *Power*: a relationship among social actors in which one social actor, A, can get another social actor, B, to do something that B would not have otherwise done.
2. *Legitimacy*: a generalized perception or assumption that the actions of an entity are desirable, proper, or appropriate within some socially constructed system of norms, values, beliefs, definitions.
3. *Urgency*: the degree to which stakeholder claims call for immediate attention.

The various combinations of these attributes result in 7 stakeholder classes (Figure 1.2).

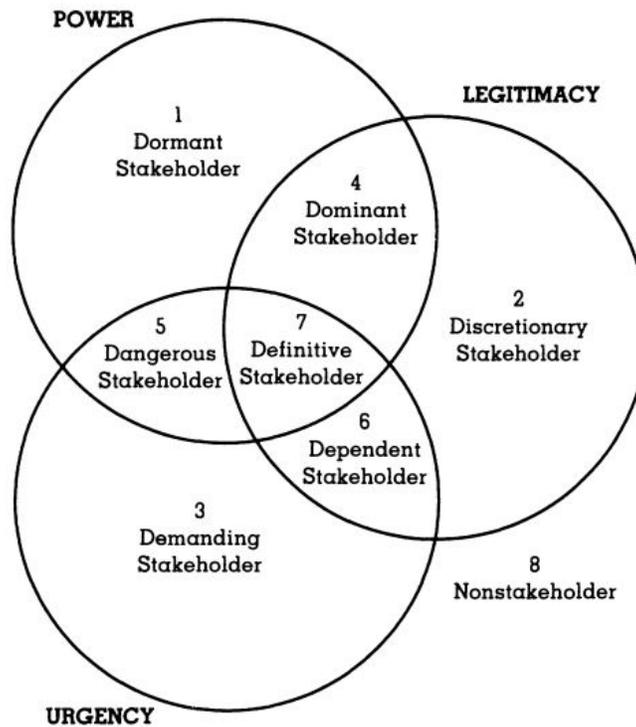


Figure 1.2: Stakeholder typology (Mitchell et al., 1997)

The stakeholder of this project are:

- Ronald Trof: as an intensivist and medical manager of the ICD, he is the problem owner and, thus, a definitive stakeholder.
- Bert Beishuizen: he is an intensivist of the ICD and he was the contact person for this project. Therefore, he is also a definitive stakeholder.
- Marc Eijsink: as head of the nurses in the ICD, he is a definitive stakeholder as well.
- Nurses: nurses in this project are dependent stakeholders.

1.4 Research goal

The aim of this project is to study how the current methodology employed for nurse staffing purposes can be improved. Namely we consider staffing based on demand but we also look for other approaches and the economic implications that come with all of them. This will allow us to come up with suggestions to improve the staffing of the nurses in the Intensive Care Department.

1.5 Scope

Hans et al. (2012) propose a framework for health care planning and control that integrates four managerial areas (medical planning, resource capacity planning, materials planning and

financial planning) and four hierarchical levels (strategic, tactical, offline operational and online operational). On the one hand, strategic planning addresses structural decision making, which has a long planning horizon. On the other hand, offline operational planning involves short-term decision making related to the execution of health care delivery processes. In between these two levels we find the tactical level, in which decisions are made at a longer term than the operational level and at a shorter term than the strategic level. Regarding personnel resources, strategic planning entails workforce planning, offline operational planning involves workforce scheduling and, thus, since tactical planning is in between, it entails staffing. Therefore, this project implies a change at the tactical level. Hence, the scope of this project is staffing, which involves the determination of the number of nurses to employ. Scheduling of the nurses is beyond the scope.

1.6 Research question and sub-questions

In order to achieve the goal of the project, we want to answer the following research question:

“ In what way can the current capacity based staffing be improved in the Intensive Care Department of the MST?”

Research sub-questions are formulated so as to clarify the research objective. These are:

1. *What has been the situation of the ICD regarding patient admissions and length of stay from 2012 to 2015?*

An important aspect in understanding the past situation is to have an overview of the health care demand in the ICD by the patients. Important factors we analyze help us to have such vision. These are patient admissions (analyzing them depending on the unit, day of admissions and from where the patient comes) and the length of stay. We address this question in Chapter 2.

2. *Which are the bed occupancy rates of the Thorax and General units from 2012 to 2015?*

Another important factor that gives us insight about the demand for health care and the utilization of the ICD is the bed occupancy rates. This factor will allow us to have a first idea of whether a demand based staffing approach would make sense or not. This is also analyzed in Chapter 2.

3. *What does literature say about demand forecasting in the health care sector?*

In order to study the manager’s idea to staff based on demand it is necessary to perform demand forecasting. We check on the literature how other researchers have tackled demand forecasting problems in the health care sector and what their results have been. The literature review is done in Chapter 3.

4. *Which other methodologies can be used for staffing?*

Apart from focusing on demand based staffing, it is important to study whether other methodologies can contribute to the improvement of the current capacity based staffing. This implies that we do not limit our selves to study just one possible solution (staffing based on demand), but to consider more methodologies and to assess the advantages and disadvantages of each of them. We carry out the assessment of these other methodologies in Chapter 6.

5. *How accurately can the demand of the Intensive Care Units be predicted?*

If we want to study an approach of staffing based on demand, it is necessary to assess how good the demand can be forecasted. In Chapter 5 we determine it by comparing the prediction with the actual demand there was.

6. *How large is the pool of nurses needed in the ICD based on demand, capacity and the hybrid approach?*

The project sets out the possibility of staffing using a different method from the current one. Hence, determining the appropriate number of nurses to employ is one relevant aspect in the project. We calculate the magnitude of the pool of nurses based on different staffing approaches and scenarios. We analyze the differences in Chapter 6.

7. *Which are the cost differences between staffing based on capacity, demand and the hybrid approach?*

It is essential to assess the impact of staffing, using different approaches, in economical terms. This allows us to evaluate whether or not each approach and scenario is efficient and if it would be worth to implement it. Chapter 7 deals with the financial part.

8. *Which are the waste, financial and practical viewpoints of the different methodologies considered for staffing?*

In order to make suggestions on how to improve the current staffing methodology, it is important to assess the alternative methodologies proposed based on different points of view. This allows us to make sure that we do not only consider a cost-savings staffing method, but one that is also appropriate to implement in terms of patient assistance quality and management of resources. The viewpoint are defined in Chapter 6 and we carry out the assessment in Chapter 8

1.7 Research method

Study design

In order to find an appropriate way for answering the research question, we have to take into account that the ICD underwent an important expansion in 2016 and, thus, the current situation diverges from the original (Figure 1.1). The only available data of the new ICD is from January 2016 to March 2016, which is not enough to analyze the new situation and extract conclusions about possible changes. Apart, for our purposes, statistical methods for demand forecasting cannot be used with this little amount of data. Instead, we should use judgmental forecasting, but this approach does not meet the manager's interests. Moreover, during our research, the only available data from 2016 turned out to be not reliable since we found an error in the hospital information system. Therefore, it is not possible to determine how to improve the current staffing methodology using data from the new ICD. To overcome this problem, together with the manager, the following assumption was agreed upon: if an approach would have been effective in the old ICD, it will be effective in the new ICD as well. Accordingly, this will be a retrospective study using data from 2012 to 2015 extracted from the hospital information system.

Data

As mentioned above, we base our study on the old ICD, which consisted of a Thorax unit and a General unit. The data used in this study is validated by the hospital and covers

the period from 1st January 2012 to 31st December 2015. It is daily data that consists of: patient number, unit (Thorax unit or General unit), admission date, admission time, origin (admission from the operating room or not), discharge date and discharge time.

The data in this form is only useful to answer the first research sub-question but not for the remaining. For the other sub-questions we will consider the number of beds occupied. This is because when planning and staffing based either on capacity or demand, the number of nurses needed does not depend on the number of patients that pass every day through the ICD but on the number of beds. Concretely, when based on capacity, the pool of nurses depends on the total capacity or beds opened, whereas when based on demand, the number of nurses does not directly depend on the number of patients admitted during the day but on the number of beds occupied at the same time. To make it easier to understand, let's make an example. Considering the staffing based on demand, let's imagine the case in which each nurse takes care of one bed. Moreover, let's say that at the end of one random day there has been a total of 3 patients. Patient 1 has stayed the whole day and thus, he has occupied one bed. Patient 2 stays from 00:00 to 16:00 and patient 3 arrives at 22:00 and stays until the next day. Patient 2 and 3 can be allocated in the same bed because of their admission and discharge times and, therefore, only one nurse is needed for them. Although there have been 3 patients in total, only 2 nurses are needed at once. In order to staff based on demand, it is not enough to predict the demand based on the patients admissions, but the length of stay plays also an important role as seen in the aforementioned example.

In order to do the study taking into account the patients admissions and the length of stay, a useful way to put this information together is by calculating the number of beds occupied every day. For that, we have developed a code using Rstudio software (version 3.2.3). Giving as input the patient number, admission date and time, and discharge date and time, the program calculates every six minutes (from 00:00h to 24:00h), and for every day, the number of beds occupied. The output is a matrix with 241 rows (corresponding to the hours of the day every six minutes) and the columns are all the days from 1st January 2012 to 31st December 2015. Each element of the matrix indicates the number of beds occupied in a specific day and time. The code can be found in the Appendix A.

Since we perform all the study based on the three different shifts (night, day and evening) we want to obtain the number of beds occupied for each shift per day. Obviously, the beds occupied vary during the day and within each shift. One first way of dealing with this variation could be to calculate the maximum value per day for each shift. By using this approach, only the maximum number of beds occupied would be considered, regardless of the total time this maximum has been present. If we staff based on the maximum, and it is present only during a small period of time, then we could be over staffing. For this reason we have thought about another way to deal with the variation, which is to smooth the data and to consider not only the maximum but also the mode (most frequent value) and the mean. For that, we have done a small change in the hours per shift, already implemented in the ICD, so as to simplify the calculations. We have considered that every shift consists of 8 hours and the division is done according to:

- Night: from 00:00h to 8:00h
- Day: from 8:00h to 16:00h
- Evening: from 16:00h to 00:00h.

	<i>Maximum</i>	<i>Mode</i>	<i>Mean</i>
Date	1/1/12	29/2/12	2/2/12
Shift	Day	Day	Day
Mean	5.79 beds	7.72 beds	7.32 beds
Mode	6 beds	8 beds	7 beds
Maximum	6 beds	9 beds	9 beds
Hours mode	6.32h	4.74h	2.67h
Hours maximum	6.32h	1.28h	1.19h
Beds occupied	6 beds	8 beds	7.32 beds

Table 1.2: Example of how the program chooses the maximum, mode, or mean according to the criteria explained

To calculate the beds occupied smoothing the data, we have used the same procedure for all the shifts. We have assumed that, so as to just consider the maximum value, the latter has to be present at least 2 hours. If this condition is not met, then we analyze the mode. If the mode is present more than 4 hours (half of a shift), then this will be the number of beds occupied. Otherwise, the mean will be the chosen value. In order to make it clearer, Table 1.2 depicts an example of real values and how the number of beds occupied is determined according to the different criteria. In the second column we see that the hours the maximum value is present are 6.32h. Since it is more than 2 hours, the maximum is chosen. In the third column we see that we do not take the maximum into account because it is present less than 2 hours. Then, we study the mode. In this case the mode is present more than 4 hours, therefore we choose it. In the last column, we do not take into consideration neither the maximum nor the mode because they are present less than 2 and 4 hours respectively. Thus, we choose the mean.

Methods for the research sub-questions

During almost all the study we analyze separately the Thorax and the General units of the old ICD. We do not pool their data together because then, the results are not representative of each of them and we can get a wrong idea. However, in the end, we pool the results to have a conclusion of all the ICD.

1. *What has been the situation of the ICD regarding patient admissions and length of stay from 2012 to 2015?*

This question is answered analyzing patients admissions according to the unit, origin, and day of the week admission, and the length of stay using descriptive statistics in Excel.

2. *Which are the bed occupancy rates of the Thorax and General units from 2012 to 2015?*

This question is addressed by calculating the bed occupancy rates for each day using the number of beds occupied and the actual capacity. To do so for one specific date, first we determine the rates for each shift using Excel and then we average them to obtain just one value representative of the whole day. We do this because the capacity differs depending on the month, day of the week and shift. The capacities used are the ones explained in Section 1.1.

3. *What does literature say about demand forecasting in the health care sector?*

We find scientific papers of our interest using the browser of the library of the University

of Twente (FindUT). With it we have access to high quality journals and the possibility of having the full paper. The words used to find them are combinations of: “Demand Forecasting”, “Emergency Department”, “Intensive Care”, “Time Series”, “ARIMA models” and “Health Care”. We also use the references of the selected papers to extend our literature review and to have more insight.

4. *Which other methodologies can be used for staffing?*

Apart from the demand based approach, we define two more approaches: capacity and an hybrid approach (mixing demand and capacity based staffing). Within each approach we define different scenarios, which correspond to the methodologies mentioned in the question. The different methodologies imply different ways of staffing according to one main approach. Inside the demand approach, we consider perfect prediction, naive forecast, moving averages and seasonal indexes. Within the capacity approach, we include staffing based on 100% of the capacity, on 85%, on 80%, on 75% and on 70%. The hybrid approach contemplates the possibility of staffing based on an 85%, 80%, 75% and 70% of capacity and fill the rest with demand.

5. *How accurately can the demand of the Intensive Care Units be predicted?*

We tackle this question using a statistical approach and heuristics. In both cases we use the number of beds occupied. For the statistical approach, according to the type of data that we have, we use time series analysis. Namely, we use ARIMA (Box-Jenkins) models. The model is fitted with a training set (data from 1st week of 2012 to the 26th week of 2015) and it is tested with the remaining data (from the 27th to the 52nd week of 2015). To assess the performance of the model, we use the Root Mean Squared Error (RMSE) criteria. The whole study is addressed using Rstudio version 3.2.3 and the package “forecast”. Regarding the heuristics, we carry out a naive forecast, an adaptation of moving averages and seasonal indexes forecast to predict 2015. The performance of this methods is also assessed using the RMSE.

6. *How large is the pool of nurses needed in the ICD based on demand and capacity?*

We answer this question calculating the pool of nurses needed in 2015 based on the approaches and scenarios presented in sub-question 4 using Excel. For each scenario we first calculate the beds occupied per shift every day and then the number of nurses needed per shift every day as well. We determined it by dividing the number of beds occupied by the corresponding ratios (mentioned in Section 1.1). Once we have the number of nurses needed every day and shift we calculate the total hours worked in 2015 according to the nurses needed. Then we divide the total nurse hours by 1683.6 h which is the annual hours worked by a full-time employee in 2015. At the end we obtain the pool of nurses based on full-time employees depending on each scenario.

7. *Which are the cost differences between staffing based on capacity, demand and the hybrid approach?*

We address this sub-question calculating with Excel the salaries that have to be paid to the employed nurses in 2015, the extra hours worked and the cost of the unnecessary hours worked due to overstaffing in certain moments. The fees used are the nurse average salary. All the calculations are done without taxes in order to be able to make comparisons with the ICD annual budget.

8. *Which are the waste, financial and practical viewpoints of the different methodologies considered for staffing?* We tackle this question in Chapter 8 by assessing the results

obtained from a waste point of view (Chapter 6) and a financial point of view (Chapter 7). Moreover, we also introduce in Chapter 8 a practical point of view. With this, we are able to analyze the advantages and disadvantages of staffing based on each approach from different perspectives and, finally we choose one to make our recommendation.

1.8 Verification and Validation

In this project we have done verification and validation at each step developed to make sure we are doing the things correct and the correct things.

Regarding the verification, in order to verify that we do not have errors in the calculations, the software codes are correct and we are applying the statistics methods properly, we have turned to several experts. These have been:

- Frans van Geer: professor at the Geosciences Department of the University of Utrecht and expert in time series. He helped in the application of ARIMA models.
- Henk Broekhuizen: PhD student at the Health Technology and Services Department of the University of Twente. He assisted in the development of the code using R software explained in the Data subsection.
- Josep Lluís Carrasco: professor and biostatistician at the Public Health Department of the University of Barcelona (Spain). He advised to use time series analysis for the analysis of the given data.
- Karin Groothuis-Oudshoorn: biostatistician and assistant professor at the Health Technology and Services Department of the University of Twente. She guided in the election of the ARIMA models.
- Gjerrit Meinsma: professor at the Department of Applied Mathematics of the University of Twente. He helped on the understanding of regression models.
- Job van der Palen: clinic epidemiologist at the Medisch Spectrum Twente. He helped us at giving a first insight to the given data.

For the validation, we have involved our stakeholders at every step of the project to make sure we work with the correct and appropriate data and our study is aligned with their objective and interests. The main stakeholders involved in the validation are the medical manager Ronald Trof and the head of nurses Marc Eijnsink.

It is important to remark that Peter Schuur, the first supervisor of this thesis, has provided assistance and guidance for both, verification and validation, at every step.

1.9 Deliverables

The deliverables of this research are:

- Advisory report: it consists of a report with all the research explained, the results, the conclusions and the recommendations.

- Occupied beds calculator: this is a tool developed with Rstudio that translates the admission and discharge dates and time of the patients into number of beds occupied. As input, the admission date, admission time, discharge data and discharge time are needed. The output is a matrix in which the rows show the time of the day (from 0 to 24h every 6 minutes, thus 241 rows) and the columns correspond to the dates (as many as desired). Therefore, each cell shows the number of beds occupied on a specific date at a specific time.
- Purified database: this database contains the results of the processing and ordering of the original data. First, the original data is processed using the “Occupied beds calculator”. Second it is further processed doing more calculations and ordering it. It has information regarding the number of beds occupied and the bed occupancy rates for every shift of every day for all the years. It also includes the results of each studied scenario.

Chapter 2

Past situation

In this chapter we analyze the situation of the Intensive Care Department from 2012 to 2015 focusing on the demand side. We study the patient admissions and flow (Section 2.1), the length of stay (Section 2.2) and the bed occupancy rates (Section 2.3). These are important characteristics of the demand that will help us having a clearer idea of the ICD workload.

2.1 Patient admissions

As explained in Chapter 1, the old ICD consisted of 2 units, namely the Thorax and General units. The number of total admissions per unit and year from 2012 to 2015 is represented in Figure 2.1. We can see how the number of admissions in both units has stayed very stable throughout the years. In average, 40% of the patients go to the General unit and 60% are admitted in the Thorax unit. This last unit receives more patients despite having less beds. This can be either because patients in the Thorax unit spend shorter lengths of stay and, therefore, there is more turnover of patients, or the occupancy rates in the General unit are low. This issue will be solved later when analyzing the length of stay and bed occupancy rates.

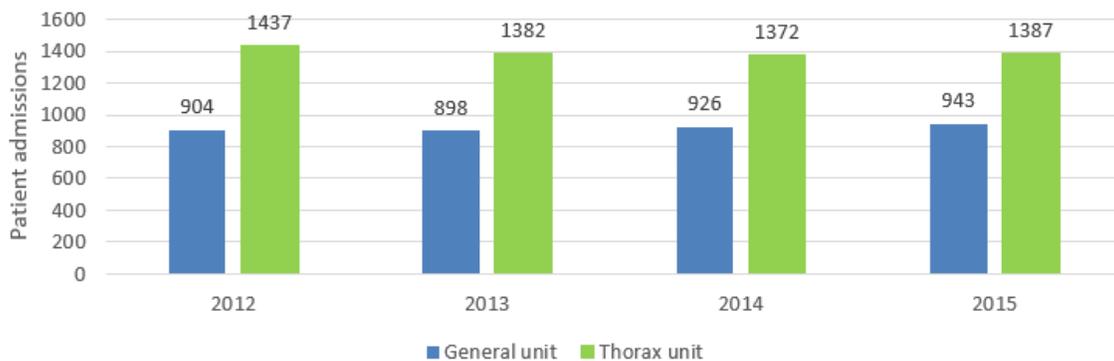


Figure 2.1: Total patient admissions per unit and year

When it comes to the origin of the patients, they can come from the operating room (OR), which is considered an elective admission, or from anywhere else, considered an emergency admission because it is not planned. In Figure 2.2 we can see the admissions percentages for both units depending on the origin. Regarding the Thorax unit (green lines) about 90% of

the admissions are elective, whereas, more or less, 10% are emergencies. Meanwhile, in the General unit (blue lines), 30% of the admissions are elective and 70% are emergencies.



Figure 2.2: Percentage of patient admissions per unit depending on the origin (coming from the operating room or not)

We also want to study the number of admissions throughout the week. Figure 2.3 illustrates the percentage of patients admitted on each unit depending on the day of the week. The percentage is calculated over each year. As we can see, the General unit has a stable number of admissions throughout the week and years. Tuesdays and Fridays are the days with the most admissions, and in weekends less patients come. However the difference is very small. With respect to the Thorax unit, it is clear that there is a huge difference between weekdays and weekends. Mondays and Tuesdays are the days that have the most admissions, while during the weekends a few number of patients come. This makes sense if we recall that 90% of the patients admitted in the Thorax unit had undergone surgery, and in weekends no surgeries are planned.

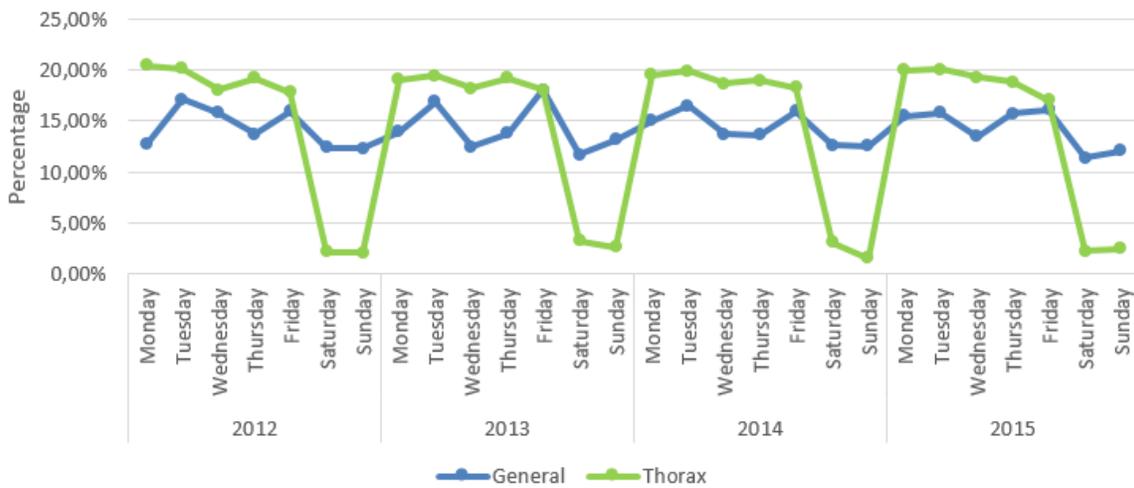


Figure 2.3: Percentage of patient admissions per unit and day of the week

2.2 Length of stay

Apart from the number of admissions, it is also important to study how long patients stay in each unit. The Length Of Stay (LOS) is calculated as the difference in days between the admission and discharge dates plus one ($LOS = (Discharge\ date - Admission\ date) + 1$). This means that if a patient comes and leaves the same day, it counts as an entire day.

The descriptive analysis of the LOS is shown in Table 2.1. The values reflect the average, standard deviation, minimum and maximum length of stay for the 4 years of historical data that we have (from 2012 to 2015). As we can see, the average LOS is shorter for the Thorax unit (2.94 days) than for the General unit (6.62 days). Moreover, patients coming from the operating room stay, in average, less days than patients coming from anywhere else (considered emergencies), regardless of the unit. Furthermore, the maximum LOS is also shorter for the Thorax unit, independently of the origin.

<i>Unit</i>	<i>Origin</i>	<i>Average LOS</i>	<i>Standard deviation</i>	<i>Minimum</i>	<i>Maximum</i>
Thorax unit	<i>OR</i>	2.73	2.54	1	72
	<i>Emergency</i>	4.98	5.03	1	42
	<i>Both</i>	2.94	2.95	1	72
General unit	<i>OR</i>	5.18	8.01	1	95
	<i>Emergency</i>	7.24	11.06	1	163
	<i>Both</i>	6.62	10.29	1	163
Both units	<i>Both</i>	4.40	7.11	1	163

Table 2.1: Length Of Stay analysis

Nevertheless, the average LOS might not be very representative of the real situation that there was in the ICD because infrequent high LOS values divert the mean. Then, it is more appropriate to look the frequency of each LOS. Figure 2.4 and Figure 2.5 illustrate the histogram of the LOS for the Thorax and General units respectively. The left y-axis represents the total number of patients from 2012 to 2015 that stayed, in each unit, the respective LOS. The right y-axis is the cumulative percentage. As we can observe, patients in the Thorax unit (Figure 2.4) were discharged mostly the day after being admitted ($LOS = 2$). 70% of all the patients in the Thorax unit were discharged within 2 days and almost 90% were discharged within the first 4 days. Only 1% of all the patients from 2012 to 2015 stayed more than 15 days in the Thorax unit. With regard to the General unit (Figure 2.5), as well, most of the patients left the unit the day after being admitted ($LOS = 2$) but, unlike the Thorax unit, a greater percentage compared the the Thorax unit were discharged the same day ($LOS = 1$). In this case, 46% of the patients were discharged within 2 days and 89% left the unit within 15 days. 11% of the patients from 2012 to 2015 stayed more than 15 days.

Recalling the fact that there were more admissions in the Thorax unit, less beds and now we know that 70% of all the patients were discharged within 2 days compared to 46% in the General unit, we can fathom that there is a higher turnover of patients in the Thorax unit.

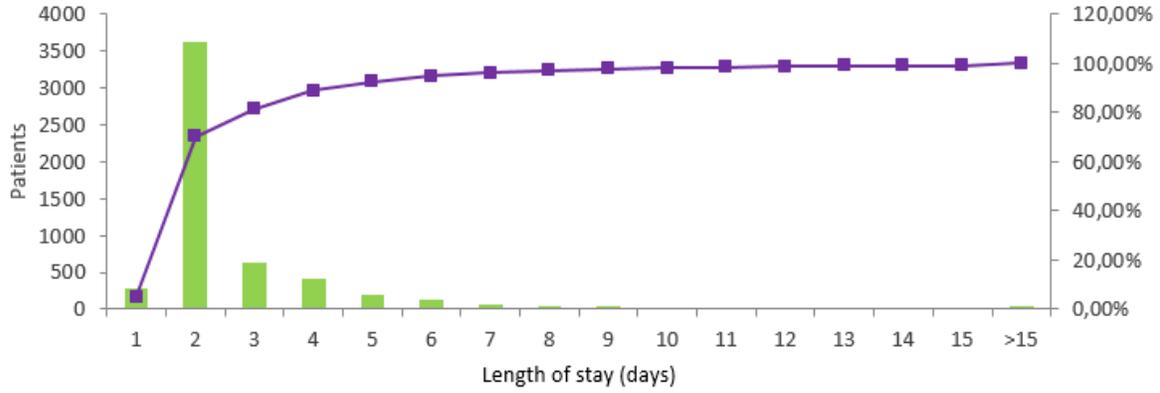


Figure 2.4: Length of Stay histogram for the Thorax unit

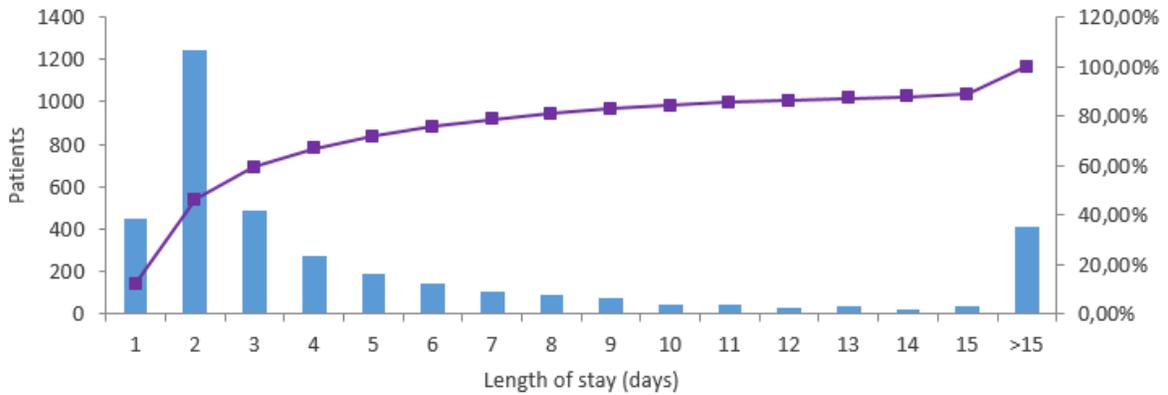


Figure 2.5: Length of stay histogram for the General unit

2.3 Bed occupancy rates

The bed occupancy rates are a sensitive indicator to assess the utilization of the Intensive Care Department and units. The bed occupancy is calculated as the ratio between the number of beds occupied and the beds available (Expression 2.1).

$$Bed\ Occupancy = \frac{beds\ occupied}{beds\ available} * 100 \quad (2.1)$$

First we compute the number of beds occupied for each day and shift as explained in the research method in Chapter 1. Then, we calculate the respective occupancy rates by dividing the beds occupied by the beds available (which correspond to the capacities explained in Chapter 1) and multiplying by 100 (to obtain the percentage) as Expression 2.1 indicates. Since we are considering the three shifts, we obtain three bed occupancy rates per day for four years.

First we want to have an overview of how the daily occupancy rates have varied throughout the years. For that, we average the rates of the three shifts for each day and we plot the

histogram for each year in Figure 2.6 (Thorax unit) and Figure 2.7 (General unit). Both figures depict the number of times (days) per year that each range of occupancy rates has been present.

In Figure 2.6 we see how throughout the years the rates have diminished. In 2012 the most frequent rates were between 91% and 100%. Then, in 2013 it fell to the range between 71% and 80% and rose again in 2014 between 81% and 90%. Finally in 2015 it drop to between 71% and 80%, although rates between 81% and 100% were also usual. Moreover, it seems that 2012 was a different year from the rest as the rates tend to be really high compared to the others. We can observe some rates over 100% which do not make sense because there cannot be more beds occupied than beds available. We have discused this issue with the manager and we do not find an explanation for this. However, this is what data reveals. It could be possible that the discharge time of a patient was registered later than the real time and, by that time, another patient would have occupied the same bed. If that was the case, in the reality there would not have been more patients than beds, but in the system it would seem so. Nevertheless, this is just and hypothesis of what could have happened.

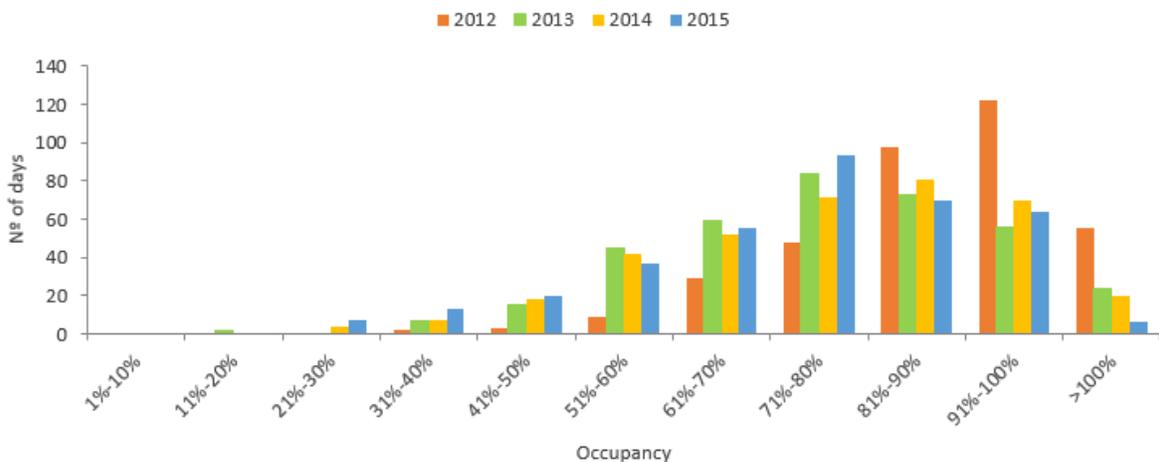


Figure 2.6: Daily average occupancy rates in the Thorax unit for each year from 2012 to 2015

Regarding the General unit, Figure 2.7 illustrates the same behavior as the Thorax unit. In 2012 and 2013 the most frequent occupancy rate was between 81% and 90%. In 2014 and 2015 it decreased to between 71% and 80%, although in both cases rates between 81% and 90% were also common.

We have also analyzed the bed occupancy rates depending on the shift for each unit throughout the four years. The histograms show that the Thorax unit tends to have a slightly higher occupancy in the evenings but the difference with the other shifts is not relevant. The occupancy rates in the General unit are very similar for the three shifts. The histograms can be found in the Appendix B.

The values obtained show that the bed occupancy rates have tended to decrease along the years. Rates between 71% and 90% are the most frequent ones. High occupancy rates would make us think that maybe a demand based staffing approach would not make sense. However,

with the observed rates there is still some room for resource allocation improvement and, then, talking about a demand based approach would make sense.

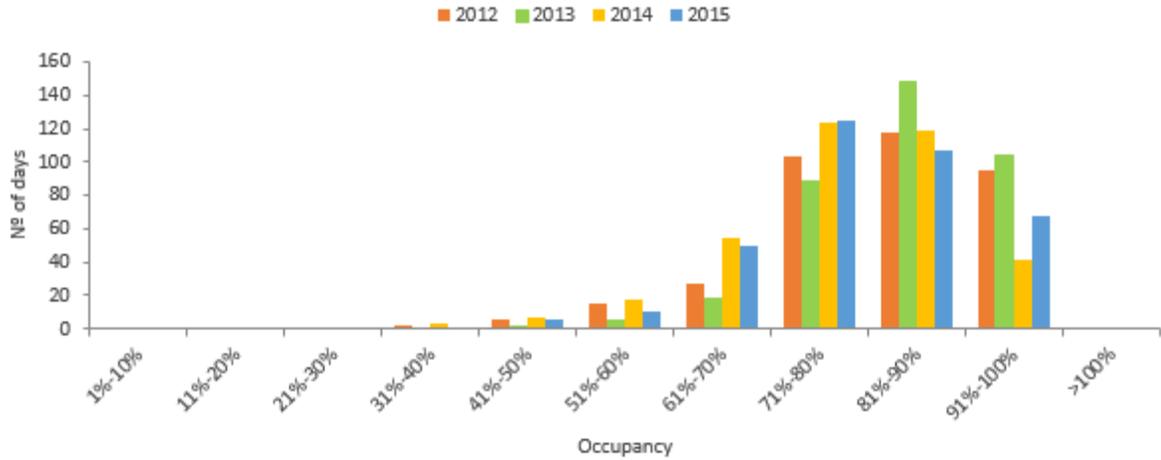


Figure 2.7: Daily average occupancy rates in the General unit for each year from 2012 to 2015

2.4 Conclusions

In this chapter, we have analyzed the situation of the Intensive Care Department in the old MST building from 2012 to 2015. Regarding the patient admissions and length of stay we can conclude that more patients are admitted in the Thorax unit than in the General unit and they stay shorter because they are mainly elective admissions for which the evolution of the patient is, in general, less uncertain than in an emergency admission. With reference to the bed occupancy rates, they have tended to decrease along the years. The most common ones are between 71% and 90%. This range of values makes us think that a demand based staffing approach could make sense.

Chapter 3

Literature review

One of the main focus of the manager is to study what would happen if we would staff based on demand. This involves the use of forecasting methods. In this chapter we review the literature about demand forecasting in the health care sector to see how other researches have handled similar problems in the past, the techniques they have used and the results they have obtained.

3.1 State of the art

As Soyiri and Reidpath (2013) state, forecasting is about predicting future events based on a foreknowledge through a systematic process or intuition. In the forecasting literature one can find two widely used approaches to forecast: judgmental (qualitative) and statistical (quantitative) forecasting. Judgmental methods are used when there is a lack of data, but given abundant data, statistical approaches are preferred (Armstrong et al., 2005; Hyndman and Athanasopoulos, 2014). From the quantitative forecasting approaches, the vast majority use either cross-sectional data or time series data. Cross-sectional data consists of data collected at a single point of time, whereas time series data is a sequence of observations collected at regular intervals over a period of time (Hyndman and Athanasopoulos, 2014). Time series forecasting is, therefore, the use of past values to predict future values. To forecast time series, different methods can be used, and none of them is superior to the others (Jones et al., 2008). The different existing methods for time series prediction are, among others, linear regression, Autoregressive Integrated Moving Averages (ARIMA), exponential smoothing, dynamic regression and artificial neural networks (Hyndman and Athanasopoulos, 2014).

One can find a wide range of literature about forecasting methods applied to the business context to improve efficiency. However, when it comes to health care departments, there are a few examples and mainly focused on the Emergency Department (ED). With regard to the Intensive Care Department, we have just found one Australian study. This study conducted by Corke et al. (2009) predicted the future long-term demand in ICU bed-days for intensive care in Australia and New Zealand with the objective of planning the allocation of resources. They applied ARIMA models to 11 years of data from both countries which included various Intensive Care Departments. The results obtained indicated an increase in intensive care demand of over 50% by 2020 going from 471,358 bed-days in 2007 to a prediction of 643,160 in 2020.

Even though we only found one study about demand prediction in an intensive care unit, we found several examples about the emergency department. According to the systematic review performed by Wargon et al. (2009), the objectives of the different studies focused on the ED were basically the adjustment of staffing patterns to visits and the allocation of resources. For that, different studies sought to predict hourly, daily, monthly and annual patient visits. The authors reported that ARIMA models are one of the most frequently used time series methods to forecast. Regarding the accuracy of the predictions, they identified three main criteria used in the ED to assess the model performance: the percentage of variability in regression analysis (R^2), the mean absolute error (MAPE) and the root mean square error (RMSE).

Abdel-Aal and Mangoud (1998) carried out a time-series analysis using two different methods to model and forecast the monthly patient volume at a walk-in clinic in Arabia Saudi. The study was conducted using 11 years of historical data, 9 to train the models and 2 to evaluate them. To assess the accuracy of the predictions they used MAPE. On the one hand they used ARIMA models and, on the other hand, they tried a much simpler method which consisted in extrapolating the growth curve of the annual means of the patient volume. The best ARIMA model resulted to be a seasonal model with a MAPE of 2.82% for one year prediction. The extrapolation method using a polynomial fit proved to be more accurate with a MAPE of 0.55%, also for one year prediction, but it also performed better in longer-term forecasts. They attributed the better performance of the second approach to the regular nature of the data considered and the lack of a stronger random component.

In a study conducted by Jones et al. (2002) in UK, they forecasted the daily number of beds occupied due to emergency admissions and the number of emergency admissions at Bromley Hospitals NHS Trust. The authors calculated the number of beds occupied as the addition of patients admitted as emergencies and patients who occupied a bed at midnight on the respective date. They used 6 years of data to develop a seasonal ARIMA model fitted using the Box-Jenkins technique and assessed the accuracy of the predictions using the RMSE. The results indicated that they found a good seasonal ARIMA model to forecast the beds occupied with a RMSE of 12.6. The model did not perform that good when there was bed crisis, which is when the model is needed the most. However, the work that such model implied was worth it because assuming a naive forecast would have given worse results. Moreover, they also studied whether the addition of two external factors such as the temperature and the influenza illness rate could improve the predictions, but these factors only increased the forecasting error. Regarding the patient admissions, the predictions obtained were not as accurate as for the beds occupied. They remarked that using the mean of the admissions was almost as good as the prediction model they developed. The hospital managers indicated that knowing the beds occupied was more useful for planning improvement.

In an Australian study, Champion et al. (2007) used time series to forecast the monthly number of patients presentations at the emergency department in a regional Victoria hospital. Their analysis was based on 5 years of data. The statistical forecasting methods used were exponential smoothing and Box-Jenkins methods (ARIMA models). They fitted the models with 5 years of data and tested them for the following 5 months. To assess the performance they used the RMSE. For their particular case, they found that the best exponential smoothing model corresponded to a simple seasonal model with a RMSE of 3.3, whereas the optimal ARIMA model was a non-seasonal moving average model with a RMSE of 3.9. They stated

both models predicted the patient presentation numbers quite successfully although the exponential smoothing method performed a little bit better.

Jones et al. (2008) used several statistical forecasting methods to predict daily ED patient presentations in three different US hospitals. The data collected consisted of daily patient arrivals from January 2005 to March 2007. The authors studied the performance of seasonal ARIMA models, exponential smoothing, time series regression and artificial neural network models and they compared them to the benchmark method used (multiple linear regression with calendar variables). To conduct the study they fitted the models with the first 2 years of data and they tested them with the remaining 3 months. They used the MAPE to evaluate the forecast accuracy. The results obtained revealed that all time series methods studied offered a better in-sample goodness-of-fit when compared to the benchmark methods. Nevertheless, when doing a post-sample analysis, the forecast accuracy of the seasonal ARIMA models, exponential smoothing and artificial neural networks was not good. Other models based on time series regressions, as the benchmark model, provided only small improvements. They concluded that multiple linear regression based on calendar variables is a good approach for forecasting daily patient volumes in the ED.

In an Indonesian study, Earnest et al. (2005) predicted the number of beds occupied during the Severe Acute Respiratory Syndrome (SARS) outbreak in Singapore. The authors applied ARIMA models to make real-time predictions on the number of beds occupied using data of 3 months and calculating the MAPE to assess the forecast accuracy. The model was fitted with one month of data and tested with the remaining data. The best ARIMA model they found was a non-seasonal model with a MAPE for the training set and test set of 5.7% and 8.6% respectively. This results translated to an error rate represent ± 7 beds and ± 13 beds respectively. They considered the results reasonable for use in the hospital setting.

There are also more studies focused on the ED that have applied the forecasting techniques mentioned so far and others. For example, Boyle et al. (2012) used ARIMA, regression and exponential smoothing methods to predict the ED presentations and subsequent admissions. McCarthy et al. (2008) developed a study with the utilization of Poisson regression as a function of temporal, climatic and patient factors to predict hourly ED presentations. Farmer and Emami (1990) predicted the surgical beds required with the use of ARIMA models and regression methods. Hoot et al. (2008) developed a discrete event simulation of patient flow to predict ED operating conditions in the next 8 hours. Finally, Reis and Mandl (2003) fitted ARIMA models to predict daily ED presentations.

Various report that the results obtained in one hospital might not be applicable for other hospitals due to different predominant circumstances. Nevertheless, the methodology used might serve as an example for other hospitals who want to conduct similar studies. Although the above studies claim to have obtained good prediction results for the Emergency Department, none of them has considered a transition from capacity to demand based staffing. Their predictions have only been used as an additional help for planning.

3.2 Conclusions

In this chapter we have discussed the state of the art regarding demand forecasting in health care departments. In this context, one prefers a statistical rather than a judgmental forecast. From the statistical approach, the analysis of time series is the most common one. For that, one of the most widely-used methods is the application of ARIMA models. The different studies report to have obtained good forecasting results but none of them has considered a transition from capacity to demand based staffing. The forecasting is only an additional tool to improve resource allocation and planning.

Chapter 4

Theory and concepts

In this chapter we explain the theory needed to be able to understand the steps followed and the notation used in Chapter 5. Section 4.1 introduces time series, Section 4.2 explains the theory behind ARIMA models, Section 4.3 summarizes the Box-Jenkins technique and, finally, Section 4.4 describes how to assess the forecast accuracy.

4.1 Time series

As described in Chapter 3, a time series is a sequence of observations collected at regular intervals over a period. According to Ozcan (2005) and Hyndman and Athanasopoulos (2014), the analysis of a time series can identify behaviors in terms of:

- **Trend:** a trend exists when there is a gradual long-term upward or downward movement in the data. It does not have to be on just one direction, it can change from an increasing trend to a decreasing trend and vice versa.
- **Seasonality:** seasonality exists when there are short-term variations in the data, which are related to seasonal factors such as the weather, holidays or others. The period of a seasonal pattern is always known. For example, it can be daily, weekly, monthly, quarterly...
- **Cycles:** there are cyclic patterns when rises or falls occur in a non-fixed period, usually of several years. They are often related to the economic conditions of the moment.

A time series comprises three components: a seasonal component, a trend-cycle component and a remainder component. The remainder contains anything else not included in the other two components. If we denote the time series as y_t , then we can write:

$$y_t = S_t + T_t + E_t , \tag{4.1}$$

where S_t is the seasonal component at time t , T_t is the trend component and E_t the remainder or error. Expression 4.1 represents the additive model. A multiplicative model can also be used and it is written as:

$$y_t = S_t * T_t * E_t \tag{4.2}$$

The use of the additive or multiplicative model depends on the magnitude of the seasonal fluctuations or the variation around the trend-cycle. If the seasonal fluctuations or variation

is proportional to the level of the time series, i.e., it increases as time passes, a multiplicative model is more appropriate. If they do not vary, then an additive model is the adequate one.

4.2 ARIMA models

As mentioned in the literature review (Chapter 3), ARIMA models are one of the most widely-used approaches for time series forecasting (Wargon et al., 2009; Champion et al., 2007; Jones et al., 2002; Hyndman and Athanasopoulos, 2014) and we have found several examples of their application in health care departments. This is why we have decided to use ARIMA models for our forecasting.

ARIMA models represent a flexible resource to model and forecast a wide range of time series by describing the autocorrelations in the data (Jones et al., 2002; Hyndman and Athanasopoulos, 2014). An ARIMA model is characterized by three processes that are associated to three different parameters (p , d and q). The first process corresponds to the autoregressive part (AR), in which the variable of interest is treated as a linear combination of p past values of the variable. The second process is the Integration part (I) which is the reverse of differencing. Differencing consists in computing the differences between consecutive observations in order to make the time series stationary. The degree of differencing is represented by d . Finally, the moving average process (MA) assumes that each value depends on q past forecast errors. The full ARIMA model can be written as:

$$y'_t = C + \sum_{i=1}^p \phi_i y'_{t-i} + \sum_{i=1}^q \theta_i \varepsilon_{t-i} + \varepsilon_t, \quad (4.3)$$

where y'_t is the differenced series (it may have been differenced more than once), C is a constant, ϕ accounts for the coefficients of the autoregressive process, θ for the coefficients of the moving averages process and ε is the error (white noise). The transformation from the original to the differenced series is done as follows:

$$y'_t = (1 - L)^d y_t \quad (4.4)$$

where y_t is the original series, d is the degree of differencing, and L is a backshift operator. This means that L operating on y_t has the effect of shifting the data back one period. Two applications of L on y_t shift the data back to periods.

Using the Box-Jenkins notation, we can call this model an $ARIMA(p, d, q)$, where p is the order of autocorrelation, d is the degree of differencing and q is the order of the moving averages part.

Seasonal ARIMA models are used when data are subjected to a strong seasonal pattern. The basis is a non-seasonal ARIMA model in which additional seasonal terms are added. In this case P and Q represent the seasonal AR and MA parts respectively. The full seasonal ARIMA model is described by:

$$y'_t = C + \sum_{i=1}^p \phi_i y'_{t-i} + \sum_{i=1}^P \Phi_i y'_{t-mi} + \sum_{i=1}^q \theta_i \varepsilon_{t-i} + \sum_{i=1}^Q \Theta_i \varepsilon_{t-mi} + \varepsilon_t, \quad (4.5)$$

where y'_t is the differenced time series (using seasonal and first differences), Φ accounts for the coefficients of the seasonal AR part, Θ represents the coefficients of the seasonal MA part and m is the number of periods per season. The transformation from the original seasonal series to the differenced one is done as follows:

$$y'_t = (1 - L)^d(1 - L)^D y_t \quad (4.6)$$

where D represents the seasonal differences.

With the Box-Jenkins notation we write the model as $ARIMA(p, d, q)(P, D, Q)_m$.

4.3 The Box-Jenkins technique

The Box-Jenkins technique is commonly used to fit ARIMA models. It consists of four main steps (Jones et al., 2002; Jones et al., 2008):

1. First of all the time series has to be stationary, which implies that the statistical properties do not vary over time. If it is not, then it has to be differenced.
2. Once the time series is stationary, the Autocorrelation (ACF) and Partial Autocorrelation Functions (PACF) are plotted. The ACF is a correlogram and it shows the correlation of the variable for different lags, i.e, for various time differences. For example, if lags are represented by weeks and the correlation of lag 1 is 0.6 it means that all the data with 1 week difference is 60% correlated. The PACF is also a correlogram but it shows the relationship between one observation and the other one after removing the effects of other time lags. Both correlograms are used to find the appropriate order of the process (determine p , d and q).
3. Once a model is found, one needs to check the residuals, which are the remaining component after deleting the trend and seasonality of the time series. The residuals have to be like white noise, i.e., random, for the model to be adequate.
4. If the residuals do not meet the requirement of appearing like white noise, then alternative models have to be considered.

4.4 Forecast accuracy

Regarding the accuracy of the forecastings, we have seen in Chapter 3 that the researchers have used mainly two criteria: the Mean Absolute Percentage Error (MAPE) and the Root Mean Square Error (RMSE). The use of the diverse criteria and their disadvantages has been discussed several times and everyone recommends a different one. We have chosen to work with the RMSE because of the disadvantages reported by Hyndman and Athanasopoulos (2014) about the MAPE. The RMSE is a scale-dependent error which means that the error is on the same scale as the data. This entails the impossibility of comparing series with different scales. The RMSE is calculated by:

$$RMSE = \sqrt{\frac{1}{n} \sum_{i=1}^n (y_i - \hat{y}_i)^2}, \quad (4.7)$$

where y_i is the i^{th} observation, \hat{y}_i is the corresponding forecast and n is the total number of observations.

4.5 Conclusions

In this Chapter we have reviewed the theory needed to understand time series forecasting and ARIMA models. A time series is a sequence of observations collected at regular intervals that can capture different behaviors: a trend, seasonality and cycles. The time series comprise three different components: a seasonal pattern, a trend-cycle component and the remainder. To describe the series either the additive and the multiplicative models can be used but it depends on the type of data we have. ARIMA models are widely used to forecast. They consist of three process: autoregressive process (AR), integration process (I) and moving averages process (MA). There are non-seasonal ARIMA models, which are described by $ARIMA(p, d, q)$ and seasonal ARIMA models, which consist in the addition of seasonal terms to the non-seasonal ARIMA. They can be written as $ARIMA(p, d, q)(P, D, Q)_m$. The Box-Jenkins technique describes four steps to fit an ARIMA model. Regarding the forecast accuracy, we chose to work with RMSE which is a scale-dependent error.

Chapter 5

Feasibility of demand prediction

In this chapter we try various ways to predict demand and we assess their performance. In Section 5.1 we do a statistical prediction using ARIMA models, whereas in Section 5.2 we try easier ways to predict using heuristics.

5.1 Statistical prediction

For the statistical prediction we apply ARIMA models. According to Hyndman and Athanasopoulos (2014), the steps based on the Box-Jenkins technique for the forecasting process of a time series using ARIMA models are the following:

1. Plot the data and identify any unusual observations
2. Transform the data to stabilize the variance if necessary
3. Make data stationary
4. Examine the ACF and PACF plots and decide which $AR(p)$ or $MA(q)$ model is the most appropriate one
5. Try the chosen model(s) and use the AIC_c and RMSE to search for the best model
6. Check the residuals from the chosen model. If they are not adequate, try a modified model
7. Calculate the forecasts

There are automated algorithms that take care of steps 3 to 5, but the researcher has to carry out the other steps. We will follow these steps to find the best ARIMA model to forecast the demand.

The ICD manager wants the forecast of the demand for each shift (day, evening and night) and day. For this, we are going to predict the number of beds occupied. The number of beds occupied is similar across the shifts and years (see figure in Appendix B). Hence, the analysis to find the forecasting model will be done using just one shift (night). If the model is fitted correctly and it is possible to accurately predict the demand, then the other shifts will be analyzed and fitted to an ARIMA model as well.

We fit the models with data from 2012 to the second quarter of 2015 (training set) and we test them with the third and fourth quarters of 2015 (test set). Despite the fact that 2015 has 53 weeks, we assume it has 52 weeks to simplify the calculations. We will carry out steps 1 to 4 with the training set and step 5 with the test set.

5.1.1 Thorax unit

Step 1: Plot the data

First of all, we plot in Figure 5.1 the data to see how it looks like and which is the weekly behavior throughout the years. In Figure 5.1 we can see the daily average number of beds occupied per week from the 1st week of 2012 to the 26th week of 2015. Generally speaking, we can say that summer weeks have less beds occupied than winter weeks, except for the last week of each year, where in 2013 and 2014 we observe a dramatic fall.

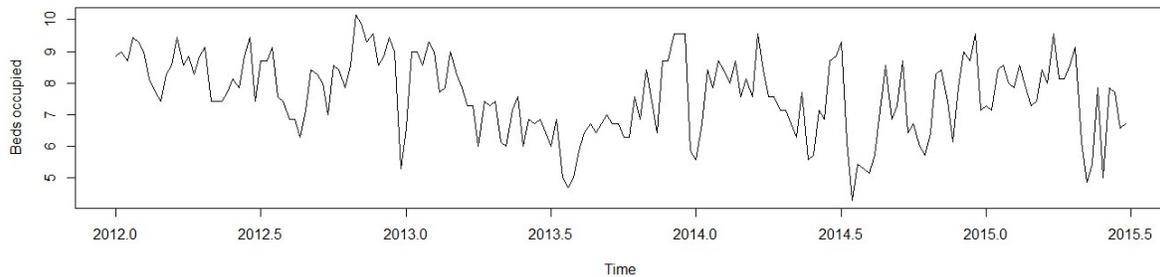


Figure 5.1: Daily average of beds occupied per week in the Thorax unit

As explained in Chapter 4, a time series consists of a trend, a seasonality and a remainder. If we add these three components, we obtain the original time series. We use an additive model because the variations in the data are not proportional to the level of the time series. We can decompose our original data to have an idea of how these components look like and to know more about our time series. To do that, we use the `stl()` function (in Rstudio) which does a seasonal decomposition of time series. The main parameter to be set in this function is the seasonal window (`s.window`), which we have set it to be periodic (i.e., identical across years). Hence, we will see that the seasonal component will be the same for all the years.

Figure 5.2 shows the decomposition of the time series using the `stl()` function for the Thorax unit. The first row shows the original data (same as Figure 5.1), the second one shows the seasonal component, the third one the trend and in the last place we see the remainder. The grey bars to the right of each row show the relative scales of the components. Each bar represents the same length but due to the different scales of the different plots, the bars vary in size. Focusing on the second panel, we see the seasonal component. According to prof. dr. ir. F.C. van Geer, an expert in time series from the Utrecht University, we would expect a smooth seasonal pattern but the plot shows a spiky one, which looks more like a coincidence due to the short period of data that we have. Hence, according to him, the seasonal component is not reliable. Regarding the trend component, its contribution is really small as we can see in the size of the corresponding bar. Although it is a description of the data, we cannot use this type of trend in a forecast according to prof. van Geer. The trend is also not reliable because of the lack of data to properly determine it. He suggests that we need longer time series (for example 10 years) for a reliable estimate. On the other hand, the remainder is the

random part of the time series that remains after removing the trend and seasonality. In this case it has an important contribution to the time series as the bar size is quite similar to the one of the original data.

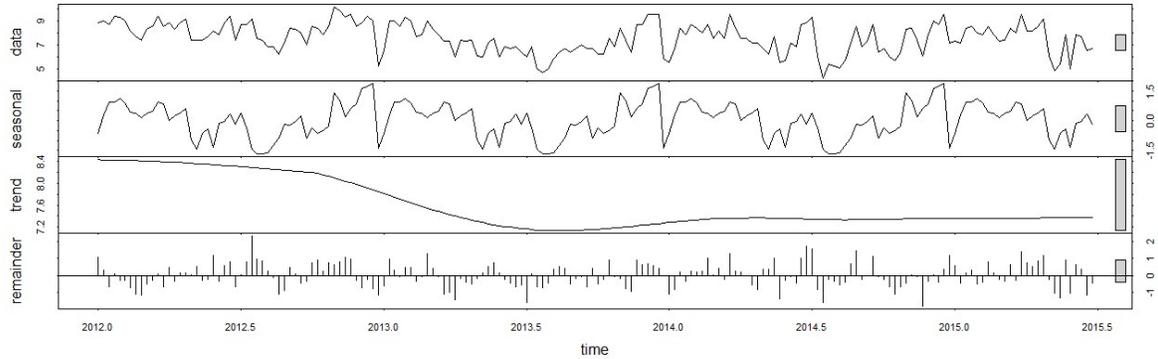


Figure 5.2: Time series decomposition for the Thorax unit

Step 2: Data transformation

The next step consists in transforming, if necessary, the data using Box-Cox transformations to stabilize the variance. This is usually necessary in those cases which the variation appears to be proportional to the level of the time series (it increases over time). Nevertheless, in this case the variation does not vary with the level of the time series, so no transformations are needed.

Step 3: Is data stationary?

Most time series forecasting methods are based on the assumption that time series can be rendered approximately stationary, which means that the statistical properties such as the mean, variance and autocorrelation do not vary over time. Therefore, the third step consists in doing our time series stationary. If data is not stationary, then it needs differencing, which consists in computing the differences between consecutive observations (known as first differences) or between an observation and the corresponding observation from the previous year (seasonal differences). This helps stabilizing the mean of a time series.

The stationary state of the data can be checked using different methods. A visual way to do it is by using the ACF plot, which is called correlogram and it shows the correlations for different lags, i.e., for different time differences: lag 1 means one week of difference, lag 2 means two weeks of difference and so on. For stationary time series, the correlations of the ACF plot drop to zero relatively quickly, while for non-stationary data they decrease slowly. Moreover, for non-stationary data, the value of the first lag is usually large and positive. On the left plot (ACF) of Figure 5.3 we can see that the correlations of our data do not drop immediately to zero but they also do not decrease really slow. The value of the first lag is positive and more or less large (0.6, being 1 the maximum). To be sure if we need differencing

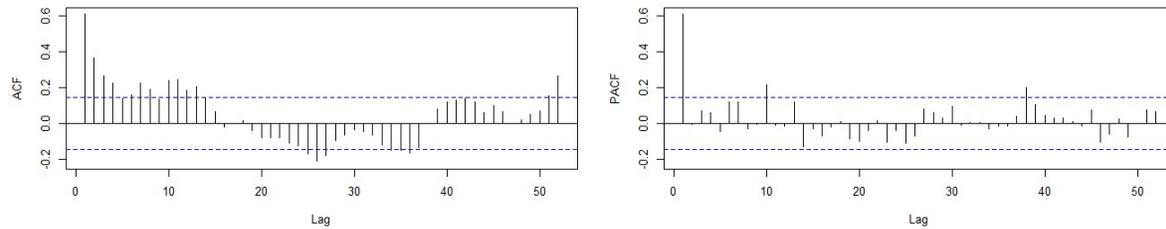


Figure 5.3: Correlation plot (ACF) (left) and partial correlation plot (PACF) (right) of the weekly data in the Thorax unit

we will use statistical tests.

A more objective way to determine if differencing is needed is using the union root tests. The most famous tests are the ADF and KPSS tests. Rstudio has two commands (*ndiff()* and *nsdiff()*) that use these and other tests to determine statistically if differencing is necessary. In this particular case, the program says that our data needs one order of first differences, which means to compute just one time the difference between one observation and the next for the whole time series. This way, our data will become completely stationary.

Step 4: Examine the ACF and PACF plots

By examining the ACF and PACF plots of the differenced data, sometimes it is possible to determine visually which model is suitable for the data. However, before thinking about the ARIMA models, we can study the ACF and PACF plots of the non-differentiated data and explain a few things about the predictions.

By observing the ACF plot (left) of Figure 5.3, we can see the correlations for different lags. The plot is done under the null hypothesis that the data is uncorrelated. The dashed lines are the 95% confidence interval, therefore correlations above or below these lines suggest that the data is not uncorrelated. Hence, from the figure we can say that lags 1 to 14 are correlated and some lags further as well. However the correlations are weak except from lag 1, which is 0.6. One might think that maybe it would be possible to predict 14 weeks ahead because until lag 14 data is correlated. Nevertheless, apart from the correlation being low, this thought is wrong. What actually happens is that there might be some lags with significant correlation because previous lags had also significant correlations. The partial correlations plot (PACF) allows us to overcome this problem by showing the relationship between one observation and another after removing the effects of other time lags. So now, we can look the PACF plot on the right of Figure 5.3 and we can see that only lag 1 has a significant and quite high correlation. From this observation it seems that to obtain the best possible prediction, the forecast has to be done for next week. One might also think that although making a prediction for next week, the forecasting can be done in a long term. This is true, because you can forecast based on another forecast, but the forecasting error will increase a lot every time and the accuracy of the prediction will decrease considerably.

Going back to determining the parameters for the ARIMA models, we examine the ACF and PACF plots of the differenced data shown in Figure 5.4. Usually the parameters can be quite easily estimated when one of the plots is exponentially decaying or sinusoidal. However, in this

case neither the ACF nor the PACF plots have such shape, but we do know that parameter d is 1 as we calculated in step 3. Therefore, it is not possible to define p and q parameters visually. To overcome this problem we use the automated algorithm (*auto.arima()*) which intends to calculate the best model, although it is not always like this. The algorithm takes some short-cuts in order to speed up the computation, but to obtain a better calculation we turn the short-cuts off.

Since we are not convinced about the seasonality of the data, we first consider non-seasonal ARIMA models ($ARIMA(p,d,q)$). The automated algorithm suggests an $ARIMA(2,1,2)$. Despite not being very confident about the seasonality, we are curious about what would happen if we would contemplate a seasonal ARIMA model ($ARIMA(p,d,q)(P,D,Q)_m$). We set the algorithm to take into account seasonality and it advocates an $ARIMA(1,1,1)(1,0,0)_{52}$. These two models are our non-seasonal and seasonal current models.

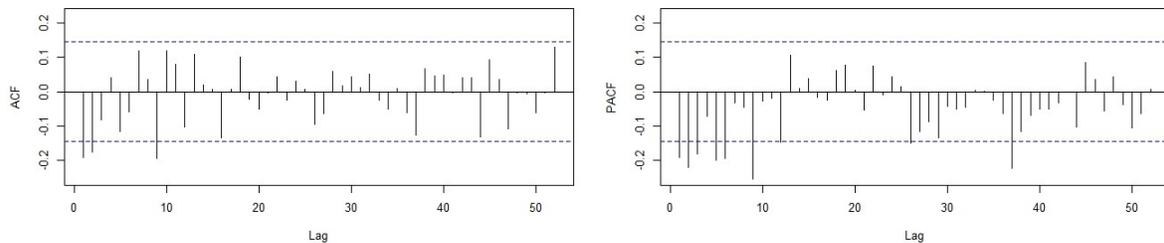


Figure 5.4: Correlation plot (ACF) (left) and partial correlation plot (PACF) (right) of the differenced weekly data in the Thorax unit

Step 5: Try the model

In this step we try both current models as well as other models by varying p , q , P and Q of the current models by ± 1 . To try means to run the model and compute the Akaike's Information Criterion (AIC). AIC is useful to determine the order of an ARIMA model by looking its value and choosing the model with the smallest AIC number. There are also the BIC number and the corrected AIC (AIC_c). Good models are obtained by minimizing either of them. We will chose minimizing AIC_c as Hyndman and Athanasopoulos (2014) do. However, the AIC_c approach is only valid with models that have the same order of differencing. If the orders differ, the AIC_c values are not comparable. Then, we have to use RMSE (Root Mean Squared Error) test set comparisons so that it does not matter how the forecasts were produced, the comparisons are always valid. RMSE shows the root mean squared error between the predicted points and the actual observations, hence we want to select the model with the lowest RMSE. Although in our case we do not modify the d and D parameters, we also compute the RMSE. To do so we select different models by changing the parameters, as explained above, and we do two sets: the training and test sets. We use data from 2012 to the second quarter of 2015 to fit the model and data from the last two quarters of 2015 to test it.

Table 5.1 and Table 5.2 show the tested models and the corresponding AIC_c and RMSE values for the non-seasonal and seasonal models respectively. Although none of them have different orders of differencing within each category, we will use both the AIC_c and the RMSE to select a model. Regarding the non-seasonal models in Table 5.1, we must say that $ARIMA(2,1,2)$, which is the selected model by the automated algorithm, gives some values that

are not numbers (NaNs). Since statistically this model gives some problems, we discard it. Checking the AIC_c , the best model is actually ARIMA (3,1,3) with an AIC_c of 511.6 and RMSE of 1.0944. Nevertheless, together with Dr. Karin G.M. Groothuis-Oudshoorn from the University of Twente we agreed this model is too complicated because a p and q order of 3 is high, which makes it difficult to interpret and, therefore, we should go for simpler models. (Abdel-Aal and Mangoud, 1998) and (Earnest et al., 2005) state that a good model is parsimonious, which means that uses the smallest number of coefficients and thus, the lowest order to explain the data. The next best models are ARIMA (2,1,1) and ARIMA (1,1,2) with an AIC_c of 517.91 and 517.90 and a RMSE of 1.0441 and 1.0667 respectively. The performance of these models is very similar so it does not make a huge difference if we choose one or the other. Again with Dr. Groothuis-Oudshoorn, we chose ARIMA (1,1,2). This is because if we assume that the PACF plot in Figure 5.4 has more or less kind of a sinusoidal shape and the ACF has 2 significant spikes (and we ignore the significant spike 9 because there is a probability that 1 out of 20 spikes is significant), then we roughly have a model ARIMA(0, d ,2). This means that checking the plot we interpret a higher order of moving averages (q) than the autorregressive part (p). Therefore we choose ARIMA (1,1,2) instead of ARIMA (2,1,1).

Non-seasonal			
<i>Model</i>	<i>AIC_c</i>	<i>RMSE</i>	<i>p-value residuals</i>
ARIMA (2,1,2)	515.28	1.1537	0.0705
ARIMA (2,1,1)	517.91	1.0441	0.1145
ARIMA (1,1,2)	517.90	1.0667	0.1169
ARIMA (1,1,1)	516.22	1.1121	0.1732
ARIMA (2,1,0)	537.30	0.9066	0.0074
ARIMA (0,1,2)	519.59	0.9431	0.0074
ARIMA (0,1,1)	537.50	0.9030	0.0007
ARIMA (1,1,0)	544.32	0.8961	0.0005
ARIMA (0,1,0)	549.08	0.8962	0.0006
ARIMA (3,1,2)	521.84	1.0180	0.0368
ARIMA (2,1,3)	521.57	1.0274	0.0435
ARIMA (3,1,3)	511.60	1.0944	0.1617
ARIMA (3,1,1)	519.91	1.1162	0.0800
ARIMA (1,1,3)	519.78	0.9912	0.0534

Table 5.1: Non-seasonal ARIMA models for the Thorax unit. The smaller the AIC_c and the $RMSE$, the better. A high p-value is desired because it indicates that the residuals behave like white noise.

With respect to the seasonal models, apart from the ones shown in Table 5.2, we also studied more models but they gave errors, so we do not take them into account. These are: ARIMA (1,1,1)(2,0,1)₅₂, ARIMA (1,1,1)(1,0,2)₅₂, ARIMA (1,1,1)(2,0,2)₅₂, ARIMA (2,1,1)(1,0,0)₅₂ and ARIMA (1,1,2)(1,0,0)₅₂. Going back to Table 5.2, we can see that the three smallest AIC_c (506.63, 508.37 and 508.37) are from models that have at least one parameter of order 2. Since the RMSE values are quite similar among the different models, we decide to choose ARIMA (1,1,1)(1,0,0)₅₂ because it has the fourth smallest AIC_c (508.56, really close to the others) and it is the simplest model.

Seasonal			
<i>Model</i>	<i>AIC_c</i>	<i>RMSE</i>	<i>p-value residuals</i>
ARIMA (1,1,1)(1,0,0) ₅₂	508.56	1.0868	0.2502
ARIMA (1,1,1)(1,0,1) ₅₂	512.92	0.9465	0.0582
ARIMA (1,1,1)(0,0,1) ₅₂	510.92	1.0819	0.2323
ARIMA (1,1,1)(2,0,0) ₅₂	508.37	1.0609	0.2307
ARIMA (1,1,1)(0,0,2) ₅₂	508.33	1.1278	0.1701
ARIMA (2,1,2)(1,0,0) ₅₂	506.63	1.4987	0.4112
ARIMA (2,1,2)(0,0,1) ₅₂	515.04	1.0877	0.1580
ARIMA (2,1,1)(0,0,1) ₅₂	512.90	1.0557	0.1835
ARIMA (1,1,2)(0,0,1) ₅₂	512.91	1.0604	0.1870
ARIMA (2,1,2)(1,0,1) ₅₂	510.12	0.9228	0.0866

Table 5.2: Seasonal ARIMA models for the Thorax unit. The smaller the AIC_c and the $RMSE$, the better. A high p-value is desired because it indicates that the residuals behave like white noise.

For the Thorax unit analysis we choose the following models to continue with the further steps: ARIMA (1,1,2) and ARIMA (1,1,1)(1,0,0)₅₂. If we calculate the coefficients with data from 2012 to the second quarter of 2015, the models are mathematically written as Expression 5.1 for ARIMA(1,1,2) and as Expression 5.2 for ARIMA (1,1,1)(1,0,0)₅₂.

$$y'_t = 0.4257y'_{t-1} + 0.8044\varepsilon_{t-1} + 0.1104\varepsilon_{t-2} + \varepsilon_t \quad (5.1)$$

$$y'_t = 0.4714y'_{t-1} + 0.2886y'_{t-52} + 0.9352\varepsilon_{t-1} + \varepsilon_t \quad (5.2)$$

where ε denotes white noise. The standard deviation of the white noise in Equation 5.1 is $sd = 0.9945$, whereas in Equation 5.2 it is $sd = 0.9568$.

Step 6: Check the residuals

After selecting the models, the next step is to study their residuals. The residuals are the remaining component after deleting the trend and seasonality. If they appear to be random (behave like white noise), then the model is appropriate. If not, we must try another model. The randomness of the residuals can be checked visually with the ACF plot and more objectively with a Portmanteau test.

The Portmanteau test that we perform is a Ljung-Box test. The settings for the R function are important since we can obtain really different p-values. One parameter that affects completely the final p-value is the number of lags (h) for which the test is performed. According to Rob J Hyndman, in his personal web page where he keeps updating findings about forecasting, the latest approach to determine the optimal number of lags is the following:

- For non-seasonal data, $h = \min(10, T/5)$
- For seasonal data, $h = \min(2m, T/5)$

where h is the number of lags, T is the length of the time series and m is the period of seasonality. In our case, $T = 52 * 3 + 26 = 182$ and $m = 52$. Another important parameter to set is the degrees of freedom (DOF). They should be set according to:

- For non-seasonal data, $DOF = p + q$
- For seasonal data, $DOF = p + q + P + Q$

If the test returns a large p-value, then it suggests white noise. It is also possible that the models do not pass the residuals test but they can still be useful.

On the one hand, the Portmanteau test settings for ARIMA (1,1,2) are $h = 10$ and $DOF = 3$. The resulting p-value of the residuals is 0.1169 as shown in Table 5.1. The p-value is higher than 0.05, which suggests that the residuals are behaving like white noise and that the model can be used for the forecasting. On the other hand, the settings for the seasonal ARIMA (1,1,1)(1,0,0)₅₂ are $h = 36.4$ and $DOF = 3$. The p-value returned shown in Table 5.2 is 0.2502, which also indicates a random behavior of the residuals and therefore, the seasonal model can also be used for the prediction.

Step 7: Forecast

Now that we have selected two models, we calculate forecasts with both of them and we will see which is the most appropriate one.

Figure 5.5 shows the forecast from ARIMA (1,1,2) model. As we can see, it predicts a smooth line with confidence intervals. The dark grey shadow shows the 80% confidence interval, while the light grey one is the 95% confidence interval. The shape of the prediction coincides with the reasoning done by Hyndman and Athanasopoulos (2014) for this kind of model. As they state, models that have a constant equal to zero ($c = 0$) and the differencing order is one ($d = 1$), the long-term forecasts go to a non-zero constant which in this case is 7.26. Therefore, since the forecasts go to a constant and the confidence intervals are very wide causing uncertainty, this information is not useful enough for what we want. We expected a forecast with fluctuations of a magnitude similar to the ones of the known data and confidence intervals adjusted to the fluctuations, not covering the whole y-axis range. We can see that the prediction is the same for all the year except for the first predicted point, which is different from the last observation and from the next predictions. This is because, as mentioned before, the best prediction is in one week ahead and, as shown, long term forecast are not reliable.

We can see in Figure 5.6 the forecast from ARIMA (1,1,1)(1,0,0)₅₂ model. The main difference with the other prediction (Figure 5.5) is that in this one we see a variation in the prediction. This is due to the seasonal parameters from the model. We might think that this prediction is better than the other one at a first sight because it allows us to see fluctuations in the demand. However, as before, the confidence intervals are very broad resulting in uncertainty. The predicted points vary from 6 to 8 beds occupied approximately, which provides us almost the same information.

Figure 5.7 shows the long term forecast of both models compared to the actual observations (black line). The solid red line represents ARIMA (1,1,2) whereas the solid green line is

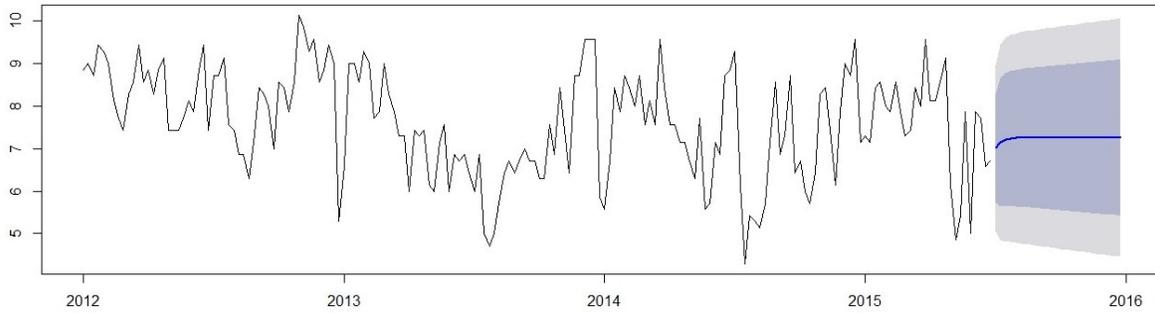
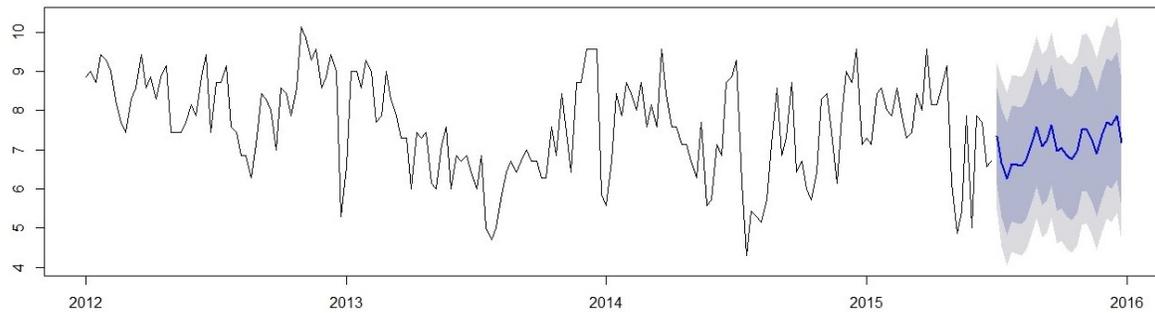


Figure 5.5: Forecast from ARIMA (1,1,2) model for the Thorax unit

Figure 5.6: Forecast from ARIMA (1,1,1)(1,0,0)₅₂ model for the Thorax unit

ARIMA (1,1,1)(1,0,0)₅₂. The dashed lines correspond to the 80% confidence intervals with the respective colors. As we can see, none of the models reflect precisely the actual variations and the confidence intervals are wide.

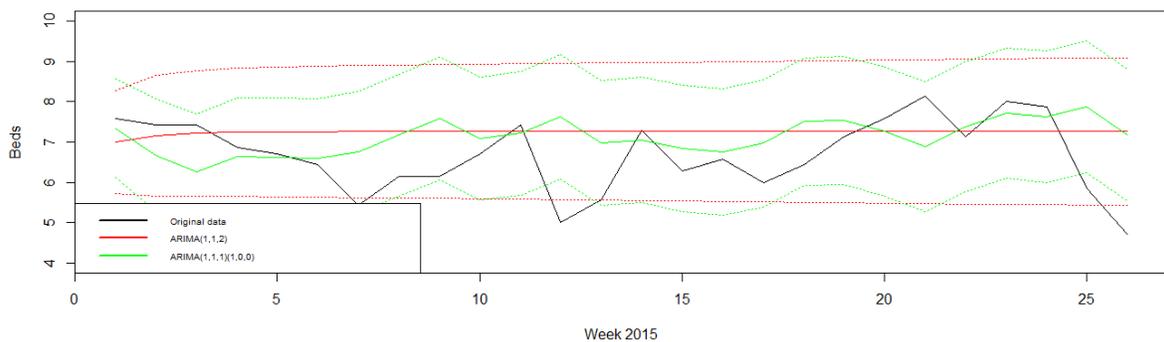


Figure 5.7: Forecast for the third and fourth quarters of 2015 for the Thorax unit. Black line: original data. Red line: ARIMA (1,1,2). Green line: ARIMA(1,1,1)(1,0,0)₅₂. Dashed lines: 80% confidence intervals.

5.1.2 General unit

For the general unit we follow the same steps as for the Thorax unit. In this subsection we just explain Step 1 (plot the data) and Step 7 (forecast). Steps 2, 3, 4, 5 and 6 can be found in the Appendix C.

Step 1: Plot the data

First of all we plot the data in Figure 5.8 to see how it looks like and which is the weekly behavior throughout the years. In Figure 5.8 we can see the daily average number of beds occupied per week from the 1st week of 2012 to the 26th week of 2015. We can see that summers of 2012 and 2014 show a dramatic drop, but for the rest we cannot say that summer weeks have less beds occupied than winter weeks as in the Thorax unit.

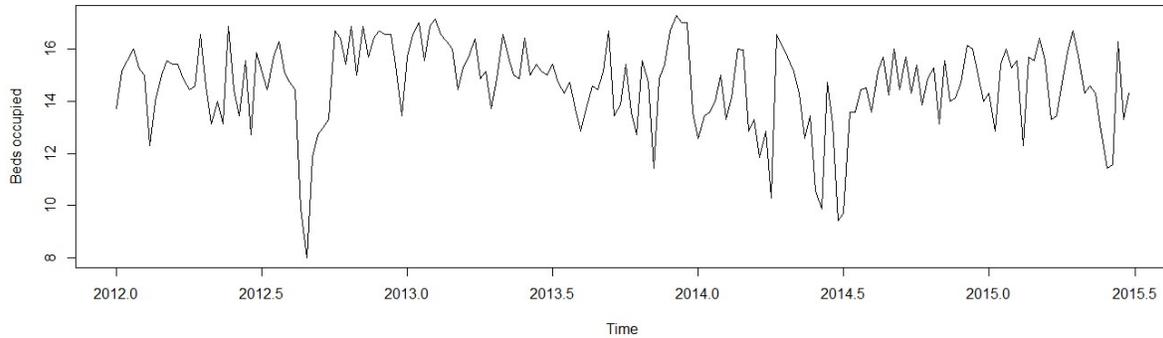


Figure 5.8: Daily average per week of the number of beds occupied in the General ICU

We can decompose our data into a seasonal, trend and remainder components, which is done in Figure 5.9. We can see that the most contributing component is the remainder. The seasonal component also contributes to the time series but, as before, according to prof. dr. ir. F.C. van Geer this seasonality is that spiky that it cannot be trusted. Finally, the trend component is not very important in this time series and, as well, more years of data are needed to determine accurately the trend and seasonality.

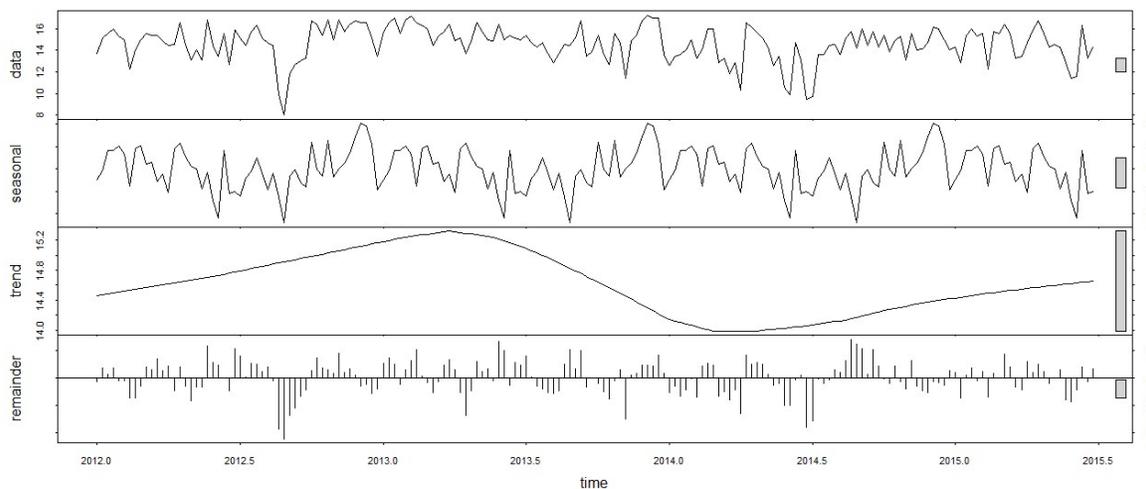


Figure 5.9: Time series decomposition for the General unit

Steps 2, 3, 4, 5, 6

The intermediate steps can be found in the Appendix C. The models chosen for the General unit are ARIMA(1,0,0)^c and ARIMA(1,0,0)(1,0,1)₅₂^c. Both of the models contain a constant and the respective formulations can be found in Expression 5.3 for the non-seasonal model and in Expression 5.4 for the seasonal model. The non-seasonal model has an AIC_c of 652.80 and a RMSE of 1.3999. The seasonal model has an AIC_c of 653.74 and a RMSE of 1.3294.

$$y'_t = 14.6101 + 0.4816y'_{t-1} + \varepsilon_t \quad (5.3)$$

$$y'_t = 14.6115 + 0.4828y'_{t-1} + 0.912y'_{t-52} + \varepsilon_t \quad (5.4)$$

Step 7: Forecast

Figure 5.10 shows the forecast for the General unit from the ARIMA (1,0,0)^c model. It predicts a smooth and constant line of 14.6 beds occupied, together with 80% and 95% confidence intervals that vary from 12.27 to 16.52 and from 11.14 to 17.64 respectively. The shape of the forecasting makes sense because according to Hyndman and Athanasopoulos (2014) a model with non-zero mean ($c \neq 0$) and without first differences ($d = 0$), the long-term forecasts go to the mean of the data, which in this case is 14.6. Thus, long-term forecasts are not useful and the only usable point would be the first prediction. Furthermore, the confidence intervals are very broad causing uncertainty.

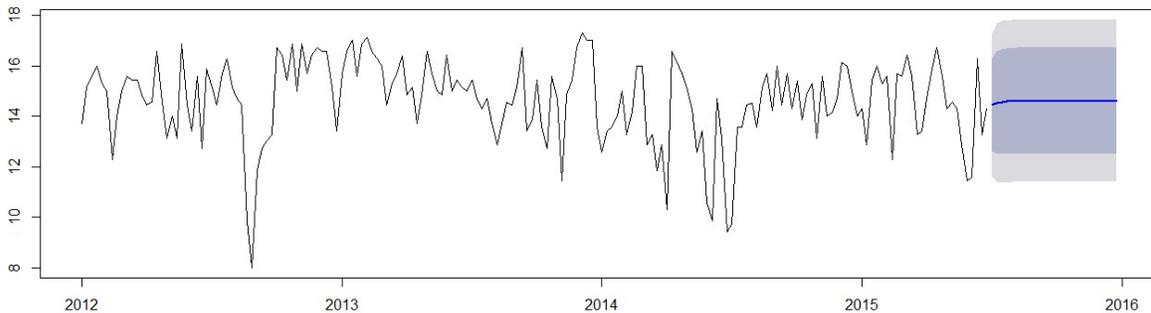


Figure 5.10: Forecast from ARIMA (1,0,0) model for the General unit

Figure 5.11 shows the forecast from ARIMA (1,0,0)(1,0,0)₅₂^c model. Although the predicted line has some variations, it is very stable and, hence, it is like having a straight line as before. Moreover the confidence intervals are very wide, so there is a lot of uncertainty as well. The prediction taking into account the seasonal terms provides the same information, in this case, as the non-seasonal model.

Figure 5.12 illustrates the forecast of the models compared to the actual observations. As we can see, both of them do almost the same prediction and they are not able to predict the fluctuations of the original data. Apart, the 80% confidence intervals are wide and include almost all the original data inside. All this means that long-term predictions are useless in this case. We explained before that the best possible prediction was for next week. Nevertheless, the first predicted point is not really close to the observation.

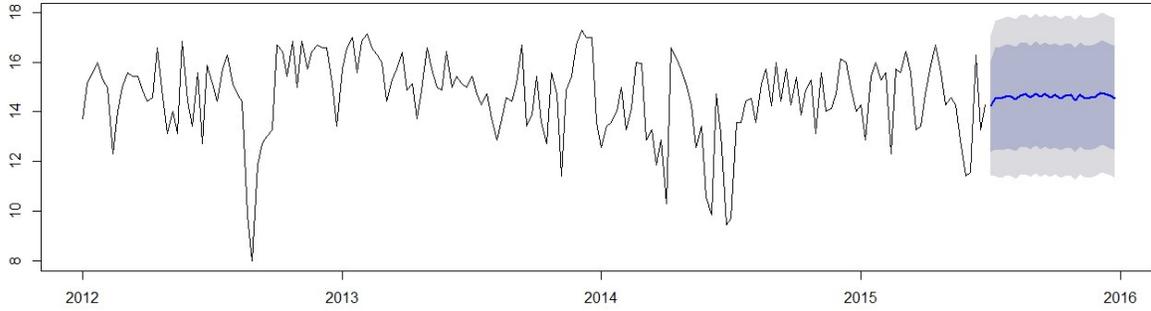


Figure 5.11: Forecast from ARIMA $(1,0,0)(1,0,0)_{52}^c$ model for the General unit

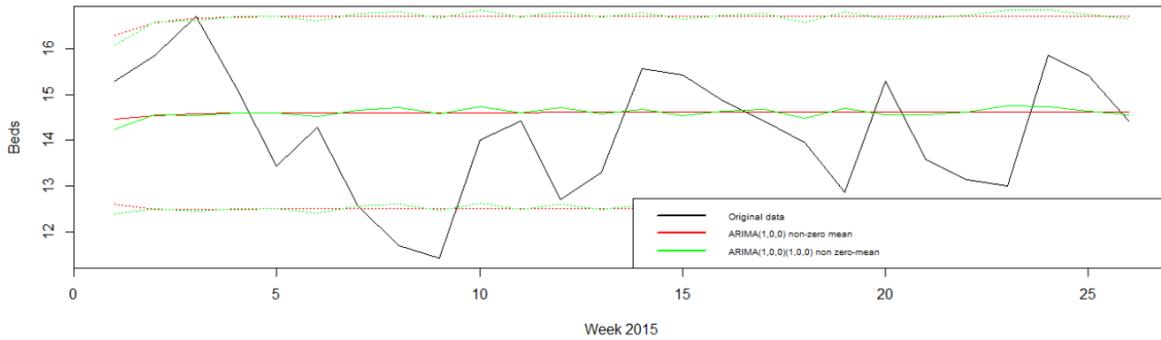


Figure 5.12: Forecast for the second quarter of 2015 for the General unit. Black line: original data. Red line: ARIMA $(1,0,0)^c$. Green line: ARIMA $(1,0,0)(1,0,0)_{52}^c$. Dashed lines: 85% confidence intervals.

Although the accuracy of our forecasts is mathematically good because the RMSE indicates that the forecasts are mistaken in an average of 1.07 beds in the Thorax unit and in an average of 1.41 beds in the General unit, the confidence intervals are still wide and this is what makes our forecast not that useful.

In our case, the use of ARIMA models results in a best possible prediction of one week ahead. This implies that ARIMA models could not be implemented in the Intensive Care Department of the MST because, according to the law, nurses need to know at least in four weeks in advance their schedule. This means that even though ARIMA models were easy to apply, this kind of forecasting is not compatible with the hospital planning and scheduling tasks. Moreover, it appears that the forecasts are close to the average of the data. Therefore, for that former and the latter reasons, we think that there might be easier ways to forecast that can lead to similar results and that could be applicable to the MST situation. We study this in the next section.

5.2 Heuristics prediction

In Section 5.1 we apply ARIMA models to statistically forecast the demand (number of beds occupied) in the General and Thorax units of the ICD. However, we have seen that ARIMA models are not easy to apply due to the statistical difficulty, all the involved steps and the deep analysis that has to be done. Moreover, we mentioned that due to the law, they are not

applicable because nurses need to know their schedule in four weeks in advance. Therefore, we want to try simpler ways to forecast (doing heuristics) that might lead to similar results and that could be applicable to the MST situation.

In this section, we use three different approaches: naive forecast of next week, moving averages forecast and seasonal indexes. First we describe them (Section 5.2.1), then we compare them with the ARIMA approach (Section 5.2.2) and, finally, we compare the three heuristics approaches among them (Section 5.2.3).

We do the forecasting for the entire year 2015. However, to compare them to ARIMA models, we will use the predictions of the third and fourth quarters of 2015.

5.2.1 Description of the approaches

Naive forecast for next week

Despite the fact that the forecasting should be done at least in four weeks in advance, we want to try a naive forecast for next week prediction to see how good this easier approach performs in comparison to ARIMA models. For this, we assume that the week we want to forecast is going to behave exactly as the previous week. So, following the example in Figure 5.13, if we want to forecast the daily average of beds occupied in the third week of April (highlighted in orange), we calculate the daily average of the previous week (highlighted in blue) and we apply it to the week being predicted. This results in a forecast for next week and just a one week prediction because we need to update the data to perform further forecastings.

APRIL

Mo	Tu	We	Th	Fr	Sa	Su
		1	2	3	4	5
6	7	8	9	10	11	12
13	14	15	16	17	18	19
20	21	22	23	24	25	26
27	28	29	30			

Figure 5.13: Example of the naive forecast for next week

Moving averages forecast

The second approach that we want to try is very similar to the previous one but it was two main differences. The first one is that the prediction is done in five weeks in advance and, the second one, is that we assume that the predicted week will behave like the two most recent weeks of data that we have (we use a moving average of 2 weeks). We make this assumption because if we only use data from the most recent week (that would be a naive forecast of five weeks ahead), we are assuming that what is going to happen will be the same as five weeks ago. However, we have seen that, for example, summer and winter weeks may differ. Then by doing an average of two weeks we smooth this effect. For an easier explanation

of the approach we will follow the example in Figure 5.14. If we want to forecast the daily average of beds occupied for the fourth week of May (highlighted in orange), we go five weeks backwards and we calculate the average of the daily number of beds occupied of the last two more recent weeks (highlighted in blue). Then we assume that the daily average of beds occupied in the week being predicted will be the calculated average. This results in a five weeks ahead forecast and just one week prediction because, as before, we need to update the data for further predictions.

APRIL							MAY						
Mo	Tu	We	Th	Fr	Sa	Su	Mo	Tu	We	Th	Fr	Sa	Su
			1	2	3	4	5				1	2	3
6	7	8	9	10	11	12		4	5	6	7	8	9
13	14	15	16	17	18	19		11	12	13	14	15	16
20	21	22	23	24	25	26		18	19	20	21	22	23
27	28	29	30					25	26	27	28	29	30
								31					

Figure 5.14: Example of the naive forecast in five weeks ahead

Seasonal indexes forecast

The third approach that we use consists in employing seasonal indexes. These indexes recognize the seasonal variation (weekly) during the year of the variable studied. The steps to perform this kind of forecasting (Cote and Tucker, 2001) are:

1. Average of historical data: consists in computing the average of the variable of interest of the historical data available. In this case, we compute the average of the daily average of beds occupied per week from the 1st week of 2012 to the 52nd week of 2014. Hence, we do a 156-week average.
2. Calculate seasonal indexes: the seasonal indexes are calculated for each week of the year (52 indexes in total). To calculate them we divide the average of all the weeks x by the average of the historical data calculated in Step 1. For example, if we calculate the seasonal index of week $x = 2$, we do the average of week 2 over they years (2012, 2013 and 2014) and we divide it by the average of historical data. The formula for this step is shown in Expression 5.5.

$$seas.index\ week\ "x" = \frac{(week\ x_{2012} + week\ x_{2013} + week\ x_{2014}) / 3}{Average\ of\ historical\ data} \quad (5.5)$$

3. Deseasonalize historical data: in this step we remove the seasonality of the data by dividing each week's value by its corresponding seasonal index.
4. Calculate trendline based on deseasonalized data: this step consists in doing a regression analysis with all the deseasonalized data (from 2012 to 2014) to obtain the trendline coefficients (intercept and slope).
5. Calculate forecast: we calculate the daily average of beds occupied per week using Expression 5.6. Numbered week of forecast refers to the number of the week ranging from 1 to 52 (depending on which week we are calculating).

$$y_t = \text{seasonal index} * [\text{intercept} + (\text{slope} * \text{numbered week of forecast})] \quad (5.6)$$

The result of this approach is a forecast that can be done in one year ahead and we can predict an entire year.

5.2.2 Comparison with ARIMA models

Table 5.3 shows the RMSE values obtained from the different ways of forecasting studied so far. The RMSE values correspond to the root mean squared error of the forecasting for the night shift from the 27th to the 52nd week of 2015. As we can see, ARIMA models are the ones that perform better compared to the new ways we have studied since the error is smaller. Nevertheless, as we explained at the introduction of this section, due to the difficulty that the forecast are close to the mean of the data, we discard ARIMA models and we will only focus, from now on, on the naive, moving averages and seasonal indexes approaches.

<i>Approach</i>	<i>Thorax unit</i>	<i>General unit</i>
Non-seasonal ARIMA	1.06	1.39
Seasonal ARIMA	1.08	1.42
Naive forecast	2.04	2.04
Moving Averages forecast	2.20	2.72
Seasonal indexes	2.10	2.04

Table 5.3: RMSE values for the different methodologies used to forecast. Night shift from 27th to 52nd week of 2015

5.2.3 Comparison among the heuristics

Once ARIMA models are rejected, we focus on the other approaches with which we have done the entire forecasting for 2015. We have predicted the daily average of beds occupied per week from the 1st of January to the 31st of December of 2015. As we can see in Table 5.4, the seasonal indexes approach is the one that performs the best for both the Thorax and General units with an average RMSE of 2.04 and 2.03 respectively. For example, a RMSE of 2 implies that everyday there is an average mistake of 2 beds.

<i>Approach</i>	<i>Thorax unit</i>				<i>General unit</i>			
	Night	Day	Evening	Av.	Night	Day	Evening	Av.
Naive forecast	2.13	1.96	2.22	2.10	2.24	2.18	2.22	2.22
Moving Averages forecast	2.31	2.12	2.38	2.27	2.52	2.43	2.51	2.49
Seasonal indexes	2.10	1.88	2.15	2.04	2.05	1.98	2.06	2.03

Table 5.4: RMSE values for the naive, moving averages and seasonal indexes forecasts. Period from 1st January to 31st December 2015. Av.=Average.

Regarding the feasibility of introducing this approach in the Intensive Care Department, the naive and moving averages heuristics present a big obstacle. These two heuristics can

be introduced at the operational level, but not at the tactical level because decision about staffing are made at with more time in advance. Moreover, if this was not the case, the naive forecast would not be able to be implemented anyway. This is because it makes a forecast for the next week, which, by law, it is not allowed to plan and schedule nurses in such a small period of time. At least they have to be scheduled in four weeks in advance. Although these two heuristics are not suitable for our situation, they might be useful for other situations, thus the ICD manager's decision is to continue the study with both approaches as well. With respect to the seasonal indexes approach, it seems to be the most suitable one because it has the lowest RMSE, it is easy to apply, it allows us to forecast in one year ahead and to do the prediction for an entire year.

5.3 Conclusions

In Chapter 5 we have used statistical and heuristics methods to predict the daily average of beds occupied per week in the Thorax and General units. Regarding the statistical approach, despite the fact that we found models that are mathematically correct, none of them has resulted to be as useful as we thought for forecasting since they have not captured the variations of our variable as we expected. Although we have only done the analysis for the night shift, as the other shifts were very similar, we can assume that these conclusions are also valid for the day and evening shifts. With respect to the heuristics approaches, they perform a little bit worse than the ARIMA models, but they are easier to apply and they do not involve as much time. The naive forecast and moving averages are not implementable methods at a tactical level, but the seasonal indexes is and it performs better than these two. Hence, we can conclude that the use of ARIMA models implies a lot of difficult and time demanding steps that are not worth the results obtained. Contrarily, the seasonal indexes heuristics is easy to implement and performs quite good taking into account that it's an heuristic. In the next chapter, we see how the heuristic methods are translated into number of nurses and the performance they would have if they would be implemented.

Chapter 6

Desired magnitude of the pool of nurses

Calculating the magnitude of the pool of nurses needed is one of the most important parts of this project. This will indicate us the amount of personnel resources (in terms of direct-care nurses) we need for an effective and efficient care.

In the previous chapter we have discussed about the feasibility of demand forecasting for a demand based staffing. Surprisingly, statistics models have not resulted to be as useful as we expected, whereas the heuristics used seem to be quite good for demand prediction and easier to apply. Nevertheless, we believe it is important to not only limit the research to demand based staffing, but also to investigate other approaches that could offer different advantages and disadvantages. Therefore, in this chapter we determine the pool of nurses (total number of direct care nurses to employ) that was needed in 2015 based on different demand and capacity approaches. We consider three of them: staffing based on demand, staffing based on capacity and staffing based on both demand and capacity (hybrid approach). Within each approach we study different scenarios, which are explained later.

Section 6.1 gives an overview of the three approaches, scenarios and the steps we go through to calculate the nurses. Sections 6.2, 6.3 and 6.4 explain in more detail the steps followed for each approach and, finally, Section 6.5 presents the results and discussion.

6.1 Overview of the approaches, scenarios and steps

Step 1: Beds occupied

The first step consists in calculating the number of beds occupied depending on the approach and scenario used. These are:

- Demand: the demand scenarios predict the daily average number of beds occupied per week and shift as explained in Chapter 5.
 - Perfect Prediction: we use the real number of beds that were occupied in 2015.
 - Naive forecast: we predict for next week using the average of the current week.
 - Moving averages forecast: we predict in four weeks ahead using the data of the last two weeks.

- Seasonal indexes: we capture the weekly seasonality of the data and we forecast based on it.
- Capacity: the capacity scenarios calculate the number of beds occupied based on a percentage of the capacity.
 - 100% of capacity: we assume that all the beds are occupied
 - 85% of capacity: we assume that 85% of the beds are occupied
 - 80% of capacity: we assume that 80% of the beds are occupied
 - 75% of capacity: we assume that 75% of the beds are occupied
 - 70% of capacity: we assume that 70% of the beds are occupied
- Hybrid approach: it is a combination of capacity and demand staffing. We calculate the beds occupied based on a percentage of the capacity. Apart we use seasonal indexes for demand prediction. If the prediction says that more beds are occupied than the ones covered with the percentage of the capacity, then we substitute the current value for the higher value. Otherwise we leave it as it is.
 - 85% of capacity & 15% seasonal indexes: we assume that 85% of the beds are occupied for sure and the rest depends on the demand prediction
 - 80% of capacity & 20% seasonal indexes: we assume that 80% of the beds are occupied for sure and the rest depends on the demand prediction
 - 75% of capacity & 25% seasonal indexes: we assume that 75% of the beds are occupied for sure and the rest depends on the demand prediction
 - 70% of capacity & 30% seasonal indexes: we assume that 70% of the beds are occupied for sure and the rest depends on the demand prediction

Step 2: Nurses needed

With the next step we determine for each scenario the nurses needed every day and shift based on the calculations in Step 1. For that we need to take into consideration the ratios nurse-patients explained in Chapter 1. We divide the beds occupied in the specific shift by the corresponding ratio and we round the result in order to have an integer number of nurses giving care.

$$Nurses\ needed = round\left(\frac{Beds\ occupied}{Ratio}\right) \quad (6.1)$$

Step 3: Nurse hours

In this step we calculate the nurse hours worked in 2015 by adding all the nurses determined in Step 2 and multiplying them by 8 (Expression 6.2) because we assumed every shift implied 8 hours of work (explained in Chapter 1).

$$Nurse\ hours = 8 * \sum_{i=1}^{365} \sum_{j=1}^3 Nurses\ needed_{ij} , \quad (6.2)$$

where i represents the day and j the shift.

Step 4: Pool of nurses

This step consists in determining the total number of direct care nurses to employ in 2015 depending on each scenario. For that, we consider nurses as aggregate demand. Actually nurses have contracts with different hours, but to simplify the calculations and the understanding of the results, we only work with full-time employees (FTE). One FTE in 2015 had to work a total of 1,683.6 hours (holidays and others are already excluded). So, we divide the total nurse hours by the hours to work in 2015:

$$Pool\ of\ nurses = \frac{Nurse\ hours}{1,683.6} \quad (6.3)$$

Step 5: Extra and unnecessary hours

In this step we calculate the extra and unnecessary hours worked for each scenario. First we have to know in which cases we have had more or less nurses than needed. To find it out, we have to determine the actual ratios (Expression 6.4), which means to divide the real beds occupied in 2015 by the forecasted nurses specified with each scenario. We do this calculation for every day (i) and shift (j).

$$Actual\ ratios = \frac{Real\ beds\ occupied_{ij}}{Forecasted\ nurses\ needed_{ij}} \quad (6.4)$$

Then, based on the ratios' thresholds we agreed with the manager, we can say if we have a lack or excess of nurses, which will be translated into extra hours (if we have a lack) and into unnecessary worked hours (if we have an excess). The ratios thresholds are shown in Table 6.1. For instance, 1:2 means 1 nurse for 2 patients.

	<i>Night</i>	<i>Day</i>	<i>Evening</i>
Maximum	1:3	1:2	1:2.5
Optimal	1:2	1:1.5	1:1.75
Minimum	1:1.75	1:1	1:1.25

Table 6.1: Nurse-patient ratios' thresholds

6.2 Staffing based on demand

Figure 6.1 shows an example of the type of calculations we do to obtain the pool of nurses with the demand approach. This is the seasonal indexes scenario for the Thorax unit. We only show the values of two random weeks to make it easier to understand.

First, we forecast the beds occupied in 2015 using seasonal indexes. This corresponds to columns E , F and G . Then we calculate the nurses needed per day and shift using the optimal ratios mentioned in Table 6.1. This is shown in columns H , I and J . In order to calculate the nurse hours, we have to do the summation of all the values in columns H , I and J for all the year. Then we multiply the result by 8 and we obtain the total nurse hours. With this scenario, all the nurses together would have worked a total of 33,992 hours in 2015. Since each nurse had to work 1,683.6 hours in 2015 by law, to obtain the pool of nurses we just divide the nurse hours by the hours to be worked in 2015 for one nurse. The result suggests

Date	REAL BEDS OCCUPIED 2015			FORECASTED BEDS			FORECASTED NURSES			ACTUAL RATIOS		
	Night	Day	Evening	Night	Day	Evening	Night	Day	Evening	Night	Day	Evening
28/9/2015	3,00	3,00	6,00	6,17	6,36	5,82	3	4	3	1,00	0,75	2,00
29/9/2015	6,00	6,00	9,00	6,17	6,36	5,82	3	4	3	2,00	1,50	3,00
30/9/2015	9,00	9,00	8,00	6,17	6,36	5,82	3	4	3	3,00	2,25	2,67
1/10/2015	8,00	8,00	9,00	6,17	6,36	5,82	3	4	3	2,67	2,00	3,00
2/10/2015	9,00	9,00	9,00	6,17	6,36	5,82	3	4	3	3,00	2,25	3,00
3/10/2015	9,00	9,00	7,00	6,17	6,36	5,82	3	4	3	3,00	2,25	2,33
4/10/2015	7,00	7,00	5,00	6,17	6,36	5,82	3	4	3	2,33	1,75	1,67
5/10/2015	5,00	5,00	4,00	5,92	5,98	5,93	3	4	3	1,67	1,25	1,33
6/10/2015	4,00	4,00	5,00	5,92	5,98	5,93	3	4	3	1,33	1,00	1,67
7/10/2015	6,00	6,00	7,00	5,92	5,98	5,93	3	4	3	2,00	1,50	2,33
8/10/2015	7,00	7,00	9,00	5,92	5,98	5,93	3	4	3	2,33	1,75	3,00
9/10/2015	8,00	8,00	8,00	5,92	5,98	5,93	3	4	3	2,67	2,00	2,67
10/10/2015	8,00	8,00	6,00	5,92	5,98	5,93	3	4	3	2,67	2,00	2,00
11/10/2015	6,00	6,00	6,00	5,92	5,98	5,93	3	4	3	2,00	1,50	2,00

TOTAL NURSE HOURS	33992,00	EXTRA HOURS	800,00
HOURS TO WORK 2015	1683,60	UNNECESSARY HOURS	1976,00
POOL OF NURSES	20,19		

Figure 6.1: Example of a demand scenario (seasonal indexes) for the Thorax unit

that to be able to cover the demand we had in the Thorax unit in 2015 we needed 20.19 FTE nurses according to this scenario.

Besides, in order to calculate the extra and unnecessary worked hours, we need to take a look at the actual ratios we would have had if we would have implemented this scenario. For that we divide the real beds occupied in 2015 (columns *B*, *C* and *D*) by the forecasted nurses (columns *H*, *I* and *J*). The ratios are shown in columns *K*, *L* and *M*. Then we check which are the ratios outside the thresholds. In the case of Figure 6.1, within the period shown, we have many examples of nurse shortage (highlighted in red) and surplus (highlighted in green). To determine the lack or surplus we have to convert the ratios into nurses and then hours. An example of nurse surplus is the night shift of 6/10/2015. We forecasted that 3 nurses were needed but there were 4 beds occupied, which means that actually we needed 2 nurses. In this case, this supposes a surplus of one nurse, which converted into hours (multiplied by 8) it is 8 hours of work that could have been avoided and we have paid. An example of nurse shortage is the day shift of 30/9/2015. During that day shift, 9 beds were occupied, which means that we needed 6 nurses. However, we only forecasted 4 nurses. To have the optimal ratio, we need two more nurses, but actually with just one more nurse the ratio is within the thresholds because 9 beds divided by 5 nurses results in a ratio of 1.8 which is acceptable. Therefore, with one more nurse we would solve the issue but the ratio would not be optimal. Regarding the shortage of nurses we are just considering the quantity of nurses to make the ratios acceptable but not optimal. If we do this same calculation for all the ratios beyond the thresholds for all 2015, we obtain that we had a lack of nurses several times which implies a total of 800 extra hours and we had nurses that we did not need, which entailed 1,976 avoidable hours of work.

6.3 Staffing based on capacity

Figure 6.2 represents the capacity approach. Namely, in this case we are staffing on a 85% capacity in the Thorax unit. The steps followed for the capacity approach are the same ones as in the demand approach (Steps 2, 3, 4 and 5) except for the calculation of beds occupied (Step 1). Thus, we only explain how we determined the beds occupied because the rest has

already been explained in Section 6.2. Although the example is to staff based on 85% of the capacity, the other capacity scenarios are done with the same methodology. We already explained that the capacity (beds opened) changes depending on the month of the year, day of the week and shift. In the Thorax unit, the total capacity on Saturday evening, all Sunday and Monday night is 8 beds whereas for the rest is 10 beds. Therefore, we calculate 85% of 10 and 8 beds, which is 8.5 and 6.8 respectively (columns *E*, *F* and *G*). Then we determine the nurses needed and the nurse hours to be able to specify the pool of nurses. With this approach and scenario, 24.11 nurses would have been needed in 2015. Regarding the lack and surplus of nurses, there are a total of 8 extra hours and 3,944 hours of unnecessary work.

Date	REAL BEDS OCCUPIED 2015			FORECASTED BEDS			FORECASTED NURSES			ACTUAL RATIOS		
	Night	Day	Evening	Night	Day	Evening	Night	Day	Evening	Night	Day	Evening
12/1/2015	6,00	7,00	7,00	6,80	8,50	8,50	3	6	5	2,00	1,17	1,40
13/1/2015	7,00	7,00	10,00	8,50	8,50	8,50	4	6	5	1,75	1,17	2,00
14/1/2015	9,00	9,00	10,00	8,50	8,50	8,50	4	6	5	2,25	1,50	2,00
15/1/2015	10,00	10,00	10,00	8,50	8,50	8,50	4	6	5	2,50	1,67	2,00
16/1/2015	10,00	10,00	9,00	8,50	8,50	8,50	4	6	5	2,50	1,67	1,80
17/1/2015	9,00	9,00	8,00	8,50	8,50	6,80	4	6	4	2,25	1,50	2,00
18/1/2015	8,00	8,00	6,00	6,80	6,80	6,80	3	5	4	2,67	1,60	1,50
19/1/2015	7,00	7,07	8,00	6,80	8,50	8,50	3	6	5	2,33	1,18	1,60
20/1/2015	8,00	8,00	9,00	8,50	8,50	8,50	4	6	5	2,00	1,33	1,80
21/1/2015	9,00	9,00	9,00	8,50	8,50	8,50	4	6	5	2,25	1,50	1,80
22/1/2015	9,00	9,00	10,00	8,50	8,50	8,50	4	6	5	2,25	1,50	2,00
23/1/2015	10,00	10,00	10,00	8,50	8,50	8,50	4	6	5	2,50	1,67	2,00
24/1/2015	10,00	10,00	7,00	8,50	8,50	6,80	4	6	4	2,50	1,67	1,75
25/1/2015	7,00	8,00	8,00	6,80	6,80	6,80	3	5	4	2,33	1,60	2,00

TOTAL NURSE HOURS	40592,00	EXTRA HOURS	8,00
HOURS TO WORK 2015	1683,60	UNNECESSARY HOURS	3944,00
POOL OF NURSES	24,11		

Figure 6.2: Example of a capacity scenario (85% of capacity) for the Thorax unit

6.4 Staffing based on demand and capacity

In Figure 6.3 we can find an example of the hybrid approach. Concretely it is the 85% capacity & 15% seasonal indexes scenario.

Date	REAL BEDS OCCUPIED 2015			FORECASTED BEDS (85% cap.)			FORECASTED BEDS (S.I)			FORECASTED BEDS (Total)			FORECASTED NURSES (Total)			ACTUAL RATIOS		
	Night	Day	Evening	Night	Day	Evening	Night	Day	Evening	Night	Day	Evening	Night	Day	Evening	Night	Day	Evening
12/1/2015	6,00	7,00	7,00	6,80	8,50	8,50	7,87	8,03	7,74	7,87	8,50	8,50	4	6	5	1,50	1,17	1,40
13/1/2015	7,00	7,00	10,00	8,50	8,50	8,50	7,87	8,03	7,74	8,50	8,50	8,50	4	6	5	1,75	1,17	2,00
14/1/2015	9,00	9,00	10,00	8,50	8,50	8,50	7,87	8,03	7,74	8,50	8,50	8,50	4	6	5	2,25	1,50	2,00
15/1/2015	10,00	10,00	10,00	8,50	8,50	8,50	7,87	8,03	7,74	8,50	8,50	8,50	4	6	5	2,50	1,67	2,00
16/1/2015	10,00	10,00	9,00	8,50	8,50	8,50	7,87	8,03	7,74	8,50	8,50	8,50	4	6	5	2,50	1,67	1,80
17/1/2015	9,00	9,00	8,00	8,50	8,50	6,80	7,87	8,03	7,74	8,50	8,50	7,74	4	6	4	2,25	1,50	2,00
18/1/2015	8,00	8,00	6,00	6,80	6,80	6,80	7,87	8,03	7,74	7,87	8,03	7,74	4	5	4	2,00	1,60	1,50
19/1/2015	7,00	7,07	8,00	6,80	8,50	8,50	7,77	8,05	7,98	7,77	8,50	8,50	4	6	5	1,75	1,18	1,60
20/1/2015	8,00	8,00	9,00	8,50	8,50	8,50	7,77	8,05	7,98	8,50	8,50	8,50	4	6	5	2,00	1,33	1,80
21/1/2015	9,00	9,00	9,00	8,50	8,50	8,50	7,77	8,05	7,98	8,50	8,50	8,50	4	6	5	2,25	1,50	1,80
22/1/2015	9,00	9,00	10,00	8,50	8,50	8,50	7,77	8,05	7,98	8,50	8,50	8,50	4	6	5	2,25	1,50	2,00
23/1/2015	10,00	10,00	10,00	8,50	8,50	8,50	7,77	8,05	7,98	8,50	8,50	8,50	4	6	5	2,50	1,67	2,00
24/1/2015	10,00	10,00	7,00	8,50	8,50	6,80	7,77	8,05	7,98	8,50	8,50	7,98	4	6	5	2,50	1,67	1,40
25/1/2015	7,00	8,00	8,00	6,80	6,80	6,80	7,77	8,05	7,98	7,77	8,05	7,98	4	5	5	1,75	1,60	1,60

TOTAL NURSE HOURS	41000,00	EXTRA HOURS	8,00
HOURS TO WORK 2015	1683,60	UNNECESSARY HOURS	4192,00
POOL OF NURSES	24,35		

Figure 6.3: Example of a hybrid scenario (85% of capacity and 15% seasonal indexes) for the Thorax unit

As before, the only step that differs from the other approaches is Step 1. With this approach, first we calculate the beds based on the 85% of the capacity to make sure that at least we will plan nurses to cover this capacity percentage (columns *E*, *F* and *G*). Then we use the seasonal indexes forecasting (columns *H*, *I* and *J*) and we compare the beds forecasted with both methodologies. If the seasonal indexes predicts more beds occupied than the capacity approach, then we substitute the lower value for the higher one, otherwise we leave the beds occupied with the capacity approach. This is shown in columns *K*, *L* and *M*. To make it more understandable we explain an example. For instance, on the one hand, in the night shift of 12/1/2015, 6.8 beds are assumed to be occupied with the capacity approach (column *E*). However, the seasonal indexes predicts that 7.87 are occupied (column *H*). So, at the end we say that 7.87 beds are occupied in that specific shift and day (column *K*). On the other hand, in the night shift of 13/1/2015, the capacity approach covers 8.5 beds (column *E*) and the seasonal indexes says 7.87 beds are going to be occupied (column *H*). Since the capacity approach value is higher, we leave this one (column *K*). Once we have the total forecasted beds (columns *K*, *L* and *M*) we determine the nurses needed and the nurse hours. Then we get that the pool of nurses with this scenario should be 24.35. At the end of the year a total of 8 extra hours would have been worked and we could have avoided working 4,192 hours.

6.5 Results and discussion

6.5.1 Definition of the viewpoints

In the end of the project we have to decide which is the best scenario for the ICD. However, this decision is made based on the analysis of different viewpoints. These are:

- *Waste point of view*: in any process, waste is not desirable. Likewise, in this project we want to see which scenarios present the least waste and, thus, allocate the resources better. We represent waste as a function of the extra hours and the unnecessary hours worked. We consider that having extra hours is worse than having unnecessary hours worked, therefore, we choose to penalize the extra hours with a weight (a factor 2). Expression 6.5 shows the function used to determine the waste. This point of view is discussed in this chapter.

$$Waste = 2 \cdot Extra\ hours + Unnecessary\ hours \quad (6.5)$$

- *Financial point of view*: Every organization tries to reduce their costs as much as possible. In Chapter 7 we discuss the achievable cost savings with each scenario and, with this, we determine the best scenario from a financial point of view.
- *Practical point of view*: apart from what numbers and theory say, we have to take into account that our project is set in a highly complex environment. In order to carry out this project we have simplified some aspects but the reality is much more complex. In Chapter 8 we discuss qualitatively the best scenario from a practical point of view.

Even though in the end we have to decide which is the best scenario for the whole ICD, because we cannot staff using different approaches and scenarios for each unit, we first analyze each unit separately.

6.5.2 Thorax unit

Figure 6.4 shows the summary of the nurse hours, pool of nurses, extra hours and unnecessary hours for the Thorax unit. We can see that the scenarios with the least waste (column $2 * EH + UH$) are the 75% and 70% capacity scenarios (both with the same values). The waste, as formulated in Expression 6.5, adds up to 2,744. In the case of the Thorax unit, it does not make a difference to staff based on 75% or 70% capacity because although the number of beds differs, the calculation of the nurses rounded gives the same amounts. The second scenarios with the least waste are the 75% and 70% hybrid scenarios, both with the same values as well. The waste of both is 3,152. It is important to remark that meaning of the extra hours (EH). For example, 8 extra hours mean that on one specific day in one specific shift there was a shortage of one nurse. More extra hours can mean either there was a shortage of more than one nurse on a specific day and shift, that several times during the year there was a shortage of one nurse, or both things at the same time.

		THORAX UNIT						
		Scenario	Total nurse hours	Pool of nurses (FTE)	Extra hours (EH)	Unnecessary hours (UH)	$2*EH+UH$	Reason for EH
DEMAND	Perfect prediction		36,680	21.79	0	0	0	-
	Naive forecast		36,120	21.45	624	2,720	3,968	OVC/Inc
	Moving Averages		36,952	21.95	496	3,272	4,264	OVC/Inc
	Seasonal Indexes		33,992	20.19	800	1,976	3,576	OVC/Inc
CAPACITY	100% capacity		48,512	28.81	0	8,896	8,896	-
	85% capacity		40,592	24.11	8	3,944	3,960	OVC
	80% capacity		37,672	22.38	88	3,080	3,256	OVC
	75% capacity		34,752	20.64	384	1,976	2,744	OVC/Inc
	70% capacity		34,752	20.64	384	1,976	2,744	OVC/Inc
HYBRID	85% capacity & demand		41,000	24.35	8	4,192	4,208	OVC
	80% capacity & demand		38,368	22.79	72	3,392	3,536	OVC
	75% capacity & demand		36,296	21.56	272	2,608	3,152	OVC/Inc
	70% capacity & demand		36,296	21.56	272	2,608	3,152	OVC/Inc

Figure 6.4: Summary of the nurse hours, pool of nurses, extra and unnecessary hours and the waste obtained with the different scenarios for the Thorax unit in 2015. OVC= Overcapacity. Inc= incapacity to cover demand

Scenarios with extra and unnecessary hours reflect a non-optimal nurse allocation. However, analyzing the data we have found that in several cases it was not that the scenarios did not do an appropriate nurse allocation but that there was overcapacity. Recalling the bed occupancy rates for the Thorax unit in Chapter 2, several times the occupancy rates were above 100%, which we already said it was not possible but this is what our data shows. Therefore, some of the extra hours -that the different scenarios- have are not due to a bad performance, but to the fact that they do not take into account the possibility of having more beds occupied than the total capacity, and they do not cover this type of demand. Nevertheless, we cannot omit the fact that some of the extra hours are triggered by the incapacity that the scenarios sometimes have to cover the actual demand. The last column of Figure 6.4 shows the reason why there are extra hours. “OVC” refers to overcapacity and “Inc” refers to the aforementioned incapacity.

Regarding the capacity approach, in the case that there were not extra hours due to overcapacity, we can demonstrate that there are some scenarios perfectly prepared to cover the demand even though the occupancy is 100% and we are staffing based on less capacity. Table 6.2 shows

the ability of the capacity scenarios to cover a demand of 100%. The second column of the table shows the capacity (number of beds) on which we are basing our staffing. In the third column we can see the number of beds occupied. Thus, the rate of occupancy is 100%. The next three columns show the nurses needed based on the scenarios and, finally, the last three columns reflect the ratios we would find with all the beds occupied and our calculated nurses. We can see that for all the scenarios, except for the evening shift of the 75% and 70% capacity scenarios, all the ratios are within the thresholds. This means that although we are staffing based on a reduced capacity, we still have the ability to cover the maximum possible demand. This results are also applicable to the hybrid approach since it has the same number of nurses as the capacity approach, or even higher (due to the forecasts). Due to the aforementioned reasons, all the extra hours that we find in Figure 6.4 the 85% and 80% capacity and hybrid approaches should not exist. They are only due to the overcapacity found in the data.

<i>Scenario</i>	<i>Planned capacity</i>	<i>Beds occupied</i>	<i>Nurses</i>			<i>Actual ratios</i>		
			Night	Day	Ev.	Night	Day	Ev.
100% capacity	10	10	5	7	6	2.0	1.4	1.6
	8	8	4	5	5	2.0	1.6	1.6
85% capacity	8.5	10	4	6	5	2.5	1.6	2.0
	6.8	8	3	5	4	2.6	1.6	2.0
80% capacity	8	10	4	5	5	2.5	2.0	2.0
	6.8	8	3	4	4	2.7	2.0	2.0
75% capacity	7.5	10	4	5	4	2.5	2.0	2.5
	6	8	3	4	3	2.7	2.0	2.7
70% capacity	7	10	4	5	4	2.5	2.0	2.5
	5.6	8	3	4	3	2.7	2.0	2.7

Table 6.2: Ratios nurse-patient in the Thorax unit for the situation when there is 100% occupancy and the nurses are calculated with different scenarios

Despite of the overcapacity, the demand approaches are the ones with the most extra hours. This makes sense since the forecasts are not always good and we try to adjust staffing to activity as much as possible. What happens then is that the margin of error is small and we do not have enough nurses to cover the actual demand. Thus, we need extra hours. The hybrid approach has less extra hours than the capacity one because it considers more nurses. However, it also has more unnecessary hours worked.

Regarding the pool of nurses, the perfect prediction pool is 21.79 FTE. We find that this value is actually between the values of the 75% and 80% capacity scenarios (in Figure 6.4). If we recall the bed occupancy rates for the Thorax unit (Figure 2.6), in 2015 the most frequent rate was between 71% and 80%. This suggests that the demand approach captures this fact and tries to match staffing to activity. Moreover, this also suggests that staffing based on a lower capacity makes sense.

6.5.3 General unit

Figure 6.5 displays the summary of the pool of nurses, extra hours and unnecessary hours and the waste for the General unit. The scenarios with the least waste is the 75% capacity, with a waste of 832. In second place we find the seasonal indexes and 70% hybrid scenarios, with both a waste of 1,000.

GENERAL UNIT							
	Scenario	Total nurse hours	Pool of nurses (FTE)	Extra hours (EH)	Unnecessary hours (UH)	2*EH+UH	Reason for EH
DEMAND	Perfect prediction	73,136	43.44	0	0	0	-
	Naive forecast	72,720	43.19	88	1,096	1,272	OVC/Inc
	Moving Averages	73,192	43.47	72	1,960	2,104	Inc
	Seasonal Indexes	70,768	42.03	0	1,000	1,000	-
CAPACITY	100% capacity	90,024	53.47	0	8,008	8,008	-
	85% capacity	77,960	46.31	0	2,880	2,880	-
	80% capacity	72,504	43.06	0	1,168	1,168	-
	75% capacity	69,088	41.04	8	816	832	OVC
	70% capacity	61,320	36.42	624	200	1,448	OVC/Inc
HYBRID	85% capacity & demand	77,960	46.31	0	2,880	2,880	-
	80% capacity & demand	73,288	43.53	0	1,336	1,336	-
	75% capacity & demand	71,680	42.58	0	1,120	1,120	-
	70% capacity & demand	70,768	42.03	0	1,000	1,000	-

Figure 6.5: Summary of the nurse hours, pool of nurses, extra and unnecessary hours obtained with the different scenarios for the General unit in 2015. OVC= Overcapacity. Inc= incapacity to cover demand

The bed occupancy rates in the General unit surpass the 100% only two times (see Figure 2.7). Therefore, extra hours that we find in this unit are mainly due to an incapacity of the scenarios to cover the demand at certain moments. Most of the scenarios are also able to cover these two times of overcapacity. The ones that are not able to cover the overcapacity are the naive forecast, the moving averages and the 75% and 70% hybrid scenarios. For the General unit we can also show that in spite of staffing based on less capacity, the capacity based scenarios are capable of covering a demand of 100% because the ratios are acceptable (except for the 70% capacity scenario). In Table 6.3 we see how all the ratios are within the thresholds except for the day shift and evening shifts (this last one only when the capacity is 18) of the 70% capacity scenario. This is also applicable for the hybrid scenarios because they consider the same amount of nurses as the capacity ones or even more. Actually, in the General unit none of the hybrid scenarios had extra hours. Therefore, this means that they can handle a demand of 100% (all beds occupied).

With regard to the pool of nurses, the perfect prediction pool for the General unit is 43.44 FTE. We find this value between the values from the 80% and 85% capacity scenarios. Recalling the bed occupancy rates for the General unit (Figure 2.7), the most frequent rate in 2015 was between 71% and 80% but the interval between 81% and 90% were also common. This suggests, as before, that demand approaches try to match staffing to activity and that staffing based on a lower capacity makes sense.

Scenario	Planned capacity	Beds occupied	Nurses			Actual ratios		
			Night	Day	Ev.	Night	Day	Ev.
100% capacity	18	18	9	12	10	2.0	1.5	1.8
	17	17	9	11	10	1.9	1.5	1.7
85% capacity	15.3	18	8	10	9	2.25	1.8	2.0
	14.45	17	7	10	8	2.4	1.7	2.1
80% capacity	14.4	18	7	10	8	2.6	1.8	2.25
	13.6	17	7	9	8	2.4	1.9	2.1
75% capacity	13.5	18	7	9	8	2.6	2.0	2.25
	12.75	17	6	9	7	2.9	1.9	2.4
70% capacity	12.6	18	6	8	7	3	2.25	2.6
	11.9	17	6	8	7	2.8	2.1	2.4

Table 6.3: Ratios nurse-patient in the Thorax unit for the situation when there is 100% occupancy and the nurses are calculated with different scenarios

6.5.4 Intensive Care Department

So far we have analyzed both units separately, but we need to know what is the situation in the whole ICD because we cannot staff using different scenarios for the two units. Figure 6.6 displays the sum of the results from both units. For the whole ICD, the best scenario from a waste point of view seems to be the 75% capacity with a total waste of 3,576. In second place we find the 70% hybrid scenario with a waste of 4,152. We can see that either in the Thorax and the General units, the 75% capacity scenario was found among the first best ones and the 70% hybrid scenario among the second best ones. This implies that the needs of both units are similar, which eases the decision making for the whole department.

INTENSIVE CARE DEPARTMENT							
	Scenario	Total nurse hours	Pool of nurses (FTE)	Extra hours (EH)	Unnecessary hours (UH)	2*EH+UH	Reason for EH
DEMAND	Perfect prediction	109,816	65.23	0	0	0	-
	Naive forecast	108,840	64.64	712	3,816	5,240	OVC/Inc
	Moving Averages	110,144	65.42	568	5,232	6,368	OVC/Inc
	Seasonal Indexes	104,760	62.22	800	2,976	4,576	OVC/Inc
CAPACITY	100% capacity	138,536	82.28	0	16,904	16,904	-
	85% capacity	118,552	70.42	8	6,824	6,840	OVC
	80% capacity	110,176	65.44	88	4,248	4,424	OVC
	75% capacity	103,840	61.68	392	2,792	3,576	OVC/Inc
	70% capacity	96,072	57.06	1008	2,176	4,192	OVC/Inc
HYBRID	85% capacity & demand	118,960	70.66	8	7,072	7,088	OVC
	80% capacity & demand	111,656	66.32	72	4,728	4,872	OVC
	75% capacity & demand	107,976	64.14	272	3,728	4,272	OVC/Inc
	70% capacity & demand	107,064	63.59	272	3,608	4,152	OVC/Inc

Figure 6.6: Summary of the nurse hours, pool of nurses, extra and unnecessary hours obtained with the different scenarios for the ICD. OVC= Overcapacity. Inc= incapacity to cover demand

Looking at the pool of nurses, for 2015 the optimal total number of nurses to employ was

65.23 FTE, whereas staffing based on 100% capacity supposed 82.78 FTE. This suggests that a reduction of nurses could be possible. Some scenarios present a smaller pool than the perfect prediction (this is the case of seasonal indexes, 75% capacity, 70% capacity, 75% hybrid and 70% hybrid). The perfect prediction scenario ratios are very close to the optimal ones. The fact that some scenarios have a lower quantity of nurses implies that they underestimate the nurses needed compared to what would be optimal, but the ratios nurse-patient are still within the thresholds. Thus, they are acceptable. The pool of nurses calculated with the hybrid approach are bigger than with the respective capacity scenarios. This makes sense since in the hybrid approach, we make sure that at least a percentage is covered, but we also add reinforcement in the case we predict that more beds are going to be occupied than the ones covered.

In general, the demand scenarios calculate a pool of nurses quite similar to the one from the perfect prediction. This means that the ratios are close to the optimal ones. Nonetheless, the allocation of nurses does not perform that well since we need to call nurses for extra hours and some nurses are working when it is not necessary. Regarding the hybrid scenarios, the fact that a percentage of the capacity is covered and it is possible to add reinforcement if needed, without being considered as extra hours, makes this approach less risky than the demand one. The pool of nurses in the hybrid scenarios are quite similar to the perfect prediction and the unnecessary hours do not differ a lot from the demand ones. The difference comes with the extra hours, because with the hybrid approach the allocation of nurses is better done. When it comes to the capacity approaches, if we bear in mind that the pool of nurses of the perfect prediction is 65.23 FTE, we see that this value is actually in between the 75% and 80% capacity scenarios pool of nurses value. Recalling the occupancy rates from Chapter 2, the most common rates for 2015 were between 71% and 80%. Hence, it makes sense to plan on less capacity. We could say that the perfect prediction staffs based on a capacity between 75% and 80% but with a perfect allocation of nurses. The capacity approach has the advantage with respect to the others that it would be simple to implement because we do not have to do any forecast.

6.6 Conclusions

In this chapter we have studied which would be the pool of nurses for the Intensive Care Department based on different approaches and scenarios. The demand scenarios have a pool of nurses close to the perfect prediction but the allocation of nurses is not that good. The hybrid approach is less risky than the demand approach and it presents less extra hours. Finally, the capacity approach is easy to implement and involves less waste than the hybrid approach. It makes sense to staff based on a lower capacity than 100%, because the most common occupancy rates in 2015 for both units were between 71% and 80%. We can conclude that the scenario with the best allocation of nurses, together with the least waste, is the 75% capacity scenario. It suggests a pool of nurses of 61.68 FTE, 3.55 FTE less than the perfect prediction and 20.6 FTE less than staffing based on 100% of capacity. So far, money has not taken part in our analysis and discussions. In the next chapter we study the cost-savings that each scenario provides and we determine the best scenario from a financial point of view.

Chapter 7

How to achieve cost savings

In this chapter we explain the implications that the scenarios presented in Chapter 6 have in economical terms . In Section 7.1 we show the prices paid per hour and the hours that have to be paid and, in Section 7.2, we present the results of the costs for each scenario.

7.1 Fees and budget

To be able to calculate the costs of the different pool of nurses, for the different approaches and scenarios presented in Chapter 6, first, we need to define the total hours that have to be paid. According to the Collective Labour Agreement (CAO), in 2015 full-time nurses had to be paid for 1,878 hours of work, although they just worked 1,683.6 hours because of holidays, vacations and others.

Nurses earn different salaries depending on their studies and years of experience. The cost for the hospital to have one “average salary” nurse working is 30.8 €/hour including the 45% of social security costs. Extra hours have a 70% extra cost for the hospital compared to the normal hours. This implies that an extra hour is paid at 52.36 €/hour (45% incl.). When excluding the 45% social security costs, a normal hour is paid at 16.94 €/hour and an extra hour at 28.80 €/hour. Table 7.1 summarizes the fees mentioned. Together with the head of nurses we agreed in doing the calculations based on an average salary and without the social security costs, in order to be able to make comparisons of the results obtained with the annual ICD budget.

<i>Social security</i>	<i>€/hour</i>	<i>€/extra hour</i>
45% incl.	30.8	52.36
45% excl.	16.94	28.80

Table 7.1: Average fees for normal and extra hours

The budget for the ICD roughly includes: salaries (nurses, managers, nurse students, secretaries, assistants and other administrative personnel), machines and instruments, material and laboratory. The value of the ICD budget for 2015 was 7,888,414 €, from which 3,050,702 € were invested in nurses.

7.2 Costs results

Figure 7.1 shows the costs for the different scenarios presented in Chapter 6. We have not considered the costs of extra hours due to illness or other situations. We have just taken into account the extra hours due to a shortage of nurses triggered by the nurse allocation in each scenario. The extra hours due to illness would be the same for each scenario, despite the fact that in some scenarios it would not be necessary to call an extra nurse, because it could be possible to handle the situation with the present nurses. However, we assume that they would be identical. Therefore, the total costs in case of illness would be higher, but the difference between the scenarios would remain the same. With the values in Figure 7.1 we are considering the minimum cost that the hospital would have for the ICD. Looking at the fifth column “Total (P+EH)”, the most expensive situation is staffing based on 100% of capacity. It results in a cost of 2,617,600 € for the whole ICD. All the other scenarios presented are cheaper compared to this one. Regarding the perfect prediction, when allocating the nurses perfectly, the total cost results to be 2,075,183 €. However, there are cheaper scenarios like seasonal indexes, 75% and 70% capacity and 75% and 70% hybrid. This is due to the reduced pool of nurses. The hybrid scenarios are more expensive than their respective capacity scenarios because the hybrid ones also take into account demand and therefore the pool of nurses is bigger as mentioned in Chapter 6. The last column in Figure 7.1 shows the costs that we would have obtained if we would have avoided working the unnecessary hours that were worked. Considering the unnecessary hours worked as waste, it proves that reducing the waste it is possible to reduce the costs.

INTENSIVE CARE DEPARTMENT						
	Scenario	Pool of nurses (P)	Extra hours (EH)	Unnecessary hours (UH)	Total (P+EH)	P+EH-UH
DEMAND	Perfect prediction	2,075,183 €	0 €	0 €	2,075,183 €	2,075,183 €
	Naive forecast	2,056,413 €	20,504 €	64,643 €	2,076,917 €	2,012,274 €
	Moving Averages	2,081,227 €	16,357 €	88,630 €	2,097,585 €	2,008,955 €
	Seasonal Indexes	1,979,425 €	23,038 €	50,413 €	2,002,463 €	1,952,050 €
CAPACITY	100% capacity	2,617,600 €	0 €	286,354 €	2,617,600 €	2,331,246 €
	85% capacity	2,238,067 €	230 €	115,599 €	2,238,297 €	2,122,699 €
	80% capacity	2,081,864 €	2,534 €	71,961 €	2,084,398 €	2,012,437 €
	75% capacity	1,962,246 €	11,289 €	47,296 €	1,973,534 €	1,926,238 €
	70% capacity	1,815,268 €	29,028 €	36,861 €	1,844,296 €	1,807,435 €
HYBRID	85% capacity & demand	2,247,929 €	230 €	119,800 €	2,248,160 €	2,128,360 €
	80% capacity & demand	2,109,859 €	2,073 €	80,092 €	2,111,933 €	2,031,841 €
	75% capacity & demand	2,040,506 €	7,833 €	63,152 €	2,048,339 €	1,985,187 €
	70% capacity & demand	2,023,009 €	7,833 €	61,120 €	2,030,842 €	1,969,723 €

Figure 7.1: Summary of the costs for each scenario in the ICD in 2015

From all the scenarios presented, the cheapest scenario is staffing based on 70% of capacity (1,844,296 €), whereas the most expensive one is staffing based on the whole capacity (2,617,600 €). The difference is 773,304 € which implies a cost reduction of 29.5%. From the financial viewpoint, the best scenario is staffing based on 70% of the capacity. We say that this is the best from the scenarios studied because if we consider scenarios below 70% capacity what would happen is that the pool of nurses cost would decrease because of a reduced pool of

nurses and the extra hours cost would increase because we would have a nurse shortage. This would result in an increased total cost and it would be more expensive than staffing based on 70% capacity.

7.3 Comparison with the ICD budget

The calculated cost of staffing based on 100% capacity is 2,617,600 €, whereas the real amount spent was 3,050,702 €. This entails a cost difference of 433,102 €. However, this variation is understandable bearing in mind that the costs presented in this chapter are the result of a simplified reality, in which we do not take into account nurses performing other tasks than direct care. This entails that these costs might not be the real ones but they give an idea of the magnitude of the achievable costs savings and, as just shown, the real and calculated values are not that different. The important concept to extract from this chapter is the fact that changing the staffing methodology allows us to reduce the costs, compared to the current staffing methodology.

7.4 Conclusions

In this chapter we see the economical implications of each scenario presented in Chapter 6. The most expensive scenario is staffing based on 100% of capacity, whereas the cheapest one is staffing based on 70% of capacity. Thus, we conclude that from a financial viewpoint the best scenario is the 70% capacity. With it, the substantial saving that we can obtain is 773,304 €. Nevertheless, the important message is that staffing using different methodologies can help us reducing costs.

In the previous chapter we defined the best scenario from a waste point of view, which was the 75% capacity scenario. In this chapter we have concluded that from a financial viewpoint the best one is the 70% capacity scenario. These two perspectives are not aligned, even though they do not differ that much. However, in the next chapter, we discuss about this difference and we define the best scenario from a practical point of view.

Chapter 8

Discussion

In this chapter, we first make a deep analysis in Section 8.1 of all the results obtained in the study by relating the different chapters and literature. In Section 8.2 we explain the limitations encountered during the development of this project and, finally, in Section 8.3 we suggest the further research that should be done to complement the project.

8.1 General discussion

8.1.1 Comparison of our results to the literature

Regarding the forecasting techniques, professor Van Geer suggested we needed at least 10 years of data to have a reliable estimate of the trend and seasonality. We have also noticed that in the forecasting examples proposed by Hyndman and Athanasopoulos (2014), they use several years of data. Nonetheless, in the literature found about forecasting in health care departments, we have seen examples in which they just use few years or even just months. There are also examples that use a considerable amount of years. In our case we have used 4 years of data which is more or less the average.

Now, let us compare our results to those from the literature. As basis for our comparison we use the RMSE, a forecast accuracy measure. Jones et al. (2002) and Champion et al. (2007) use ARIMA models and RMSE to assess the forecast accuracy, as we do. We have to bear in mind that the RMSE is a scale dependent error, therefore we cannot compare our results to other literature that has used RMSE but has analyzed the data in another scale. This is the case of Champion et al. (2007) in which they forecast the monthly number of patients whereas we forecast the number of beds occupied. Hence, because of the aforementioned reason, our studies are not comparable. Jones et al. (2002) do forecast the beds occupied and the results they obtain are worse than ours in terms of the forecasting error. With 11 years of data they obtain a RMSE of 12.6 for a seasonal ARIMA model. With 4 years of data we obtain RMSE of 1.08 and 1.42 for seasonal models and 1.06 and 1.39 for non-seasonal models. They say to have obtained a good forecast whereas we are not completely satisfied with ours, which is better in terms of the forecasting. We are not satisfied because, as explained in chapter 5, the forecasts are close to the mean of our data and do not show variations. We expected ARIMA methods to provide us with more interesting forecasts, rather than the mean of the data, which is different from what actually happened in 2015. Jones et al. (2002) predict the daily number of beds occupied due to emergency admissions. Hence a RMSE of 12.6 implies that everyday they commit an error of 12.6 beds, whereas we forecast the daily average number of beds

occupied per week and we make an average mistake of 1.2 beds. We do not know to what we can attribute this difference. It is surprising that despite having obtained such a good result compared to the literature, we are not satisfied. We actually wanted the ARIMA models to forecast more fluctuations with narrower confidence intervals. However, we might have had a wrong idea about our ARIMA models. The previous analysis to the forecasting suggested that the best possible prediction was one week ahead and when checking the long-term forecasting we did not see what we expected. Thus, we were reluctant to use ARIMA models. If we would have proceeded with a further analysis, as we did for the heuristics scenarios, maybe we would have obtained good results when translating the forecasted beds into nurses. Nevertheless, we think that all the work this approach requires is not worth when having other possibilities. On the one hand, apart from all the steps we have carried out to obtain the forecast, if we do other predictions later in the future, we need to redo all the steps. The ARIMA models we propose in Chapter 5 are optimal for the data we have used, but when we update the data, it is possible that the optimal models change. On the other hand, in Chapter 6 we see other approaches - discussed later - with which we consider we obtain good final results and they do not require as much work as ARIMA models. Jones et al. (2002) says that all the work to forecast with ARIMA models is worth because, in their case, naive forecasts perform much worse. However, in our case we believe our heuristics are good.

8.1.2 Consideration of other staffing approaches

In Chapter 2 we say that according to the bed occupancy rates, it could make sense to think about a demand based staffing approach, because there is room for improvement as the units are not saturated. Nonetheless, it is not until Chapter 6 that we do not think about the possibility of just staffing based on a lower capacity. After studying the demand scenarios and considering that just the demand approach could be risky, as we want to have certainty, we thought about merging capacity and demand staffing. As explained in Chapter 6, with this we make sure that at least we cover a percentage of the capacity for sure and then, using forecasting techniques, we reinforce the number of nurses if necessary. This new hybrid approach can be implemented at a tactical level. After analyzing the hybrid approach we realized that just staffing on a lower capacity was feasible and could provide good results.

8.1.3 Waste, financial and practical viewpoints

From a financial perspective, we determined in Chapter 7 that the best scenario is staffing based on 70% of the capacity. Nevertheless, taking into account organizational and quality aspects, this scenario does not provide us neither a comfortable planning nor optimal nurse-patient ratios. First, this scenario implies, for sure, the need to call extra nurses. As we demonstrated in Tables 6.2 and 6.3, the scenario is not capable of covering a 100% bed demand. This means that the managers might have to be constantly aware of calling extra nurses because of the “gap” that this scenario presents. Second, related to the first reasoning, the nurse-patient ratios are not optimal, actually they graze and even exceed the thresholds when there is demand for all the beds. Therefore, although the 70% capacity scenario could work, we refuse it for the aforementioned reasons.

From the waste point of view, we concluded in Chapter 6 that the best scenario is to staff based on 75% of the capacity, because it has the best allocation of nurses and least waste among all the scenarios. It does not require many extra nurses as well as it does not suppose

a lot of unnecessary worked hours as other scenarios. This implies that the nurses are quite adjusted to the activity. With regard to the ratios, when there is 100% occupancy, they are very close to the thresholds. If we compare the pool of nurses for this scenario (61.68 FTE) with the perfect prediction (65.23 FTE), the 75% capacity scenario considers less nurses, which implies that the ratios will be above the optimal ones. Even though numerically this scenario is good because the resource allocation is appropriate and the ratios are within the intervals, from a practical point of view we think this scenario is too adjusted. In order to make all the calculations of this research, we have simplified the reality. Therefore, staffing based on 75% of the capacity might not be enough if we take into account that reality is much more complex. Moreover, we also have to take into account that if the ICD decides to staff with another methodology, it supposes a big change. Changes are difficult to carry out and need to satisfy all the stakeholders to be successful. Thus, if the decision would be to staff based on less capacity, we would recommend to start with a scenario more similar to the 100% capacity. We suggest the possibility of taking into account the 85% capacity scenario. This scenario does not present extra hours (just the ones due to overcapacity) and the unnecessary worked hours are reduced in a 59.6% compared to the 100% capacity scenario. Regarding the ratios, if we staff nurses to cover 85% of the beds, but all of them are occupied, the ratios are within the intervals and quite close to the optimal ones. Therefore, we think this is a good scenario to start with, in the case of desiring a staffing methodology change, because it seems practically feasible, it reduces the waste and it does not imply a very big change compared to the other scenarios. This scenario does not require any forecasting skills and it works as the methodology used so far, so we do not think it would be difficult to put in practice. In the case it is implemented but in practice it does not work properly, it is possible to staff based on a higher capacity or to use an hybrid approach. We consider that the demand scenarios suppose a huge change and they are risky because of the uncertainty. We have seen that the demand scenarios are more or less like staffing based on a reduced capacity so, anyway, we would recommend to start first with a capacity scenario.

8.1.4 Comparison of the pool of nurses needed

It is important to remark that when we completed the research, we noticed that planners take some latitude in applying the guidelines to calculate the direct care nurses needed. Even though the official guidelines suggest specific nurse-patient ratios, in practice it is easier for planners to slightly deviate from the guidelines. We noticed this because the number of nurses needed did not match with the nurses calculated with our 100% capacity scenario. We have had only access the data that shows the amount of direct care FTE nurses needed for 2015 in the General unit, which was 66.23. The number we obtained for the General unit with the 100% capacity scenario was 53.47 FTE. This difference is due to two reasons:

1. They do not apply the guidelines ratios strictly. When applying them for the General unit we would need every day 12 FTE in the morning, 10.28 FTE in the evening and 9 FTE at night. Instead, they determine that 13 FTE are needed in the morning, 11 FTE are needed in the evening and 11 FTE are needed at night. We have noticed that in the morning they calculate the needed nurses based on the ratios and then they add one more nurse.
2. The amount of hours that the nurses had to work in 2015 differs in the calculations. We were told that nurses were paid for 1,878 hours but, counting holidays and others,

they worked 1,683.6 hours. Instead, they calculated the nurses needed based on 1,599.36 hours of work. We talked with the head of nurses about this variation but we could not find the cause.

Due to the aforementioned reasons, their calculations on the number of direct care nurses needed is bigger than ours.

8.1.5 Forecasting an unknown year

What we have discussed so far are the best scenarios for 2015 and we have analyzed them once we have had all the data needed. But what if we want to do it for an unknown year? How can we calculate the nurses that we need? We suggest to do a naive forecast which means to assume that the next year (let's say 2017) is going to be as the previous one (2016). Thus, for 2016 we should do a similar analysis to the project, calculate the pool of nurses for different scenarios and decide which one suits the best for the next year (2017). Nevertheless, with the data that we have, it would be appropriate to do this if the ICD was the same as in the old situation. But it has changed, so it is not adequate to derive conclusions of changes for the new ICD if we only have data from the old situation. We assume that the application of the different scenarios will be the same, because the ICD works as before, but the results will depend on the new occupancy rates. Maybe they are significantly higher, lower or they remain the same. Since we use the bed occupancy rates to see if there is room for change, we would need to analyze them for the new ICD. Once we had one year of data from the new situation, we could carry out the same analysis as performed in this project. However, this is only valid for the capacity approach because one year of data is enough for the calculations. For the demand approach we basically use seasonal indexes as we refused to use the other two scenarios. Seasonal indexes need more years of past reliable data to be able to predict. Hence, neither the demand nor the hybrid approaches could be studied with just one year of data from the new ICD. Furthermore, we must say that in May 2016 we found an error in the hospital information system. It did not take into consideration patients who were changed from one unit to another. The system only recorded the last unit where the patient had been, thus it seemed the patient had stayed in this last unit since he was admitted in the ICD. The results obtained processing this data were extremely high occupancy rates above 100%. This is how we realized that the system had a problem. The technicians solved the issue, but all the data from January 2016 to May 2016 was polluted, therefore, it is not usable. At least, the ICD has to wait until May 2017 to have one year of reliable data.

8.1.6 Implementation

So far we have talked about the advantages and disadvantages of each approach. Nevertheless, it is important to think about how the chosen scenario would be implemented. We propose the following steps for the implementation:

1. Measurement of the nurses' motivation: the main reason why this research was done is because the medical manager had the feeling that there were more nurses than needed and their motivation decreased when there was not a lot of workload. The first step we suggest is to measure the motivation of the nurses using questionnaires. This will allow us to later know how the change of the methodology affects the nurses.

2. Month test: before finally deciding if this methodology change is to be implemented, a real life test should be carried out to see how it would work. We suggest to do a test during one month. For that, some preparation needs to be done, in which a steering team formed by the main stakeholders (medical manager, the head of the nurses and a leader nurse) has to motivate the nurses to try it. Even though we have introduced this methodology change at a tactical level, it is possible to make a test with the 85% capacity scenario -if chosen- and plan it at the operational off-line level. In case of choosing the 85% capacity scenario, it would only be necessary to reduce the amount of nurses that work during the specified month. We propose to perform the test during August for two reasons. The first one is that it is a month with less workload, which can be useful to reduce the stress that this test can cause or to be able to better handle possible obstacles that might appear. The second reason is because some nurses might be looking forward to have more summer vacations if they are not needed during the test. If the test is carried out during August, a total of 51.8 FTE are needed per week: 16.8 FTE for the Thorax unit and 35 FTE for the General unit. The calculations for these numbers can be found in the Appendix D.
3. Measurement of the test performance: once the test finishes, it is of high relevance to see how many corrections at the on-line operational level are needed. This means to investigate how many times there has been a shortage of nurses and extra nurses have been called. Furthermore, it is also important to analyze how adjusted the pool of nurses has been to the activity.
4. Measurement of the nurses' motivation after the test: once the test month finishes, it is important to measure the motivation of the nurses again to be able to see how this change has affected them and what they think about it.
5. Decision making: If the given feedback (test performance and nurse motivation) is positive, the test can be extended one more month. In case of having performed the test during August, it can continue during September. The positive aspect of continuing in September is that it is a good moment to introduce a new system, because a period starts again after summer and people might have fresh courage to carry it out. Moreover, the new methodology would be tested in a month of more workload. If the extended test has positive outcomes, then the change can be implemented at the tactical level. In the case the feedback is negative, then we recommend to refuse using the chosen scenario.

This change may not be palatable to the nurses because this means that the amount of nurses needed in the ICD is reduced. Thus, they can be afraid of being fired. Nevertheless, it is necessary that the steering team motivates them to perform the test. In the sequel, we present motivation arguments:

- The test (and maybe change) is carried out in order to have a better alignment between nurses and number of beds occupied. This leads to a more evenly spread workload, which implies that nurses will have a more constant work pace.
- If with the current staffing methodology extra nurses are needed when unexpected events happen, with the new scenario, which contemplates less nurses, it is more likely that even more extra nurses will be needed. The extra hours worked are better paid and this can motivate the nurses.

- Nurses may fear to be fired with this change. However, by natural cause maybe some nurses leave, some retire, others are pregnant... so it is possible that, at the end of the year, no more nurses are hired but also none are fired.
- It is possible that instead of firing nurses, they can be allocated to other departments of the hospital in which they are needed.

It is of relevance to motivate and involve the nurses in the test, because they can have a negative idea about such change, generated by the fear of being fired or because they do not understand the reason for changing. If they are reluctant to it, this will be reflected in the questionnaires performed after the test. Nonetheless, it is important to separate the concepts of how they felt with the new methodology regarding the day-to-day work (workload, work pace...) and what their fears and thoughts are about the new methodology.

8.2 Limitations

When it comes to the limitations of the project, the first one we had was the lack of data for the new ICD. However, we thought we could use the few data we had to have an insight about the new situation. The second limitation appeared when the current data was unreliable as aforementioned. Hence, we had to change the way to tackle this project. One might think that another limitation this project has, is the fact of predicting emergency and elective patients together. Even though the bed occupation by an elective patient, who comes from the operating room, can be predicted in advance, we have two reasons why we do not do such division. The first one is because there is a time difference between when the pool of nurses has to be calculated (it was to be known before the beginning of the year) and when the elective surgeries are planned. This means that when the calculations are done, most of the elective surgeries are not known. The second reason is because we know that in the MST a great percentage of elective surgeries are not done in time. Thus, these patients become much less predictable.

8.3 Further research

Since this research is set in a highly complex environment, we have simplified the reality in order to be able to perform all the calculations. This means that we have not taken into account factors that could complicate the analysis of the scenarios. As a future research it is possible to simulate all the scenarios using software such as *Plant Simulation* from Siemens. In this simulation, aspects that we have not considered in this project could be integrated. For example, it would be possible to simulate when nurses get sick, when nurses have to leave the ICD because there is a reanimation in another ward, nurses that perform other tasks apart from direct care, etc. This simulation would reflect in a more realistic way how the scenarios perform and the extent to which they are implementable.

8.4 Conclusions

In this chapter we do an extended discussion about the results obtained and the implications they have. Comparing our work to the literature we obtain better results. Moreover, the best scenario when analyzing the allocation of nurses and waste seems to be staffing based on 75% of the capacity, whereas from a financial point of view, staffing based on 70% of the

capacity is preferred. From a practical perspective, taking into account that changes are not easy and the we have done the calculations based on a simplified reality, we think that staffing based on 85% of the capacity is more appropriate because it is easy to implement, it does not involve a big organizational change and the waste and costs are reduced. Even though we have deeply focused on demand during the project, the allocation of nurses that we achieve with the demand approach is not as good as with other approaches.

Chapter 9

Conclusions and recommendations

In the framework of completing the master thesis of Health Sciences, I performed research at the Intensive Care Department (ICD) of the Medisch Spectrum Twente (MST) into studying if demand nurse staffing is feasible and which other ways of staffing could be used in order to improve the current staffing method used. This led to the following central research question:

“ In what way can the current capacity based staffing be improved in the Intensive Care Department of the MST?”

In order to answer the central research question we propose a number of sub-questions. In the sequel we first elaborate on answering the sub-questions. Next, we present -based on the answers provided- our solution to the central research question.

1. *What has been the situation of the ICD regarding patient admissions and length of stay from 2012 to 2015?.*

The situation in the old MST building was that more patients were admitted in the Thorax unit and more patients were discharged within the first 2 days than in the General unit. This makes sense recalling the fact that 90% of the admissions in the Thorax unit were patients coming from the OR, whereas in the General unit the vast majority were emergency patients.

2. *Which are the bed occupancy rates of the Intensive Care Units?*

The bed occupancy rates have tended to decrease along the years. The most frequent ones are rates between 71% and 90%, which makes us think that it is possible to do a staffing methodology change because the units are not always saturated.

3. *What does literature say about demand forecasting in the health care sector?*

In health care, the most common statistical approach used is time series analysis. One of the most-widely used methods is the application of ARIMA models. Even though the different studies examined report to have obtained good forecasting results, they only consider forecasting as an additional tool to improve resource allocation and planning.

4. *Which other methodologies can be used for staffing?*

Apart from demand forecasting we also consider staffing based on different percentages of capacity and an hybrid approach, which takes into account capacity and demand. Demand scenarios seem to be risky due to the amount of extra hours and unnecessary hours of work that they have. Staffing based on different percentages of capacity reduces

the waste the most and these scenarios are less risky than the ones of the demand approach. Finally, the hybrid scenarios present more waste than the capacity ones, but the risk of not having enough nurses is reduced compared to the demand and capacity approaches.

5. *How accurately can the demand of the Intensive Care Units be predicted?*

First we use ARIMA models to forecast the demand for 2015. A first analysis suggested that the best possible prediction was one week ahead because only data that had one week difference was correlated. However, assessing the accuracy of the long-term forecast with the RMSE we can say that numerically we have obtained forecasts which are accurate because the RMSE is low. Despite the good accuracy, at a first sight we were not satisfied with the predictions, this is why we considered other ways of forecasting. Using a naive forecast, a moving average and seasonal indexes, we have obtained worse RMSE values but they are still low. Therefore, we think all the work that has to be done for ARIMA models is not worth if we have these other forecasting techniques.

6. *How large is the pool of nurses needed in the ICD based on demand and capacity?*

To answer this question we have considered different approaches and scenarios. As explained in the introduction the pool of nurses is based on FTE and it only considers direct care nurses. In the hypothetical case of doing a perfect prediction the pool of nurses would be 65.23 FTE. If we staff based on the total number of beds, then we obtain a pool of 82.28 FTE. This implies a difference of 20.05 nurses. All the other scenarios considered are in between these values, although some of them suggest a smaller pool than the perfect prediction because they underestimate the nurses needed. With this difference we see that there are more nurses than needed and it is possible to cover the demand and work with less nurses. However, we have to bear in mind that we have simplified the reality for the calculations.

7. *Which are the cost differences between staffing based on capacity, demand and the hybrid approach?*

If we could achieve a perfect prediction, the total cost (without 45% taxes) would be 2,075,183 €. Staffing based on the maximum capacity, the cost would rise to 2,617,600 €. This implies a cost difference of 542,417 €. With the remaining considered scenarios the cost is always lower than staffing based on the maximum capacity. The hybrid approach is in general more expensive than the demand approach and, with respect to the capacity approach, the equivalent capacity scenarios are cheaper than the ones of the hybrid approach. We see that it is possible to reduce the costs if we change the staffing methodology.

8. *Which are the waste, financial and practical viewpoints of the different methodologies considered for staffing?*

We have chosen a best scenario based on each point of view. From a financial perspective, we determine that the the 70% capacity scenario is the best one due to the reduced costs. However, we discard this scenario because of its disadvantages. From a waste viewpoint, the best scenario is the 75% capacity one since it has the least waste and, thus, a good allocation of resources. Nevertheless, from a practical perspective, we think that this last scenario is too adjusted because we have simplified the reality. Therefore, taking this into account, and also that changes are difficult, we think that the 85% capacity scenarios is the best one from a practical point of view.

The answer to the research question “*In what way can the current capacity based staffing be improved in the Intensive Care Department of the MST?*” is the following:

In this project we have introduced three approaches for staffing. These are:

- (i) *Demand approach*: consists in predicting the daily average number of beds occupied per week and shift using data from the past.
- (ii) *Capacity approach*: consists in calculating the number of beds occupied based on a percentage of the capacity. The different scenarios assume that, for sure, a certain percentage of the capacity is occupied. Thus, they staff based on this percentage. However, all the beds are still open.
- (iii) *Hybrid approach*: It is a combination of capacity and demand staffing. First, staffing is done based on a percentage of the total capacity. Then, a reinforcement of this staffing is done based on demand forecasting.

We think that the *demand approach* is too risky due to the amount of extra hours and unnecessary hours of work that it has. Hence, we consider that the resource allocation is not properly done with this approach. Moreover, two of the presented scenarios cannot be applied at a tactical level. Another disadvantage of this approach is that it requires past reliable data of some years to be able to forecast. Nevertheless, it is not always possible to have it. Regarding the *capacity approach*, we have seen that it is possible to improve the current staffing method, using the capacity approach, by just reducing the percentage of capacity based on which the staffing is done. In other words, if before, the number of necessary nurses was done based on 100% of the capacity, i.e., based on all the open beds, it is possible to decrease this percentage and staff based on a lower amount of beds. However, still keeping the original number of beds open. This is because demand fluctuates over time and the units are not always full. Moreover, almost all the capacity scenarios proposed are able to handle 100% of the demand, even though they do not staff based on 100% of capacity. Using some of the presented capacity scenarios would reduce the waste and the costs. Furthermore, this approach is easy to implement. Finally, the *hybrid approach* presents more waste than the capacity approach, but the risk of not having the needed amount of necessary nurses is reduced. The disadvantage of this approach is the need of past reliable data to forecast the demand.

Therefore, with the central research question answered, and taking into account the sub-research questions, we recommend to consider the capacity approach and, from this, the 85% capacity scenario. Regarding the implementation, we suggest to make a real life test, in which during a month the amount of nurses working is reduced. It is important to measure the motivation of the nurses before and after the test using questionnaires, as well as to measure the performance of the test. This will be valuable information for decision making with regard to the change of methodology. Nevertheless, this is just a first recommendation in case of thinking about a change and its implementation. We consider that, before the real life test, further research using simulation should be done to make sure that such scenario would work in the complex reality.

Even though this analysis and recommendation are done based on the old situation of the ICD, taking into account the assumption mentioned in the introduction, it is possible to

change the staffing method in the new ICD and get benefits from it. The assumption was: if an approach would have been effective in the old ICD, it will be effective in the new ICD as well.

Bibliography

- Abdel-Aal, R. and Mangoud, A. (1998). Modeling and forecasting monthly patient volume at a primary health care clinic using univariate time-series analysis. *Computer Methods and Programs in Biomedicine*, 56(3):235–247.
- Armstrong, J. S., Green, K. C., et al. (2005). Demand forecasting: Evidence-based methods. Technical report, Monash University, Department of Econometrics and Business Statistics.
- Boyle, J., Jessup, M., Crilly, J., Green, D., Lind, J., Wallis, M., Miller, P., and Fitzgerald, G. (2012). Predicting emergency department admissions. *Emergency Medicine Journal*, 29(5):358–365.
- Champion, R., Kinsman, L. D., Lee, G. A., Masman, K. A., May, E. A., Mills, T. M., Taylor, M. D., Thomas, P. R., and Williams, R. J. (2007). Forecasting emergency department presentations. *Australian Health Review*, 31(1):83–90.
- Corke, C., de Leeuw, E., Lo, S. K., George, C., et al. (2009). Predicting future intensive care demand in australia. *Critical Care and Resuscitation*, 11(4):257.
- Cote, M. J. and Tucker, S. L. (2001). Four methodologies to improve healthcare demand forecasting. *Healthcare Financial Management*, 55(5):54–54.
- Earnest, A., Chen, M. I., Ng, D., and Sin, L. Y. (2005). Using autoregressive integrated moving average (arima) models to predict and monitor the number of beds occupied during a sars outbreak in a tertiary hospital in singapore. *BMC Health Services Research*, 5(1):1.
- Farmer, R. and Emami, J. (1990). Models for forecasting hospital bed requirements in the acute sector. *Journal of epidemiology and community health*, 44(4):307–312.
- Hans, E. W., Van Houdenhoven, M., and Hulshof, P. J. (2012). A framework for healthcare planning and control. In *Handbook of healthcare system scheduling*, pages 303–320. Springer.
- Hoot, N. R., LeBlanc, L. J., Jones, I., Levin, S. R., Zhou, C., Gadd, C. S., and Aronsky, D. (2008). Forecasting emergency department crowding: a discrete event simulation. *Annals of emergency medicine*, 52(2):116–125.
- Hyndman, R. J. and Athanasopoulos, G. (2014). *Forecasting: principles and practice*. OTexts.
- Jones, S. A., Joy, M. P., and Pearson, J. (2002). Forecasting demand of emergency care. *Health care management science*, 5(4):297–305.
- Jones, S. S., Thomas, A., Evans, R. S., Welch, S. J., Haug, P. J., and Snow, G. L. (2008). Forecasting daily patient volumes in the emergency department. *Academic Emergency Medicine*, 15(2):159–170.

- McCarthy, M. L., Zeger, S. L., Ding, R., Aronsky, D., Hoot, N. R., and Kelen, G. D. (2008). The challenge of predicting demand for emergency department services. *Academic Emergency Medicine*, 15(4):337–346.
- Mitchell, R. K., Agle, B. R., and Wood, D. J. (1997). Toward a theory of stakeholder identification and salience: Defining the principle of who and what really counts. *Academy of management review*, 22(4):853–886.
- Ozcan, Y. A. (2005). *Quantitative methods in health care management: techniques and applications*, volume 4. John Wiley & Sons.
- Reis, B. Y. and Mandl, K. D. (2003). Time series modeling for syndromic surveillance. *BMC Medical Informatics and Decision Making*, 3(1):1.
- Soyiri, I. N. and Reidpath, D. D. (2013). An overview of health forecasting. *Environmental health and preventive medicine*, 18(1):1–9.
- Wargon, M., Guidet, B., Hoang, T., and Hejblum, G. (2009). A systematic review of models for forecasting the number of emergency department visits. *Emergency Medicine Journal*, 26(6):395–399.

Appendix

Appendix A: Code

This code is used to calculate the number of beds occupied at every moment of the day from the 1st of January 2012 to 31st of December 2015. Two codes were developed; one for the Thorax unit and one for the General unit. The codes are exactly the same but only the data used changes. Therefore, different names are given to the different data. In this appendix only the code of the General unit is provided.

```
#Read file
p_data=read.table(file="Libro1.txt", header=TRUE, sep = "\t", dec=".", stringsAsFactors=FALSE)

#Create sequence of dates
Date=seq(as.Date("2011/1/1"),as.Date("2016/01/10"), by="day")

#Split file according to department
p_data_dep=split(p_data,p_data$Department)
p_data_gic=p_data_dep[[1]] #Won't be used in this script
p_data_tic=p_data_dep[[2]]

#To convert the times to indexes
timeToIndex=function(time){
  hour=as.numeric(gsub("(:[0-9][0-9])|(:[0-9])","",time))
  minute=round(as.numeric(gsub("([0-9][0-9]:)|([0-9]:)","",time))/6)
  return(10*hour+minute+1)
}

#To convert the dates to indexes (first index is 1/1/2011)
dateToIndex=function(date){

  index1=as.Date("2011-01-01")
  diff=as.numeric(as.Date(date)-index1)
  return(diff+1)
}

#Create the Vector of indexes for each patient (VI)
VI=function(patient){
  ADI=dateToIndex(as.Date(p_data_gic$Admission.date[patient])) #Admission Date Index
  ATI=timeToIndex(p_data_gic$Admission.time[patient]) #Admission Time Index
  DDI=dateToIndex(as.Date(p_data_gic$Discharge.date[patient])) #Discharge Date Index
  DTI=timeToIndex(p_data_gic$Discharge.time[patient]) #Discharge Time Index

  vi=c(ADI,ATI,DDI,DTI) #vector with the indexes (Adm.date, adm.time, disch.date, disch.time)
  return(vi)
}
```

```

#Create the matrix where the data will be sorted
B=matrix(0, nrow=length(seq(0,24,0.1)), ncol=length(Date))

#Make a for loop for all the patients
for(k in 1:length(p_data_gic$Patientnumber..anonymous.)){

  #Get the vector of patient k
  vik=VI(k)

  #If the patient k is admitted and discharged at the same day
  if(vik[3]==vik[1]){

    for (i in vik[1]:vik[3]){ #for for the columns
      for (j in vik[2]:vik[4]){#for for the rows

        B[j,i]=B[j,i]+1 #Add 1 to each specified cell
      }
    }
  }

  #If the patient k is discharged the they after being admitted
  if (vik[3]==vik[1]+1){

    VInew=c(vik[1],vik[2],vik[1],241)

    for(i in VInew[1]:VInew[3]){
      for (j in VInew[2]:VInew[4]){
        B[j,i]=B[j,i]+1
      }
    }

    VInew2=c(vik[3],1,vik[3], vik[4])

    for (i in VInew2[1]:VInew2[3]){
      for (j in VInew2[2]:VInew2[4]){

        B[j,i]=B[j,i]+1
      }
    }
  }

  #If the patient k is discharged the they after being admitted
  if (vik[3]==vik[1]+1){

    VInew=c(vik[1],vik[2],vik[1],241)

    for(i in VInew[1]:VInew[3]){
      for (j in VInew[2]:VInew[4]){
        B[j,i]=B[j,i]+1
      }
    }

    VInew2=c(vik[3],1,vik[3], vik[4])

    for (i in VInew2[1]:VInew2[3]){
      for (j in VInew2[2]:VInew2[4]){

        B[j,i]=B[j,i]+1
      }
    }
  }

  #If the patient is discharged after more than one day of being admitted
  if (vik[3]>vik[1]+1){

    VInew=c(vik[1],vik[2],vik[1],241)
  }
}

```

```
for(i in VInew[1]:VInew[3]){
  for (j in VInew[2]:VInew[4]){
    B[j,i]=B[j,i]+1
  }
}

for (i in (vik[1]+1):(vik[3]-1)){
  for (j in 1:241){
    B[j,i]=B[j,i]+1
  }
}

VInew2=c(vik[3],1,vik[3], vik[4])

for (i in VInew2[1]:VInew2[3]){
  for (j in VInew2[2]:VInew2[4]){

    B[j,i]=B[j,i]+1
  }
}
}

#Store the matrix created
write.table(B, file="All_beds_GIC.csv", row.names=FALSE, sep=" ")
```

Appendix B: Bed occupancy rates per shift

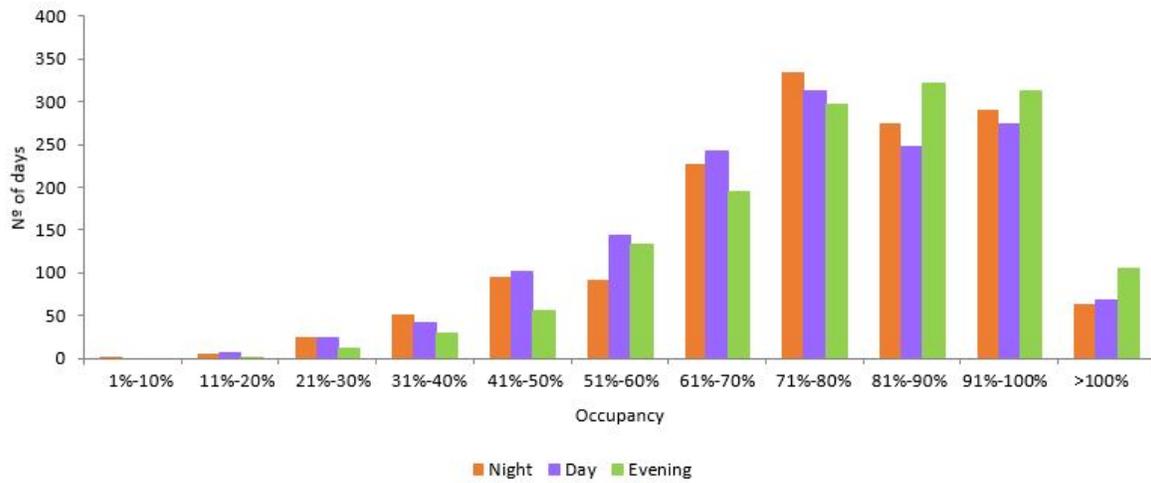


Figure 9.1: Thorax unit occupancy rates depending on the shift (sum of all the years)

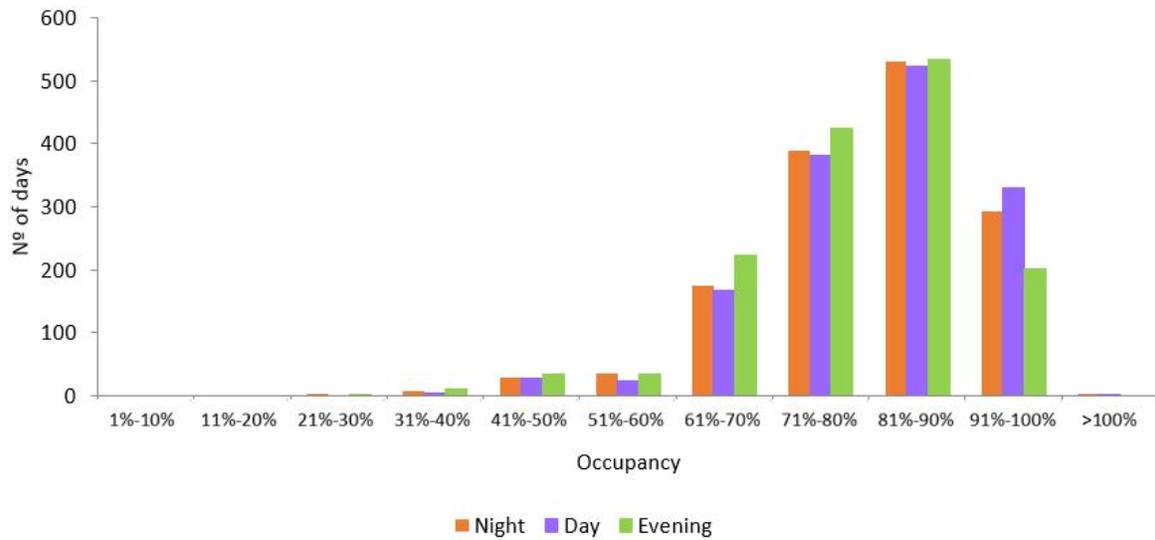


Figure 9.2: General unit occupancy rates depending on the shift (sum of all the years)

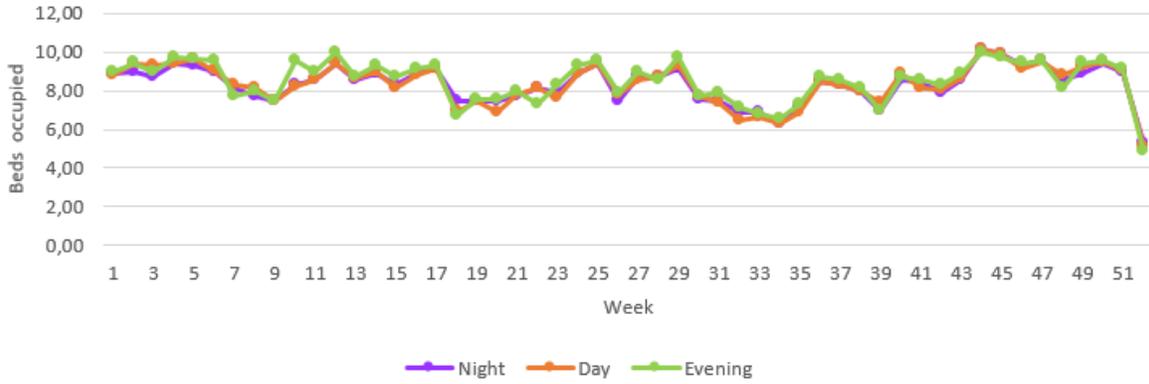


Figure 9.3: Number of beds occupied in the Thorax unit in 2012 depending on the shift

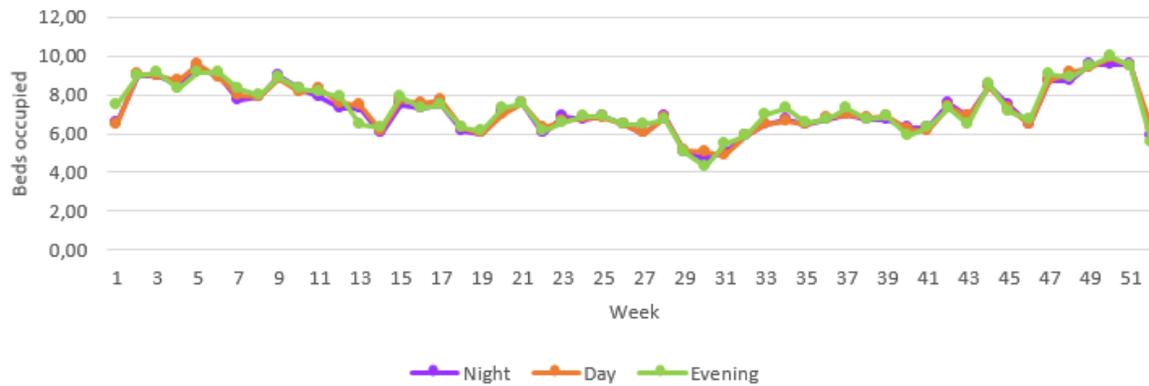


Figure 9.4: Number of beds occupied in the Thorax unit in 2013 depending on the shift

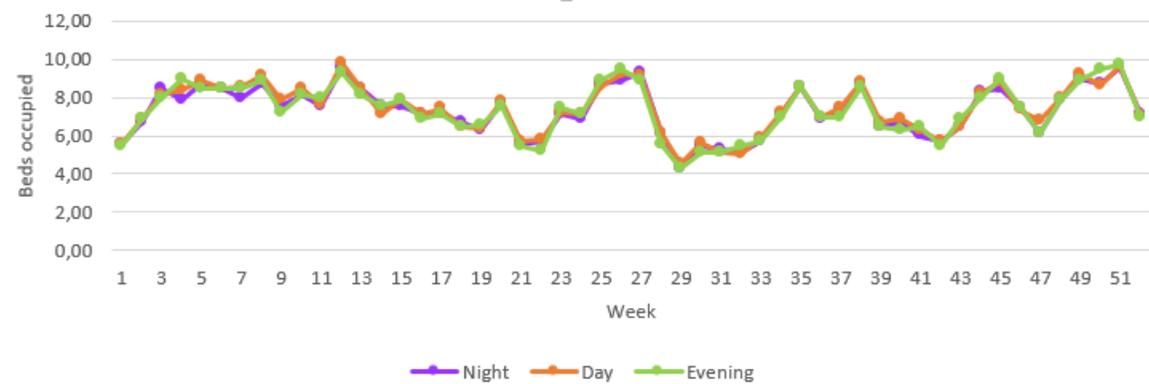


Figure 9.5: Number of beds occupied in the Thorax unit in 2014 depending on the shift

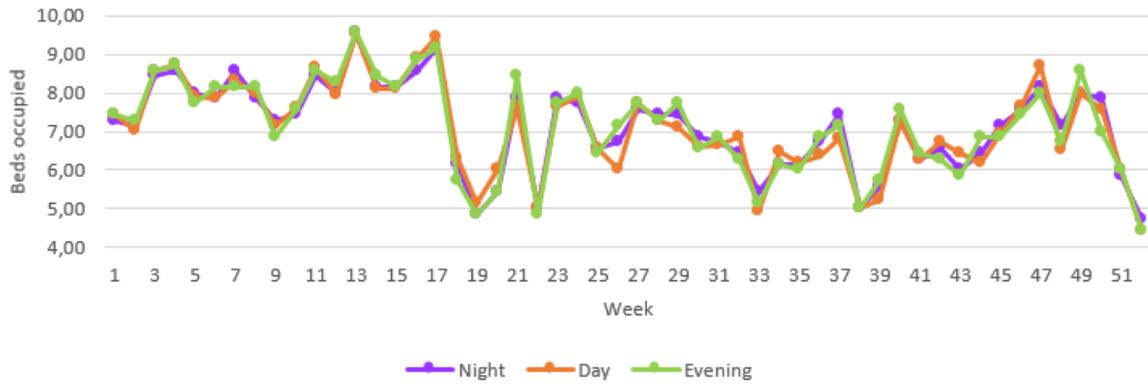


Figure 9.6: Number of beds occupied in the Thorax unit in 2015 depending on the shift

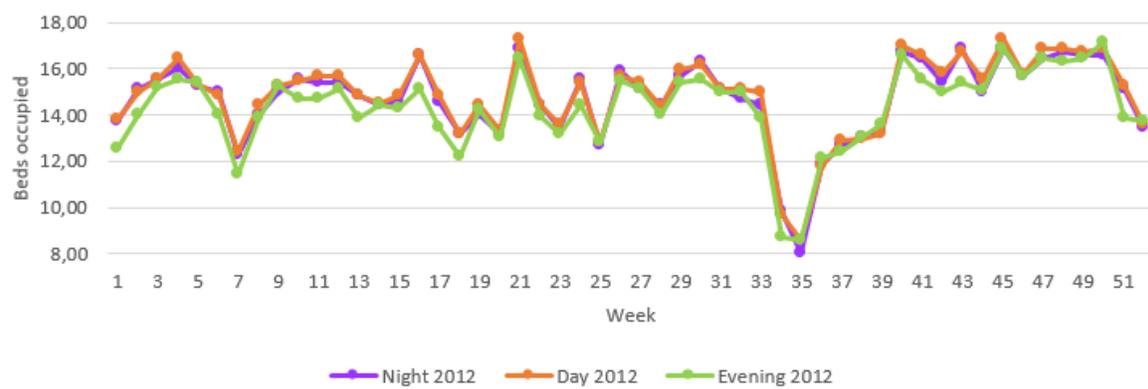


Figure 9.7: Number of beds occupied in the General unit in 2012 depending on the shift

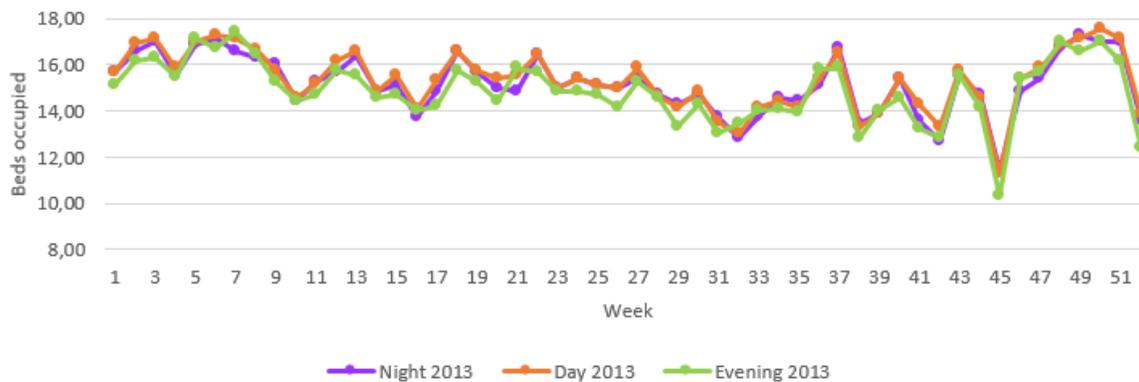


Figure 9.8: Number of beds occupied in the General unit in 2013 depending on the shift

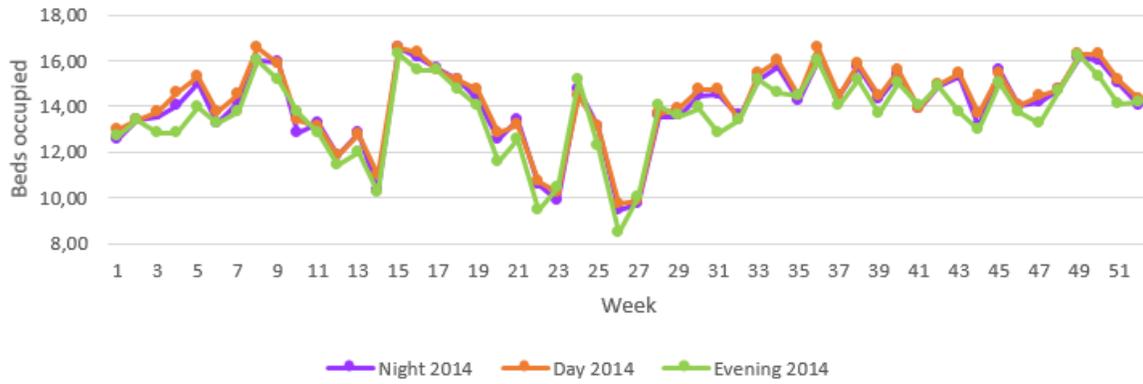


Figure 9.9: Number of beds occupied in the General unit in 2014 depending on the shift

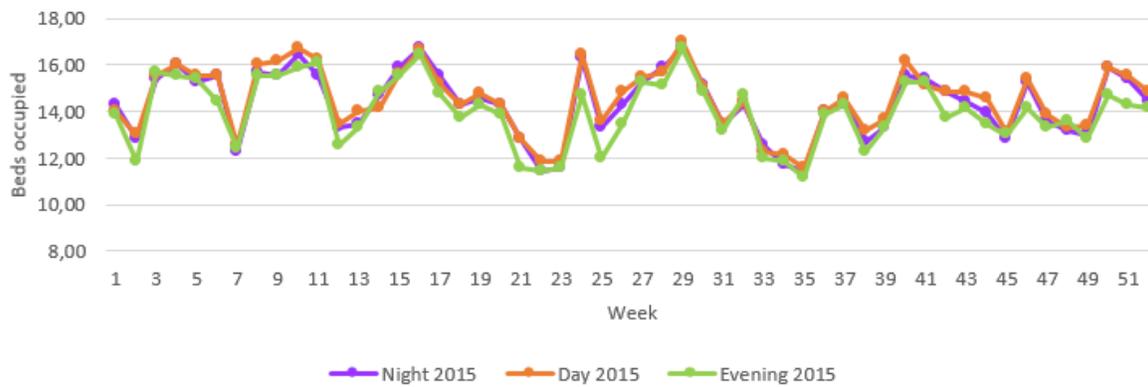


Figure 9.10: Number of beds occupied in the General unit in 2015 depending on the shift

Appendix C: General unit demand forecasting

In this section we explain steps 2 to 6 of the demand forecasting in the General unit. After analyzing the behavior of the data and decomposing it into a seasonal, trend and remainder components, we continue with the following steps.

Step 2: Data transformation

In this case, we do not need to transform the data from the General unit because the variation is not proportional to the level of the time series. Hence, we will leave it as it is.

Step 3: Is data stationary?

As for the Thorax unit, we will check the stationary state of the General unit data using the ACF plots and union root tests.

As we can see in Figure 9.11 the correlations of the ACF plot (left) drop to zero quickly. In this case, the first lag is positive but not very large. Hence, this suggests visually that the time series is quite stationary although maybe it could be more if we differentiate it. To evaluate it, we use union root tests to determine whether differencing is needed. The union root tests suggest that, in this case, data is stationary, so there is no need for neither seasonal nor first differences.

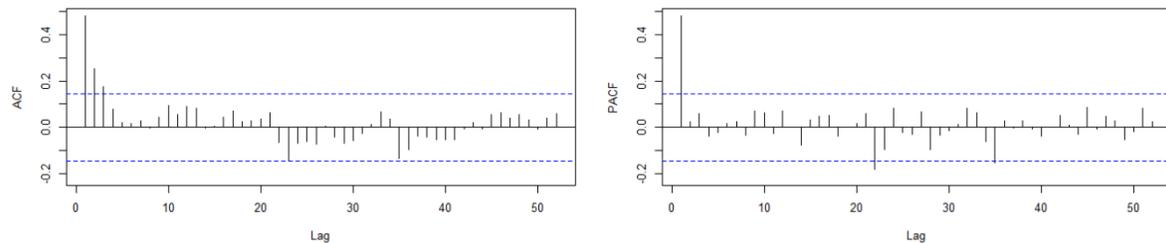


Figure 9.11: Correlation plot (ACF) (left) and partial correlation plot (PACF) (right) of the weekly data in the General ICU

Step 4: Examine ACF and PACF plots

By observing the ACF plot (left) in Figure 9.11, we can see how only lags 1, 2 and 3 are correlated and all the others are within the threshold limits. Even though the first three lags are correlated, the correlations are weak. Lag 1 shows the greatest correlation being 0.5. However, lags 2 and 3 can be correlated because of previous correlations. We confirm this by looking at the PACF plot on the right of Figure 9.11, where we see that the only significant value beyond the threshold is lag 1 (we ignore the other two slightly significant spikes as 1 out of 20 spikes can be due to chance). Hence, as in the Thorax unit, the best possible forecasting is for next week and, again, a long term prediction is possible but the forecasting error will increase dramatically every time because the correlation of lag 1 is only 0.5 and this will result in an inaccurate forecast.

We also use the ACF and PACF plot to determine the parameters of the ARIMA model. As we can see in Figure 9.11, the ACF is exponentially decaying and there is a significant spike

at lag p ($p = 1$) in the PACF, but none beyond lag p (we can ignore the other far two ones as mentioned before). Furthermore, we already know the value of d because in Step 3 we concluded that no differencing was needed. Hence, with $d = 0$, we visually determine that the appropriate model is ARIMA(1,0,0). As for the Thorax unit, we are not very convinced about the seasonality and this is why we have looked for a non-seasonal model. Nevertheless, we also want to check what happens with a seasonal model. Determining its parameters visually is too difficult, therefore, we use the automated algorithm (turning the short-cuts off). We also use the automated algorithm to see which kind of non-seasonal model it chooses. The automated algorithm suggests an ARIMA(1,0,0) with non-zero mean as the non-seasonal model and ARIMA(1,0,0)(0,0,0) with non-zero mean as the seasonal model. These will be our current models.

Step 5: Try the model

In this step we vary p , q , P and Q from the current model by ± 1 . Even though the automated algorithm suggests to use models with a constant (with non-zero mean), we also want to see how the models perform without this constant. Therefore we assess the performance of models with and without constant. From now on we refer to the models with constant as ARIMA(p,d,q)^c and ARIMA(p,d,q)(P,D,Q)_m^c. We use AIC_c and RMSE to compare the models and choose the best one. To compute the RMSE we fit the model with data from 2012 to the second quarter of 2015 and we test it with the remaining data from 2015.

Table 9.1 shows the tested non-seasonal models and their corresponding AIC_c and RMSE values. As seen, ARIMA(1,0,0)^c has the smallest AIC_c , one of the smallest RMSE and is a simple model. Therefore, we choose this one.

Non-seasonal			
<i>Model</i>	<i>AIC_c</i>	<i>RMSE</i>	<i>p-value residuals</i>
ARIMA (1,0,0) ^c	652.80	1.3999	0.9365
ARIMA (1,0,1) ^c	654.76	1.4020	0.9069
ARIMA (0,0,1) ^c	661.22	1.3896	0.1594
ARIMA (2,0,0) ^c	654.79	1.4016	0.9030
ARIMA (2,0,1) ^c	656.77	1.4023	0.8594
ARIMA (1,0,2) ^c	656.53	1.4006	0.8908
ARIMA (0,0,2) ^c	658.51	1.3955	0.4747
ARIMA (1,0,0)	709.32	1.8245	0.0109
ARIMA (1,0,1)	684.84	1.3986	0.0204
ARIMA (0,0,1)	1276.4	13.949	<0.05
ARIMA (2,0,0)	697.38	1.5952	0.0387
ARIMA (2,0,1)	664.33	1.3640	0.8247
ARIMA (1,0,2)	669.89	1.3173	0.3458
ARIMA (0,0,2)	1135.1	13.671	<0.05

Table 9.1: Non-seasonal ARIMA models for the General unit. The smaller the AIC_c and the $RMSE$, the better. A high p-value is desired because it indicates that the residuals behave like white noise.

Table 9.2 shows the tested seasonal models. Apart from these ones, we also tried ARIMA(2,0,0)(1,0,1)₅₂^c

but it gives an error, so we reject it. As we can observe, ARIMA(1,0,0)(1,0,1)₅₂^c has the smallest AIC_c and the second lowest RMSE. However, ARIMA(1,0,0)(1,0,0)₅₂^c has almost the same AIC_c and a bit higher RMSE, but it is simpler, so we decide to choose this last model.

Seasonal			
<i>Model</i>	<i>AIC_c</i>	<i>RMSE</i>	<i>p-value residuals</i>
ARIMA (1,0,0)(1,0,0) ₅₂ ^c	653.75	1.4256	0.5862
ARIMA (1,0,1)(0,0,1) ₅₂ ^c	653.82	1.4250	0.5846
ARIMA (1,0,0)(1,0,1) ₅₂ ^c	653.74	1.3294	0.5331
ARIMA (2,0,0)(1,0,0) ₅₂ ^c	655.68	1.4287	0.5000
ARIMA (2,0,0)(1,0,1) ₅₂ ^c	657.86	1.4259	0.4552
ARIMA (2,0,0)(0,0,1) ₅₂ ^c	655.74	1.4283	0.4975
ARIMA (1,0,0)(1,0,0) ₅₂	709.94	1.6508	0.0148
ARIMA (1,0,0)(0,0,1) ₅₂	710.13	1.6601	0.5846
ARIMA (1,0,0)(1,0,1) ₅₂	707.32	1.2940	0.0141
ARIMA (2,0,0)(1,0,0) ₅₂	697.41	1.4635	0.0174
ARIMA (2,0,0)(0,0,1) ₅₂	697.51	1.4669	0.0168

Table 9.2: Seasonal ARIMA models for the General unit. The smaller the AIC_c and the RMSE, the better. A high p-value is desired because it indicates that the residuals behave like white noise.

Therefore, the chosen models to continue with the further steps are: ARIMA (1,0,0)^c and ARIMA(1,0,0)(1,0,0)₅₂^c. If we calculate the coefficients with the training data, the models are mathematically written as Expression 9.1 for ARIMA(1,0,0) and as Expression 9.2 for ARIMA(1,0,0)(1,0,0)₅₂^c. The standard deviation of the white noise in Expression 9.1 is $sd = 1.4366$ whereas for Expression 9.2 it is $sd = 1.4346$.

$$y'_t = 14.6101 + 0.4816y'_{t-1} + \varepsilon_t \quad (9.1)$$

$$y'_t = 14.6115 + 0.4828y'_{t-1} + 0.912y'_{t-52} + \varepsilon_t \quad (9.2)$$

Step 6: Check the residuals

The Portmanteau test settings for ARIMA(1,0,0)^c are $h = 10$ and $DOF = 1$. The p-value returned is 0.9365 as shown in Table 9.1. Since it is higher than 0.05, this suggests that the residuals behave like white noise and, thus, it seems an appropriate model. For ARIMA(1,0,0)(1,0,0)₅₂^c the settings are $h = 36.4$ and $DOF = 2$. The resulting p-value, shown in Table 9.2, is 0.5862 which means that the residuals appear to be white noise.

Since both models have passed the tests of the residuals, we do not have to check for other models and we can proceed with the forecasting.

Appendix D: Calculation FTE for the implementation

The calculations done in this figure follows the calculation methodology explained in Chapter 6.

General Unit													
Month	Day	CAPACITY	BEDS 85% CAPACITY			NURSES NEEDED 85% CAPACITY			NURSE HOURS 85% CAPACITY				
			Night	Day	Evening	Night	Day	Evening	Night	Day	Evening		
8	Monday	17	14.45	14.45	14.45	7	10	8	56	80	64		
8	Tuesday	17	14.45	14.45	14.45	7	10	8	56	80	64		
8	Wednesday	17	14.45	14.45	14.45	7	10	8	56	80	64		
8	Thursday	17	14.45	14.45	14.45	7	10	8	56	80	64		
8	Friday	17	14.45	14.45	14.45	7	10	8	56	80	64		
8	Saturday	17	14.45	14.45	14.45	7	10	8	56	80	64		
8	Sunday	17	14.45	14.45	14.45	7	10	8	56	80	64		
									Total nurse hours		1400		
									Weekly hours FTE		40		
									FTE needed per week		35		

Thorax Unit													
Month	Day	CAPACITY			BEDS 85% CAPACITY			NURSES NEEDED 85% CAPACITY			NURSE HOURS 85% CAPACITY		
		Night	Day	Evening	Night	Day	Evening	Night	Day	Evening	Night	Day	Evening
8	Monday	8	8	8	6.8	6.8	6.8	3	5	4	24	40	32
8	Tuesday	8	8	8	6.8	6.8	6.8	3	5	4	24	40	32
8	Wednesday	8	8	8	6.8	6.8	6.8	3	5	4	24	40	32
8	Thursday	8	8	8	6.8	6.8	6.8	3	5	4	24	40	32
8	Friday	8	8	8	6.8	6.8	6.8	3	5	4	24	40	32
8	Saturday	8	8	8	6.8	6.8	6.8	3	5	4	24	40	32
8	Sunday	8	8	8	6.8	6.8	6.8	3	5	4	24	40	32
											Total nurse hours		672
											Weekly hours FTE		40
											FTE needed per week		16.8

Figure 9.12: Number of FTE needed per week per unit in case of implementing the 85% capacity scenario in August