Norwegian University of
Science and Technology

Improving the Robustness of Nurse Schedules in a Real-Life Instance - a Quantitative Analysis Based on Simulation and Rescheduling Under Uncertainty

Isabel Nordli Løyning
Line Maria Haugen Melby

Industrial Economics and Technology Management
Submission date: June 2018
Supervisor: Henrik Andersson, IØT
Co-supervisor: Anders Gullhav, IØT
Kjartan Kastet Klyve, IØT

Norwegian University of Science and Technology
Department of Industrial Economics and Technology Management

Problem Description

The purpose of this thesis is to develop tools to analyze the robustness of nurse schedules which are subject to uncertainty, as well as to identify strategies to improve this robustness. This is to be done by developing an optimization-based model which creates schedules, taking important rules and preferences into account. Furthermore, a tool should be created which reestablishes a feasible schedule during the online operational phase, when the real demand and supply of nurses is revealed. The robustness after the uncertainty realization and rescheduling can then be assessed.

This thesis is to be written in cooperation with the Department of Neonatal Intensive Care at St. Olavs Hospital.

Preface

This master's thesis concludes our Master of Science at the Norwegian University of Science and Technology (NTNU), with specialization in Applied Economics and Optimization. The work has been conducted during the Spring semester of 2018, and is written in collaboration with the Department of Neonatal Intensive Care (DNIC) at St. Olavs Hospital - Trondheim University Hospital.

We would like to express our sincere gratitude to our supervisors Professor Henrik Andersson, Postdoctoral researcher Anders N. Gullhav and PhD Candidate Kjartan Kastet Klyve, all from the Department of Industrial Economics and Technology Management (IØT) at NTNU, for their invaluable guidance and thorough feedback during our work. We would also like to thank Irene Voldset, who is the scheduling manager at DNIC, for helping us to develop a realistic case study through in-depth discussions on how the scheduling and rescheduling processes at DNIC are performed. Furthermore, we want to thank Inge Andersen, HR-analytic at St. Olavs, for his efforts in providing us with the data sets studied in the thesis.

There is an ongoing collaboration between researchers at IØT and St. Olavs Hospital, organized through the Regional Center for Health Care Development. This thesis is a part of this research collaboration, where researchers investigate whether operations such as scheduling can be carried out more efficiently using methods of operations research (RSHU, 2018). The thesis builds upon Løyning and Melby (2017), which is the preliminary work of the thesis.

Isabel Nordli Løyning and Line Maria Haugen Melby

Trondheim, June 2018

Abstract

Nurse schedules are subject to uncertainties in both the demand and supply of nurses, which often change on a daily basis. The demand depends on the number and severity of the patients admitted, while the supply depends on the number of absent employees. Schedules are usually created manually, taking only expected values and insights from previous periods into account. This makes it difficult to create schedules that always meet demand. We define a schedule's *robustness* as its *stability*, or capability to absorb unexpected events, as well as its *flexibility*, or capability to be reestablished when demand exceeds supply. If a schedule is not robust, it might be necessary to make multiple daily changes to the schedule, which is costly, time-consuming for the managers and inconvenient to the employees. It is therefore desirable to identify means to make schedules more robust.

In this thesis, we study this robustness, using real-life data from the Department of Neonatal Intensive Care (DNIC) at St. Olavs Hospital. We define a baseline MIP scheduling model used to make nurse schedules subject to the rules, regulations and preferences at DNIC. We also propose multiple proactive scheduling strategies intended to increase the schedule robustness, which are added to the baseline model. To imitate the uncertainties, we define models simulating the demand and supply of nurses, based on probability distributions calculated using historical data provided by DNIC. Both the simulation results and an initial schedule are used as inputs to a rolling horizon rescheduling model, which uses mixed integer programming to solve the daily rescheduling problem using the same rescheduling actions as are used in practice at DNIC. Finally, we evaluate the robustness of multiple schedules using various stability and flexibility measures.

Using the proactive scheduling strategies proposed, we significantly outperform the baseline model in terms of robustness in three out of four cases. We also combine the best

strategies into one model, successfully boosting the performance further. The best performing strategies include 1) assign a reasonable buffer of employees in excess of minimum demand on all shifts; 2) assign excess capacity to the shifts where employees who are extra vulnerable to absence are scheduled, and 3) allow employees to work extra weekends in exchange for extra off days. Another key insight obtained is the optimal duration of the replanning period when rescheduling. Using the robustness measures, we show that the longer the duration of the replanning period is, the more efficiently the schedule is reestablished.

Sammendrag

Turnusplanene til sykepleiere er utsatt for usikkerhet i både etterspørselen og forsyningen av sykepleiere, der begge aspektene ofte varierer på en daglig basis. Etterspørselen avhenger av antall pasienter samt deres alvorlighetsgrad, mens forsyningen er avhengig av hvor stort ansattfraværet er. Likevel utføres turnusplanlegging vanligvis manuelt, slik at usikkerheten bare tas hensyn til gjennom forventningsverdier og innsikt fra tidligere perioder. Dette gjør det vanskelig å lage turnusplaner som alltid møter etterspørselen. Vi definerer *robustheten* til en turnus som dens *stabilitet*, eller evne til å absorbere uventede hendelser, og dens *fleksibilitet*, eller evne til å bli gjenopprettet når den virkelige etterspørselen er høyere enn forsyningen. Hvis en turnusplan ikke er robust, kan det være nødvendig å gjøre flere daglige endringer på planen, noe som er kostbart, tidkrevende for de turnusansvarlige og slitsomt for de ansatte. Derfor er det ønskelig å identifisere måter å gjøre turnusplaner mer robuste på.

I denne masteroppgaven studerer vi denne robustheten, ved hjelp av virkelig data fra seksjon Nyfødt Intensiv (NI) ved St. Olavs hospital. I oppgaven definerer vi en grunnleggende MIP turnusplanleggingsmodell som generer turnusplaner underlagt de samme reglene og preferansene som turnusene ved NI. Vi foreslår også flere proaktive planleggingsstrategier som implementeres i modellen. For å imitere usikkerheten, simulerer vi etterspørselen og tilførselen av sykepleiere ved hjelp av modeller som benytter sannsynlighetsfordelinger basert på historisk data fra NI. Deretter brukes disse simuleringsresultatene samt en av de genererte turnusplanene som input til en *rolling horizon* replanleggingsmodell. Denne benytter blandet heltallsprogrammering til å løse det daglige replanleggingsproblemet ved hjelp av de samme mulige tiltakene som NI bruker i virkeligheten. Til slutt evaluerer vi robustheten til alle turnusplaner generert ved bruk av ulike måltall.

Ved å bruke de proaktive strategiene klarer vi å forbedre turnusplanens robusthet i forhold til grunnmodellen i tre av fire tilfeller. De beste strategiene er 1) sett opp en buffer av ansatte i tillegg til minimumsbemanningen på alle skift; 2) sett opp ekstra kapasitet på skift hvor ansatte som er sårbare for fravær er satt opp, og 3) la ansatte jobbe ekstra helger i bytte mot ekstra fridager. Vi kartlegger også hva som er den optimale lengden på replanleggingsperioden i replanleggingsproblemet. Ved hjelp av måltallene for robusthet viser det seg at jo lengre replanleggingsperioden er, jo mer effektivt kan turnusplanen gjenopprettes.

Contents

Problem Description **i**

Preface **ii**

Abstract **iv**

Sammendrag **vii**

1 Introduction **1**

2 Background **3**

2.1 Terminology 3

2.2 St. Olavs Hospital 4

2.3 Department of Neonatal Intensive Care 5

2.3.1 Shifts 6

2.3.2 Employees 6

2.3.3 Demand 8

2.3.4 Scheduling 9

2.3.5 Rescheduling 13

3 Literature Review **17**

3.1 Positioning Within Relevant Literature 18

3.2 Robustness 20

3.2.1 Robustness Terminology 20

3.2.2 A Methodology to Studying Robustness 21

3.3 The Nurse Scheduling Problem 22

3.3.1 Scheduling Approaches . . . 22
3.3.2 Key Aspects Within Nurse Scheduling . . . 23
3.3.3 Solution Methods for the Nurse Scheduling Problem . . . 26
3.4 Uncertainty . . . 27
3.4.1 Uncertainty in Demand . . . 27
3.4.2 Uncertainty in Supply . . . 29
3.5 The Personnel Rescheduling Problem . . . 29
3.5.1 Key Aspects Within Personnel Rescheduling . . . 30
3.5.2 Rescheduling and Robustness . . . 33
3.5.3 Solution Methods for the Personnel Rescheduling Problem . . . 34
3.5.4 Comparison of Papers on Personnel Rescheduling . . . 36

4 Problem Description 39
4.1 Description of the Scheduling Problem . . . 40
4.2 Description of the Uncertainty Realization . . . 43
4.2.1 Uncertainty in Demand . . . 43
4.2.2 Uncertainty in Absence . . . 43
4.2.3 Uncertainty in Availability . . . 44
4.3 Description of the Rescheduling Problem . . . 44
4.4 Description of the Robustness Evaluation . . . 48
4.4.1 Robustness Measures . . . 48
4.4.2 Managerial Insights . . . 49

5 Scheduling Model 51
5.1 Assumptions and Simplifications . . . 51
5.2 Definitions . . . 52
5.2.1 Indices . . . 52
5.2.2 Sets . . . 53
5.2.3 Parameters . . . 54
5.2.4 Variables . . . 55
5.3 Objective Function . . . 56
5.4 Constraints . . . 56
5.4.1 Covering Demand . . . 56
5.4.2 Weekends . . . 57

5.4.3 Work Hours 58
5.4.4 Required Rest 58
5.4.5 Shift Patterns 58
5.4.6 Other Scheduling Requirements 60
5.4.7 Variable Declarations and Fixations 61
5.5 Scheduling Example 61
5.6 Proactive Model Extensions 63
5.6.1 Extension 1: Buffer 63
5.6.2 Extension 2: Ghost 65
5.6.3 Extension 3: Absence 67
5.6.4 Extension 4: Extra Weekends 69

6 Data Analysis and Uncertainty Modelling 71
6.1 General Assumptions and Simplifications 71
6.2 Demand Uncertainty Model 72
6.2.1 Available Patient Data 72
6.2.2 Assumptions and Simplifications of the Demand Model 73
6.2.3 Objectives of the Demand Model 73
6.2.4 General Demand Model 74
6.2.5 Demand Transition Probability Distributions 75
6.2.6 Demand Simulation Algorithm 76
6.3 Absence Uncertainty Model 77
6.3.1 Available Employee Data 77
6.3.2 Assumptions and Simplifications of the Absence Model 80
6.3.3 Objectives of the Absence Model 81
6.3.4 General Absence Model 82
6.3.5 Absence Transition Probability Distributions 83
6.3.6 Absence Simulation Algorithm 87
6.4 Availability of Non-Absent Employees 88
6.5 Uncertainty Example 89

7 Rescheduling Model 91
7.1 Assumptions and Simplifications 92
7.2 Definitions 93

7.2.1 Indices 93
7.2.2 Sets 94
7.2.3 Parameters 95
7.2.4 Variables 97
7.3 Objective Function 99
7.4 Constraints 100
7.4.1 Covering Real Demand 100
7.4.2 Technical Constraints for Actions 101
7.4.3 Double Shift 101
7.4.4 Swap and Exchange 102
7.4.5 Consecutive Work 103
7.4.6 Variable Declarations and Fixations 103
7.5 Rescheduling Example 104
7.6 Reactive Model Extensions 108
7.6.1 Extension 1: Stricter Rescheduling 108

8 Case Study **111**
8.1 Case Instances 112
8.2 Data in the Scheduling Instances 112
8.2.1 S0: Base Case 113
8.2.2 S1: Buffer 116
8.2.3 S2: Ghost 117
8.2.4 S3: Absence 117
8.2.5 S4: Extra Weekends 118
8.3 Setup of the Uncertainty Models 119
8.3.1 Technical Settings 119
8.3.2 Demand Simulation 120
8.3.3 Absence Simulation 122
8.3.4 Availability Simulation 124
8.4 Data in the Rescheduling Instances 126
8.5 Technical Analysis 131
8.5.1 Testing the Scheduling Instances 131
8.5.2 Testing the Rescheduling Instances 132

9 Computational Study **135**
9.1 Test Phases . . . 135
9.2 Quantitative Robustness Analysis . . . 137
9.2.1 Phase 1 . . . 137
9.2.2 Phase 2 . . . 139
9.2.3 Phase 3 . . . 143
9.2.4 Phase 4 . . . 146
9.2.5 Validity of Results . . . 148

10 Concluding Remarks **149**

11 Future Research **151**

Bibliography **153**

Appendices **159**

A Compressed Models **161**
A.1 Scheduling Model . . . 161
A.1.1 Definitions . . . 161
A.1.2 Objective Function . . . 164
A.1.3 Constraints . . . 164
A.1.4 Proactive Model Extensions . . . 166
A.2 Rescheduling Model . . . 170
A.2.1 Definitions . . . 170
A.2.2 Objective Function . . . 174
A.2.3 Constraints . . . 175
A.2.4 Extension 1: Stricter Rescheduling . . . 177

B Examples **179**
B.1 Scheduling Example . . . 179
B.2 Uncertainty Example . . . 180
B.3 Rescheduling Example . . . 181
B.4 Calculating Real Demand . . . 183

C Data Analysis and Simulation Results **185**

C.1 Absence Probability Distributions 185
C.2 Results from Hypothesis Testing 186

D Technical Analysis of Instances `s5a-c` **189**

List of Tables

2.1 The skills included in each skill category 7
2.2 Patients treated by each skill category 9
2.3 The current scheduling process at DNIC 12
2.4 Rescheduling actions and their corresponding costs 15

3.1 Key papers on rescheduling reviewed in this thesis 36
3.2 Abbreviations used in Table 3.3 36
3.3 Comparison of this thesis ([1]) and 12 key papers within personnel and nurse rescheduling 38

4.1 Robustness measures, divided into stability and flexibility 49
4.2 Managerial insights from proactive and reactive strategies 50

5.1 The assumptions of the scheduling model 52
5.2 Examples of desirable shift patterns containing t, where $t \in \mathcal{T}^{SUN}$ 59
5.3 Examples of undesirable shift patterns ending on day t, where $t \in \mathcal{T}$ 59
5.4 Examples of undesirable shift patterns ending on day t, where $t \in \mathcal{T}^{SUN}$. 60
5.5 Examples of illegal shift patterns ending on day t, where $t \in \mathcal{T}$ 60
5.6 An example schedule for 9 nurses over 9 days, where D, E, N, F and F1 denote Day, Evening, Night, Off and the mandatory weekly off day, respectively 62
5.7 Proactive scheduling strategies and related managerial insights 63

6.1 General assumptions of the uncertainty models 72
6.2 The assumptions of the demand uncertainty model 73
6.3 The objectives of the demand uncertainty model 73

6.4 Examples of absence codes 78
6.5 The assumptions of the absence uncertainty model 81
6.6 The objectives of the absence uncertainty model 81
6.7 Results of a two-sample t-test for all combinations of days of the week for the short-term absence rate, where crosses or check marks indicate that the null hypothesis is rejected or accepted, respectively 85
6.8 Comparison of no aggregation and high- and low-risk aggregation when the data set is divided into weekdays and weekends 86
6.9 An example schedule, where the demand on Wednesday and Thursday is higher than expected, Nurse 4 is absent on Wednesday, and Nurse 8 hands in a sickness note of 7 days 89

7.1 The assumptions of the rescheduling model 92
7.2 An example schedule, where the current day is denoted by index 0 and the replanning period consists of three days. Nurse 4 is absent today, Nurse 8 is long-term absent, and the real demand today and tomorrow is higher than expected 105
7.3 Example of one possible assignment of actions when the real demand on day 0 and 1 is two, Nurse 4 is absent today and Nurse 8 has a long-term absence. A change to an employee's initial shift is illustrated by a slash cancellation and highlighted by blue 107

8.1 Definitions of case instances and their connection to model extensions and managerial insights 112
8.2 Key settings in the scheduling problem 113
8.3 Minimum and maximum demand per shift and day of the week 114
8.4 Minimum demand per skill category per shift 114
8.5 Number of employees by skill category 115
8.6 Values of the weighing parameters in `s0` 115
8.7 Key settings and weighing parameters in `s1` 116
8.8 Key settings in `s2` 117
8.9 Key settings in `s3` 118
8.10 Key settings in `s4` 119
8.11 The probabilities of transitioning between patient levels 121

8.12 Probabilities for transitioning between the states non-absent (a_N), short-term absent (a_S), and long-term absent (a_L) 123
8.13 Probabilities for accepting an extra shift depending on how far in advance the notice is given 125
8.14 Key settings in the rescheduling problem 126
8.15 Length of replanning period and its impact on $\mathcal{T}^{SN}$ and $\mathcal{T}^{LN}$ 127
8.16 Minimum demand limits for online operational staffing 127
8.17 Minimum demand per skill category per shift in the rescheduling problem . 128
8.18 Estimated need for nurses per patient per level (Halsteinli, 2017) and the skills required to treat the patients at each level, where skill 2 is Intensive Care and 3 is Monitoring 128
8.19 Values of the penalties in the objective function of the rescheduling problem, where each weight consists of a cost term and an inconvenience term . 129
8.20 Key statistics, run time and optimality gap for instances `s0-s4` 132

9.1 The four test phases of the computational study, their respective goals as well as the test instances included in each phase 136
9.2 Flexibility of instances `s0_r0_1-7`, where the lowest score on each measure is highlighted by green and the highest by red. Each measure is the total number of occurrences of the corresponding action during the full planning period of 77 days, averaged over 200 simulations 137
9.3 Flexibility of instances `s0-s4_r0_7`, where the lowest score on each measure is highlighted by green and the highest by red. Each measure is the total number of occurrences of the corresponding action during the full planning period of 77 days, averaged over 200 simulations 141
9.4 Flexibility of instances `s4_r0_7` and `s5a-c_r0_7`, where the lowest score on each measure is highlighted by green and the highest by red. Each measure is the total number of occurrences of the corresponding action during the full planning period of 77 days, averaged over 200 simulations 145
9.5 Flexibility of instances `s0_r1_7` and `s5*_r1_7`. The _post measures are averaged over the full planning period and 200 simulations, while the remaining measures are the total number of occurrences of the corresponding action during the full planning period of 77 days, averaged over 200 simulations . 147

B.1 An example schedule for 9 nurses over 9 days, where D, E, N, F and F1 denote Day, Evening, Night, Off and the mandatory weekly off day, respectively . . . 179
B.2 An example schedule, where the demand on Wednesday and Thursday is higher than expected, Nurse 4 is absent on Wednesday, and Nurse 8 hands in a sickness note of 7 days . . . 180
B.3 An example schedule, where the current day is denoted by index 0 and the replanning period consists of three days. Nurse 4 is absent today, Nurse 8 is long-term absent, and the real demand today and tomorrow is higher than expected . . . 181
B.4 Example of one possible assignment of actions when the real demand on day 0 and 1 is two, Nurse 4 is absent today and Nurse 8 has a long-term absence. A change to an employee's initial shift is illustrated by a slash cancellation and highlighted by blue . . . 182
B.5 Estimated need for nurses per patient per level (Halsteinli, 2017), where skill 2 is Intensive Care and 3 is Monitoring . . . 183
B.6 Historical patient data from February 1st 2016 . . . 183

C.1 Probabilities for transitioning between the states non-absent (a_N), short-term absent (a_S), and long-term absent (a_L) . . . 185
C.2 Results from two-tailed, two-samples t-tests with a significance level equal to 5 % for short-term absences on various days of the week. H_0 is that the means of variable 1 and 2 are equal. . . . 186
C.3 Results from two-tailed, two-samples t-tests with a significance level equal to 5 % for short-term absences on various shift types. H_0 is that the means of variable 1 and 2 are equal. . . . 187
C.4 Results from two-tailed, two-samples t-tests with a significance level equal to 5 % for long-term absences and various days of the week. H_0 is that the means of variable 1 and 2 are equal. . . . 187
C.5 Results from two-tailed, two-samples t-tests with a significance level equal to 5 % for long-term absences on various shift types. H_0 is that the means of variable 1 and 2 are equal. . . . 188

D.1 Values of the weighing parameters in instances `s5a-c` . . . 189
D.2 Key variable values, run time and optimality gap for instances `s5a-c` . . . 190

List of Figures

2.1 The seven steps in the current scheduling process at DNIC, and the actors involved in each step (Løyning and Melby, 2017) 11

3.1 A framework for health care planning and control (Hans et al., 2012) . . . 18

4.1 The overall robustness problem faced at DNIC 40
4.2 Placement of the Night shift in the scheduling problem 41
4.3 Overview of the online operational rescheduling problem at DNIC, where the dashed boxes represent inputs to the daily problems 45
4.4 Placement of the Night shift in the rescheduling problem 46

6.1 The transitions between states 0-5 for each hospital bed, where the probabilities of going from one state to the next is indicated on each arc 74
6.2 Key numbers for the employee data set 80
6.3 Transitions between the states non-absent (a_N), short-term absent (a_S) and long-term absent (a_L), as well as their respective transition probabilities . . 83

8.1 Variance in total number of patients for various simulation runs 120
8.2 Comparison of the average number of patients for simulation results and historical values for 200 simulation runs 121
8.3 Two arbitrary simulations compared to the historical average need for care per patient during the period from January 1st 2016 through April 9th 2016122
8.4 Comparison of the rate of employee absences each day for simulation results and historical values 124
8.5 Probability distribution for the number of requests for extra shifts 125

8.6 Average number of variables and constraints after presolve per iteration over the entire planning period and 200 simulations 133
8.7 Average run time per iteration over the entire planning period and 200 simulations in addition to error bars displaying the spread one standard deviation from the mean 133

9.1 Frequency of understaffing for instances `s0-s4_r0_7` 139
9.2 Frequency of understaffing for instances `s4_r0_7` and `s5a-c_r0_7` 144

List of Algorithms

6.1 Algorithm for the demand uncertainty model 77
6.2 Algorithm for the absence uncertainty model 88

Chapter 1

Introduction

Many hospitals, including St. Olavs Hospital in Trondheim, Norway, employ thousands of people with large variety in skills. Many hospital departments are manned round-the-clock, requiring that the employees perform shift work. The personnel schedules are subject to many regulations and preferences originating from both governments, trade unions, managers and employees, which must all be complied with. The result is that making high-quality personnel schedules is a very time-consuming and complex task.

The environment at hospitals is subject to uncertainties in the supply and demand for employees. These uncertainties are difficult to plan for, both in terms of unexpected absence such as sickness leave and a fluctuating number of patients admitted. Consequently, a common problem in health care systems worldwide is shortage of nursing staff combined with the fluctuating nature of patient demand (Lim and Mobasher, 2011). A small staff size leaves the schedules vulnerable to unforeseen events, where the absence of a few employees or unexpectedly high demand may cause a chain reaction of events, where many employees have to work unplanned shifts. The alternative, being understaffed, may negatively affect the health and well-being of both the patients and the employees at work. As making multiple changes to the employees' schedules is both costly to the hospitals and straining for the employees, it is desirable to have some protection against unforeseen events in the hospitals' nurse schedules. This extra protection is what we denote as *robustness*.

The purpose of this thesis is to develop suitable tools to analyze the robustness of nurse schedules, as well as to identify characteristics that improve this robustness. The tools are applied to a real-life case in cooperation with the Department of Neonatal Intensive Care, hereafter also denoted *DNIC*, at St. Olavs Hospital. The purpose is achieved through a four-step process. First, we develop a MIP model that solves the scheduling problem at DNIC. Additionally, multiple model extensions intended to increase the schedule robustness are implemented. Second, we define two models used to simulate the real demand and supply for nurses, based on probability distributions calculated using historical data from DNIC. Third, we develop a MIP rescheduling model, which takes a schedule and an uncertainty realization as input and solves the daily problem of reestablishing the schedule when the real demand exceeds the real supply, using the same rescheduling actions as used in practice at St. Olavs. Fourth and finally, we use multiple measures to assess the robustness of the underlying schedules.

To the best of our knowledge, we are the first to make a rescheduling model that takes simulation results based on probability distributions calculated from real data as input. All models proposed are based on a real-life case, and we believe that our study is among the most realistic ones to date, helping to bridge the gap between theory and practice. Key managerial insights obtained in the thesis include suggestions to DNIC on how the department best can improve the robustness of their schedules, as well as recommendations for the optimal duration of the replanning period used in the rescheduling problem.

The thesis starts with a presentation of useful background information regarding St. Olavs Hospital and DNIC in Chapter 2. The problem is put in a theoretical context in Chapter 3, where we present relevant literature. Chapter 4 contains the problem description, divided into the scheduling problem (Section 4.1), the uncertainty realization (Section 4.2), the rescheduling problem (Section 4.3) and the robustness evaluation (Section 4.4). These are further modelled in the same order in Chapters 5, 6 and 7, respectively. The case description is provided in Chapter 8, and the robustness evaluation is conducted in the computational study in Chapter 9. Finally, concluding remarks and suggestions for future research are presented in Chapters 10 and 11, respectively.

Chapter 2

Background

In this chapter, we present background information relevant to the thesis. Section 2.1 contains terminology important to understand the information provided. We then present the general scheduling situation at St. Olavs Hospital in Section 2.2, before describing the scheduling and rescheduling process at the Department of Neonatal Intensive Care (DNIC) in Section 2.3. Parts of this chapter builds upon Løyning and Melby (2017).

2.1 Terminology

The terminology presented in this section is important to understand the characteristics of DNIC. The terms listed are highlighted in italic letters the first time they occur in the subsequent sections.

- *Assistant nurse.* In the context of this paper, an assistant nurse is a certified nursing assistant with additional education within child care. In Norwegian, this is called a *barnepleier*.
- *CPAP (Continuous Positive Airway Pressure).* Treatment that uses a machine to pump air under pressure into the airway of the lungs (MedlinePlus, 2017).
- *High Flow.* For infants, High Flow is the delivery of heated, humidified and blended air/oxygen delivered via a nasal cannula at different flow rates ≥ 2 L/min (Mikalsen et al., 2016).

- *NAST (Neonatal Ambulatory Medical Care, in Norwegian called Nyfødt intensiv Ambulerende Sykepleie Tjeneste).* Medical care provided to prematures who are well enough to stay at home with their family, but whom still need extra follow-up from nurses.
- *NIV NAVA (Non-Invasive Ventilation with Neurally Adjusted Ventilatory Assist).* A mode to provide effective, appropriate non-invasive support to newborns with respiratory insufficiency (Stein et al., 2016).
- *Outpatient clinic.* A department dedicated to the diagnosis and treatment of people who at the time of visiting do not require a bed or to be admitted for overnight care.
- *Planning period.* The time period for which a schedule is created. A nurse schedule containing every shift of all the employees for the full duration of the planning period is created ahead of each period.
- *Rescheduling.* During the execution of the schedule, rescheduling is the process of reestablishing a schedule where supply meets demand after unexpected events such as employee absence or increases in the number of patients have occurred.
- *Respirator.* A mechanical ventilator that supports and provides breathing to patients that do not have the ability to breathe on their own (National Heart Lung and Blood Institute, 2011).
- *Preference scheduling.* Scheduling procedure where employees make requests for their preferred schedule ahead of the schedule creation. These requests are taken into account to the highest degree possible to ensure a fair schedule.
- *Shift pattern.* The combination of a nurse's assigned shifts over a set of consecutive days.

2.2 St. Olavs Hospital

St. Olavs Hospital is situated in Trondheim. The hospital is among the largest in Norway, with roughly 10 500 employees. The total operating costs in 2017 were $9,87$ billion NOK, with wages accounting for 64 % of the costs (Helse Midt-Norge, 2018).

Many of the wards at St. Olavs deliver services based on immediate help, making it difficult to predict the actual demand for staff at any given day. Further, several wards are open at all times, making shift work a necessity. The scheduling process at the hospital is typically done manually. The managers at each ward are free to conduct the scheduling as they think is best, given the budget and staffing level decided upon by the upper administration. The scheduling process is very time consuming, involving both trade unions, employees and managers. Further, the employees prefer knowing when they will work and when they will have holidays well in advance, resulting in a preference for using long *planning periods* for the schedules at many of the wards.

The staffing level at each shift is based on estimations of what will be the actual demand. Because the number of patients and the severity of their conditions is fluctuating, it is very difficult to obtain a good estimate. Furthermore, unforeseen events such as sickness leave frequently occur, with an average sickness leave of 7.7 % at St. Olavs in 2017 (St. Olavs Hospital, 2018). The long planning horizon combined with the fluctuating patient demand and unforeseen events results in frequent situations where wards are short on staff, spurring a need to call in substitute employees at high costs. This raises the question of whether it is possible to make the schedules more robust than what they are today, to better be able to cover demand with the currently available resources.

2.3 Department of Neonatal Intensive Care

The Department of Neonatal Intensive Care is a department at St. Olavs Hospital, whose main function is to treat sick newborns and premature babies. The treatment offered is advanced and varied, covering everything from treating critically ill newborns in incubators to assisting newborns with nutrition.

The employees at DNIC are organized into three units, depending on their skill sets. The units each cover particular patient groups, and are similar in size regarding the number of employees. The department is organized this way to ensure that the employees get sufficient continuity in their work to maintain a high standard in their specific areas of expertise. The three units are not physically divided; employees from the all units work together in the same workspace.

Besides these units, the department has an *outpatient clinic* with follow-up of extremely premature children, children in risk of deviant development up until school-age, as well as newly referred infants. DNIC also offers nutritional assistance for patients who stay at home with their family and live less than a one hour drive from the hospital through *NAST*, as well as a breast milk bank.

2.3.1 Shifts

The department is open and staffed year-round, 24 hours a day, which implies that shift work is required. Employees are assigned a work shift or an off shift each day, where a work shift can be either a Day shift, an Evening shift or a Night shift. To ensure a good information flow about the current condition of the patients, there is always some overlap between one shift and the next.

Each planning period, some employees must take on the responsibility to work some special shift types, which are executed simultaneously as the Day shifts. The tasks include working at the outpatient clinic, NAST and the breast milk bank.

Weekend shifts, Night shifts and shifts falling on holidays are far more unpopular to work than others. Consequently, there are specific requirements regarding work during these shifts, to ensure that the employees perceive the schedules as fair.

Some patterns of consecutive shifts are unpopular among the employees, either out of health reasons or because they prefer not working these patterns. The assignment of such *shift patterns* is generally avoided. Other patterns are considered favorable, and are assigned whenever possible. For example, employees prefer to be assigned off days on Fridays or Mondays when they have a weekend off to get a longer period of spare time. The scheduling manager tries to take this into account when setting up a schedule. Finally, some shift patterns are illegal according to the governmental regulations and trade union agreements. These are never assigned in the scheduling process.

2.3.2 Employees

The employees are divided into the three units according to their skill set; the Monitoring Unit, the Intensive Care Unit and the Emergency Unit. Employees are always assigned shifts at the unit they work for, but may cover the shifts of employees with a lower skill if people are absent during the execution of the schedule. A minimum number of employees from each unit must be present at all times. Additionally, the employees' needs and wishes regarding when they prefer to work, and the number of hours stated in their work contract, is important to consider.

Skills

Employees working in the Monitoring Unit are responsible for treating the patients with the least severe conditions. All new hires start working at there, and remain for at least half a year regardless of their previous work experience. They also go through a training period lasting for eight weeks when they start working at the department. Some *assistant nurses* also work at the Monitoring Unit. They are not trained in all of the Monitoring skills, and in turn, the scheduling manager tries to not assign too many assistant nurses to the same shift.

Employees who rank up from the Monitoring Unit start working in the Intensive Care Unit. These employees have been trained in using *respirators*, and are prepared to treat patients with more serious conditions than those at the Monitoring Unit. After a few years at the Intensive Care Unit, employees may rank up to the Emergency Unit. Nurses in this unit treat the infants with the most severe conditions, and are the only employees who have the skills required to receive patients straight out of the Maternity Ward. Table 2.1 contains a summary of the skills in each skill category.

Table 2.1: The skills included in each skill category

	Maternity ward	**Respirator**	**CPAP NIV NAVA Medication**	**Nutrition High Flow**
Emergency skills	x	x	x	x
Intensive Care skills		x	x	x
Monitoring skills			x	x
Assistant skills				x

Personal Requests

The scheduling manager at DNIC considers the employees' personal shift requests when planning the schedule for the next period using *preference scheduling*. Each day the employees can either request a specific work shift, an off shift, or leave it blank. An employee's requested schedule has to fulfill the requirements from preferred practices at the department, and governmental regulations. However, they may request to work the shift patterns that are usually not assigned, as long as governmental work regulations are not contradicted.

Contracted Work

DNIC employs a mix of full time and part time workers. The number of weekly hours each employee should work is defined in his or her work contract. Every employee is scheduled to work accordingly, but small deviations of a few hours are allowed out of practical reasons. The weekly workload may change from week to week, and it is only the average weekly workload over the entire planning period which must be equal to the number of hours in the contract. Employees are never scheduled to work overtime, but may take on extra shifts during the execution of the schedule, which could result in overtime pay.

DNIC mostly offers part-time work contracts. This is due do a local rule meant to ensure a fair division of weekend shifts among the nurses, stating that every nurse must work every third weekend. This rule implies that DNIC needs enough employees to be able to divide them into three separate work groups, where each group is big enough to cover the minimum demand for employees during weekends. This number of employees would result in an expensive over-coverage of the weekday shifts if every employee had a full-time position.

2.3.3 Demand

Patient arrivals at DNIC are difficult to predict, and during the patients' hospital stay the severity of their conditions may change considerably. To monitor the patients admitted to DNIC, each patient is examined and categorized into one out of five levels corresponding to the severity of their condition daily (Halsteinli, 2017). The estimated number of nurses required to treat the patients at each level can then be used to estimate the demand. The more severe the condition of a patient is, the more care that patient needs, and the more specialized skills are required to treat that patient. Table 2.2 roughly indicates which levels are treated by which skill categories.

Table 2.2: Patients treated by each skill category

Level	1	2	3	4	5
Monitoring skills	x	x	x		
Intensive Care skills				x	x
Emergency skills					x

According to the scheduling manager, there are no common patterns in the duration of stay on different levels. For example, the length of stay for patients on severity levels requiring treatment by nurses with Intensive Care skills may vary from two weeks to two months. Further, it happens that patients who are seemingly on their way to get well have a relapse, and are reclassified back to a more severe level. This makes it challenging to predict the exact demand for nurses within each skill category for each shift.

There are large variations in the average need for care per patient at DNIC. In the period January to October, 2016, there were, on average, approximately 17 patients admitted at the ward each day. The number of patients varied between 10 to 24, while the average need for care varied between 0.6 and 1.0 nurses per patient (Halsteinli, 2017). Due to this patient mix, there is a demand for employees from each skill category at each shift.

In addition to treating patients, some routine tasks must be completed during the weekday Day shifts regardless of the number of patients admitted. In turn, the minimum requirements are higher during these shifts.

2.3.4 Scheduling

The scheduling manager at DNIC creates schedules twice a year, where each schedule covers roughly 26 weeks and takes approximately 12 weeks to create. The long planning horizon is chosen due to the same practical reasons as explained in Section 2.2.

Challenges

Numerous aspects complicate the scheduling process. The most important challenges are elaborated on below.

Uncertain Demand Due to the long planning horizon and fluctuating demand for employees, it is difficult to schedule the right number of employees for each shift. To overcome this problem, a minimum requirement regarding the total number of employees working each shift, as well as requirements regarding their skill-mix, has been set, based on historical data.

Absence During the execution of the schedule, it is far from certain that all employees scheduled will show up for work. Inspecting historical data of employee absence per shift, it is clear that for almost every shift, one or more employees are usually absent. On average, the amount of sickness leave within employees working in nursing services has been stable on about 10 % the past 10 years. In 2016, the sickness leave for nurses at DNIC was 11.1 % (Halsteinli, 2017). The unpredictable absence makes it challenging to evaluate whether a planned schedule is sufficiently robust before the schedule is put into practice.

Weekend Work The demand for nurses is quite stable throughout the week (Halsteinli, 2017), meaning that the demand during the weekend and weekday shifts is similar. While the demand is the same, most employees prefer working weekdays instead of weekends. As described in Section 2.3.2, the employees are divided into three groups working every third weekend, each group big enough to just cover the minimum demand. As it is difficult to assign weekend shifts such that sufficiently many employees of each skill are present at all shifts, weekend shifts are the biggest bottleneck in the planning process.

The Scheduling Process at DNIC

The current manual scheduling process at the department is organized into seven steps, as described in Table 2.3, and illustrated in Figure 2.1. The duration of each step are estimations, and it may vary for each planning process.

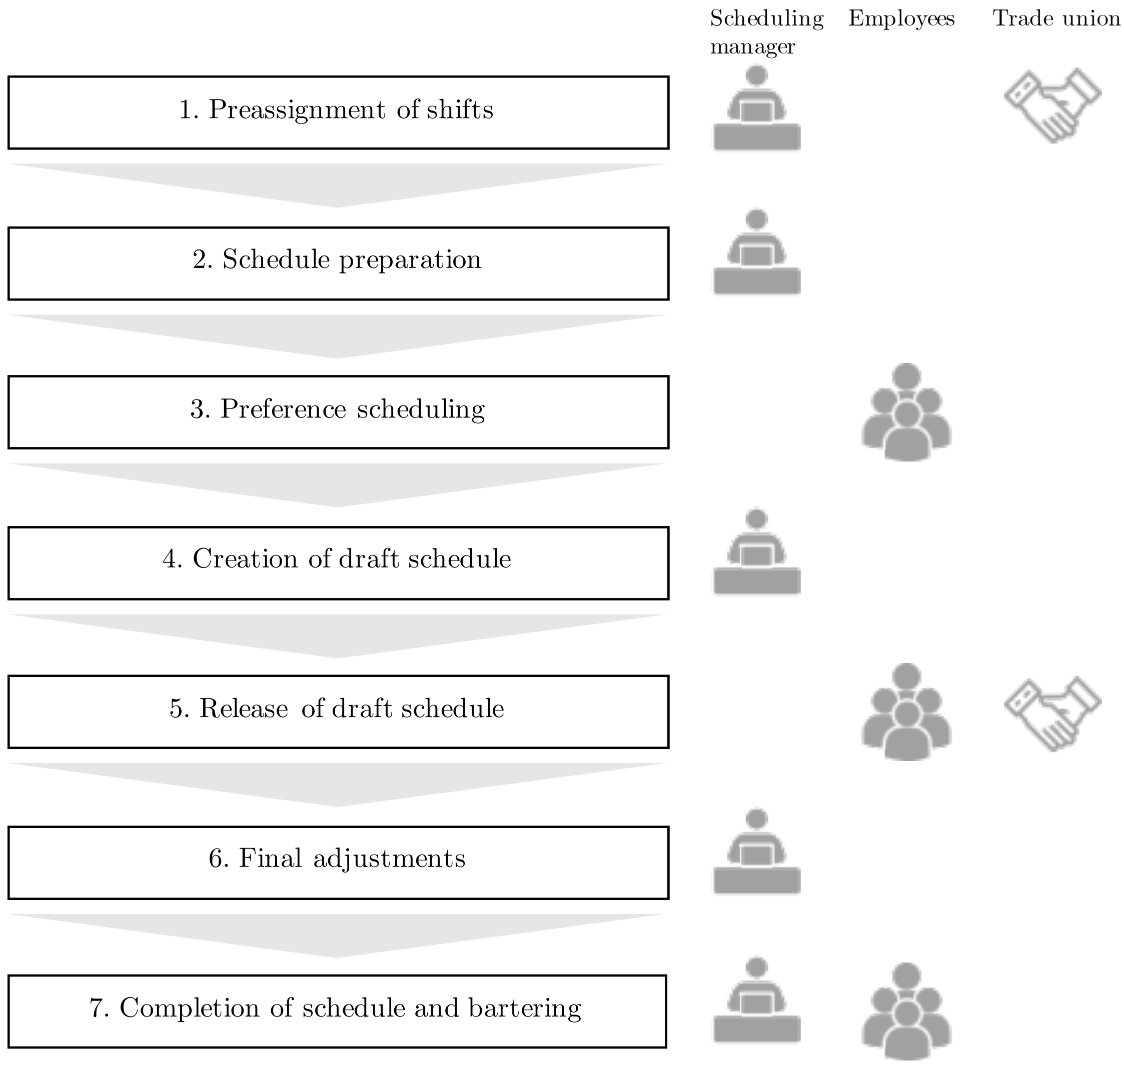

Figure 2.1: The seven steps in the current scheduling process at DNIC, and the actors involved in each step (Løyning and Melby, 2017)

Table 2.3: The current scheduling process at DNIC

Step	Name	Description	Duration
1	**Preassignment of shifts**	The scheduling manager divides the employees into three groups assigned to work the same weekends. The employees in each group must have a given set of skills to be able to meet demand on all shifts. In parallel, ground rules with the trade union are established. At the end of this step, a document containing these rules, the labor law regulations and an overview of the weekend groups is distributed to all employees.	2 weeks
2	**Schedule preparation**	The scheduling manager sets up the scheduling system and cross-checks that all work contracts registered are correct.	2 days
3	**Preference scheduling**	The employees enter their preferred schedule for the upcoming period into the scheduling system. The schedule must be in line with the ground rules and regulations from the scheduling manager.	2 weeks
4	**Creation of draft schedule**	The scheduling manager takes all employee requests into account and makes a draft schedule.	6 weeks
5	**Release of draft schedule**	The draft schedule is made available to both the trade union and the employees. Both parties have one week to make objections.	1 week
6	**Final adjustments**	The scheduling manager makes the final adjustments after inputs from the trade union and the employees.	1 week
7	**Completion of schedule and bartering**	The schedule is now complete. However, the employees are allowed to barter shifts as long as it is consistent with laws and regulations. These changes must be reported to the scheduling manager.	-

2.3.5 Rescheduling

The actual staffing and demand is gradually revealed when the schedule is put into use. The demand for employees usually changes daily, as patients are admitted, discharged or transferred from one level to another. However, the managers can normally predict approximately what the demand in the upcoming days will be. The estimate is based on the condition and treatment of the patients currently admitted, and on an ongoing communication with the Maternity Ward, where DNIC is notified in advance whenever the Maternity Ward believes that sick or premature newborns will arrive.

Employee absences are often difficult to predict, and can occur on short notice before the start of a shift. There is a difference in the duration and predictability of the absences occurring, meaning that different considerations and actions are made when each of them happen.

A short-term absence is an absence where an employee hands in a self-certification of absence, stays at home with a sick child, or in another way is unable to come to work in a way such that the duration of the absence is unpredictable, but typically shorter than a week. Common for these absences is that they are reported to DNIC the same day as they occur, often at the Day shift, and that it is uncertain when the absence will end.

A long-term absence means that an employee hands in a sick leave certificate provided by a doctor, enters maternity leave, or in another way enters a state of absence with a more predictable, but longer, duration. Although long time spans are the most common, by this definition the duration does not necessarily have to be long; sick leave certificates may also be provided for cases where the absence is predicted to only last a few days.

If the long-term absence only lasts for a few weeks, the scheduling manager typically publishes the shifts that have to be covered on an online platform, where the employees can request to work these shifts. The shifts are assigned to the employees a few days in advance. When an employee enters a state of absence with a longer duration, such as maternity leave, the manager often tries to create the schedule of employees returning from a long-term absence such that these employees may cover the shifts of the absent employee. Another option is to hire new employees, or give extra training to selected employees such that they can advance in skill level and cover the shifts of employees higher up in the skill hierarchy.

The Daily Rescheduling Process at DNIC

When an employee calls in absent or the demand increases, the scheduling manager first evaluates whether there is sufficient staff left to work the shift in terms of skills and overall staffing. If there is, no actions are taken. If there is not, several rescheduling actions are available. The possible actions and their costs are described in Table 2.4.

Table 2.4: Rescheduling actions and their corresponding costs

Action	Description	Cost
Extra shift	Schedule an employee with a day off to work an extra shift.	Contracted wage plus possible overtime costs.
Double shift	When an employee works two consecutive shifts without a break, or with a break in between which is too short to get sufficient rest according to relevant regulations; e.g. an employee who worked the Evening shift could stay to work the Night shift as well. Employees can also work partially double shifts; e.g. to stay for a few hours after a work shift is completed to help finish certain tasks before heading home.	Contracted wage, financial inconvenience compensation and possible overtime costs. The compensation depends on whether a break was provided between the shifts or not.
Swap	Ask an employee who has a work shift in the upcoming days to work another shift in which there is insufficient supply relative to demand instead. For example, assume employee A is assigned the Day shift tomorrow, and B is assigned the Evening shift today. If employee B calls in sick, employee A could be swapped to the Evening shift today in exchange for an off day tomorrow. Employees are most commonly swapped from the Day shift to Evening or Night, but the opposite also happens. This is because there usually are bigger buffers of surplus employees who can be swapped from Day shifts.	Financial inconvenience compensation which depends on how early the employee is notified; short notice is given after 12 AM the day before the swap and gives a higher compensation than if the employee is notified longer in advance. Wage cost is transferred to the new work shift.
Exchange	Exchange is a special case of swap, and can occur if an employee has specifically requested an off day on a day when she has a work shift scheduled. She can then take on a work shift another day in exchange for the requested off day, but without receiving any additional financial compensation - it is thus a form of exchanging shifts with the department.	No additional costs; wage cost is transferred to the new work shift.

Because special training is required to work at DNIC, only nurses who are employed by DNIC may work unplanned shifts. This is different from many other departments at St. Olavs Hospital, where nurses with more general competence may take on shifts at departments where they do not usually work. The decision of which employee is asked to cover the shift in question is based on a weighing of several criteria:

- *Costs.* The wage costs, the possible financial inconvenience compensations and the overtime costs of the unplanned shift depends on which employee is called in and the action taken, as indicated in Table 2.4. Swaps and double shifts incur a financial compensation regardless of how much the employee has worked that week. Extra and double shifts increase the number of hours worked that week, meaning that the employees could accumulate enough weekly work hours to receive overtime pay. Therefore, it is often cheaper to offer unplanned shifts to employees with smaller work contracts, as these are less likely to reach this limit.
- *Employee requests.* There is a list at DNIC where employees can signal that they are interested to take on specific extra shifts. The employees on this list are often the ones who are asked first when additional employees have to be called in.
- *Required skill.* The employee called in should be able to cover the tasks corresponding to the skills in deficit. New hires therefore seldom work extra shifts.
- *The employees' schedule around the shift.* Ideally, the employees should still be allowed sufficient rest before and after their other work shifts. As there is some inconvenience related to schedule changes, it is desirable that the employees do not have too many of them each week, and that a changed shift is not re-changed later.
- *Fairness.* The scheduling manager tries to distribute the additional shifts in a way the employees perceive as fair. However, the perception of fairness may vary from employee to employee. One employee may believe that offering equally many additional shifts to each employee is fair, while another finds it fair that the employees who most frequently request additional shifts are offered more shifts than those who seldom request it, or who often reject an offer.

Chapter 3

Literature Review

In this chapter, we present literature relevant to the problem of scheduling under uncertainty faced at The Department of Neonatal Intensive Care. The literature discussed provides a theoretical context for the processes of both creating baseline schedules and of reestablishing them when the schedules are executed and imbalances in demand and supply occur due to unexpected events. Some of the work in this chapter is based on the literature review by Løyning and Melby (2017). In the review, we assume that the reader is familiar with basic terminology within Operations Research.

The literature review on personnel scheduling by Van den Bergh et al. (2012) has been a great source for finding relevant literature on nurse scheduling, while the main papers on nurse rescheduling studied in this thesis were also reviewed by Mutingi and Mbohwa (2017). Furthermore, we conducted a literature search using relevant keywords on Google Scholar and Oria. A sample of the keywords searched for include *nurse/personnel*, *(re)scheduling/(re)rostering*, *robustness/robust* and *uncertainty*. While the nurse scheduling problem has been studied in depth, only a limited selection of literature on nurse rescheduling exist. The scope of our review therefore also includes personnel rescheduling.

The literature review starts by positioning the problem studied within the relevant literature in Section 3.1. We thereby define key aspects of robustness in Section 3.2. In Section 3.3 we present key aspects of the nurse scheduling problem, important to create the nurse scheduling model later in this report. Section 3.4 contains a description of how the uncertainties in demand and supply are treated in the related literature. Finally, we investigate the characteristics of the personnel rescheduling problem in Section 3.5.

3.1 Positioning Within Relevant Literature

There are numerous processes which health care organizations need to plan and control in their daily operations. Hans et al. (2012) present a framework that can be used to structure these processes in an organized manner, dividing the operations into four managerial areas and four hierarchical levels of control. The framework is displayed in Figure 3.1.

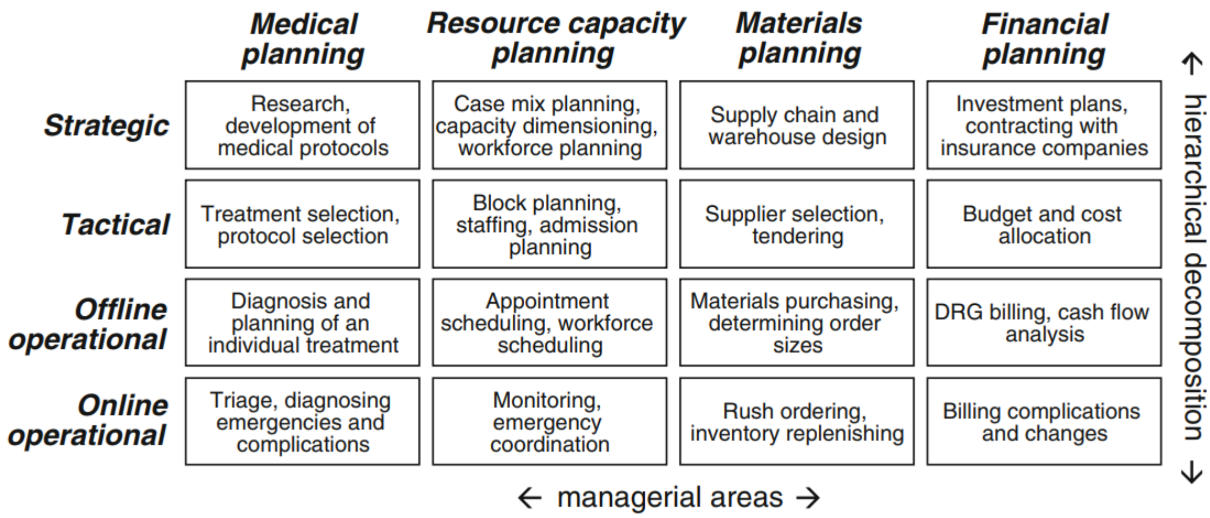

Figure 3.1: A framework for health care planning and control (Hans et al., 2012)

In the hierarchical decomposition, the strategic, tactical and offline operational level concern in advance decision making, while online operational planning concern reactive decision making. The decision horizon is longest on the strategic level. The information available on this stage is more uncertain and less available than on the other levels. As the time span between planning and execution shortens, more information becomes available. There is typically less flexibility, meaning that coordinators must adapt their short-term plans to the resources they have been given. Each managerial area has a diverse set of tasks to be performed on each level of planning, as indicated by the examples in Figure 3.1.

Personnel scheduling, or rostering, is one of the processes to be performed. This is the process of constructing work timetables of staff so that an organization can satisfy the demand for its goods or services (Ernst et al., 2004). The process can be divided into two parts; 1) determining the number of staff, and the necessary skill set, required to meet the service demand, and 2) assign individual staff members to shifts to meet the required staffing levels at the time of the shifts, and assign duties to the individuals for each shift. As resource capacity planning addresses the dimensioning, planning, scheduling, monitoring and control of renewable resources such as staff, personnel scheduling fits well within this managerial area.

Planning the size and schedules of the staff at a hospital spans multiple hierarchical levels of control. On the strategic level, which has a long-term planning horizon where decisions are based on aggregated information and forecasts, the planning is affected by decisions such as dimensioning the capacity of a new department to be able to service a given level of patients. Tactical planning has a shorter planning horizon. The employee staffing problem, which concern deciding the exact size of the staff, is solved on this level.

Operational planning involves short-term decision making, and is divided into in advance offline and reactive online planning. Staff scheduling, or to create schedules for a fixed work force, is typically performed on the offline level. This is commonly called the nurse scheduling or nurse rostering problem within health care applications. The uncertainties in demand and supply are realized on the online operational level. The action performed on this level is thus to react to unforeseen events such as sickness leave in order to reestablish a balance between the real demand at a department and the capacity of the workforce. In the related literature, the problem of reestablishing a feasible nurse schedule after unexpected events have occurred is known as the nurse rescheduling or nurse rerostering problem.

The focus of this thesis is personnel planning on the operational level, with a fixed work force consisting of nurses and assistant nurses. We therefore consider solving the staff dimensioning problem on the tactical level as outside the scope of this thesis. However, the size of the staff certainly affect our problem; in general, the higher the sum of contracted work hours at DNIC, the easier it is to cover demand, and the less vulnerable the department is to unforeseen events as more employees are scheduled on each shift. This is also supported by the related literature; both Harper et al. (2010) and Kokangul et al. (2017), who study the nurse staffing problem, stress the importance of first knowing the best size and skill-mix of nurses before the nurse scheduling problem is solved.

As we want to study how uncertainty affects the robustness of nurse schedules, we study both the offline and online operational planning level. On the offline operational level, the problem solved is to create a feasible schedule for a fixed workforce, which is best characterized as a nurse scheduling problem. The problem of reestablishing a feasible schedule as a reaction to the uncertainty realization faced, or the nurse rescheduling problem, is the problem solved on the online operational level.

3.2 Robustness

When schedules are subject to uncertainty, their robustness indicate how prepared the schedules are to cope with the possible outcomes of the uncertain elements. In this section, we further elaborate on the robustness term, as well as presenting a methodology where optimization and simulation is used to assess the robustness of schedules.

3.2.1 Robustness Terminology

Robust schedules are able to precede uncertainties and have a predefined solution for addressing those uncertainties (Lim and Mobasher, 2011). Having a robust schedule means, for example, that we avoid situations where just one person calling in sick causes a chain reaction of disruptions throughout the hospital because that person is the only scheduled person with a particular expertise. This is done by having two (or more) people with that particular expertise scheduled to be at work at the same time (Burke et al., 2004).

A *disruption* is an event leading to an unexpected imbalance between the supply and demand for personnel. The disruptions can either be absorbed by the personnel schedule, or require schedule changes (Ingels and Maenhout, 2015, 2017, 2018). Personnel schedules are robust if they are *stable* and *flexible* when disruptions occur in the online operational phase (Ionescu and Kliewer, 2011; Ingels and Maenhout, 2017, 2018). Stability is the degree to which schedules can absorb disruptions. Stable schedules thus require few adjustments when there are unexpected changes to the operating environment. A schedule's flexibility, or adjustment capability, is its capability to react efficiently to disruptions.

Thus, a schedule is flexible if there are sufficient options for schedule changes to efficiently reestablish a schedule where supply meets demand after an unexpected event.

Proactive and *reactive* scheduling strategies can be used to improve the robustness of nurse schedules. Proactive strategies are used during the offline operational scheduling phase, while reactive strategies concern making changes to the schedule on the online operational level (Maenhout and Vanhoucke, 2013b). A common proactive strategy, which also improves the stability, is to assign time or capacity buffers on the shifts of a schedule (Ingels and Maenhout, 2015, 2017). Examples of reactive strategies are to assign extra shift, overtime work, or to simply accept that demand cannot fully be met (Ingels and Maenhout, 2015).

Robust solutions should not be confused with *robust optimization*. Robust optimization was first studied by Soyster (1973), and is a well-defined procedure with clear rules about how to handle uncertainty. However, as we study a real-life case in this thesis and want managerial insights which can be used in the planning process at DNIC to be among the results of the thesis, we consider using e.g. proactive strategies as more relevant than applying robust optimization.

3.2.2 A Methodology to Studying Robustness

Ingels and Maenhout (2015, 2017, 2018) use a three-step methodology to study the flexibility and stability of personnel schedules. The methodology can be used to test the quality of different proactive and reactive scheduling strategies by applying it several times for multiple baseline rosters, each created using specific strategies. The three steps are as follows:

1. Construct a baseline personnel roster using an appropriate solution method. Possibly base the roster on specified proactive scheduling strategies.

2. Perform a discrete-event simulation of the daily variability in the supply and demand for staff, and perform necessary adjustments to the baseline roster based on the simulation results. Execute this step multiple times to obtain an accurate picture of the impact of the proactive and reactive strategies used.

3. Evaluate the adjusted roster by assessing the performance in the online operational

decision phase.

The methodology can thus be used to evaluate schedule robustness on a day-to-day basis, and we consider it an appropriate tool to imitate the real-life online operational rescheduling process, given realistic simulations of the variability in demand and supply for staff.

3.3 The Nurse Scheduling Problem

Several approaches to putting up nurse scheduling exist, where each method differs in the different stakeholders' level of involvement. No matter the approach, aspects such as satisfying demand and following relevant rules, regulations and preferences are important to take into account in order to make good schedules. In this section, we further elaborate on both the scheduling process and the most common underlying conditions as well as solution methods.

3.3.1 Scheduling Approaches

There are three main approaches to nurse scheduling; *centralized scheduling*, *unit scheduling* and *self-scheduling* (Burke et al., 2004). The main difference between them is the level of involvement of different stakeholders in the process.

In centralized scheduling, an administrative department carries out the whole scheduling process, while the responsibility is shifted towards the managers or head nurses of specific units with unit scheduling. Unit scheduling thus gives a higher local influence over the scheduling process than centralized scheduling, but is less cost efficient. Self-scheduling, on the other hand, empowers the employees, by letting them create their own schedules. A hybrid approach between unit and self-scheduling is *preference scheduling* (Bard and Purnomo, 2005b), where the employees first create their own schedules, before the head personnel oversees that all regulations are met and make the necessary changes. This approach gives a trade-off between fulfilling the individual requests of nurses, and accommodating to the scheduling preferences at the hospital. Recalling the scheduling process at DNIC described in Section 2.3, preference scheduling stands out as the most similar process.

3.3.2 Key Aspects Within Nurse Scheduling

Various aspects must be considered in the nurse scheduling process. First, the demand for employees on each shift should be satisfied. Further, the finished schedule should comply with rules and regulations regarding staff scheduling. Also, several quality measures can be considered to ensure that the nurses perceive the schedule as acceptable and fair.

Covering Demand

Hospitals are required to provide at least a minimum level of care on each shift, and often around the clock. The minimum demand can be given both in terms of total staff, and in terms of different skill categories. Hard constraints are commonly used to ensure that the minimum staffing requirement for each shift is met (Van den Bergh et al., 2012; Azaiez and Al Sharif, 2005; Rönnberg and Larsson, 2010; Bard and Purnomo, 2005b; Lim and Mobasher, 2011; Klyve and Beckmann, 2016). However, both Bard and Purnomo (2005b) and Rönnberg and Larsson (2010) acknowledge that it may be impossible to fulfill the staffing requirements at all times, allowing the use of substitute nurses at a penalty cost to ensure feasibility.

When there is demand for nurses of different skills, a common approach is to define a hierarchical set of skills, where employees belonging to higher categories are allowed to be assigned to shifts that lower skilled nurses are usually assigned to, but not vice versa (Klyve and Beckmann, 2016; Lim and Mobasher, 2011; Aickelin and Dowsland, 2004). This is similar to the skill system at DNIC. However, at DNIC, it is desirable to minimize the number of cases where employees must rank down, while Klyve and Beckmann (2016) and Lim and Mobasher (2011) are indifferent to whether nurses rank down on a shift or not. Skills are considered fixed throughout the planning period in all the papers mentioned, just as in the scheduling problem faced at DNIC.

Time-Related Considerations

A collection of constraints related to time is a common characteristic of the nurse scheduling problem. These constraints are regulated by both governmental regulations, such as the Norwegian Arbeidsmiljøloven §10 (Arbeids- og sosialdepartementet, 2017), and by preferred practices at the workplace. In the literature, they are found treated as both hard and soft. Common examples are to assign at most one work shift per nurse per day, and to put a maximum limit on the number of allowed consecutive work shifts (Rönnberg and Larsson, 2010; Azaiez and Al Sharif, 2005; Bard and Purnomo, 2005b; Klyve and Beckmann, 2016) and consecutive night shifts (Rönnberg and Larsson, 2010; Klyve and Beckmann, 2016).

Most employees prefer not working during weekends. Therefore, Bard and Purnomo (2005b), Azaiez and Al Sharif (2005), Rönnberg and Larsson (2010) and Klyve and Beckmann (2016) state that a nurse shall either have a *work weekend*, meaning that he or she should work both Saturday and Sunday, or have both days off. However, not all of them model this as a hard constraint. Furthermore, Rönnberg and Larsson (2010) also include the Friday Evening and Monday Day shifts in the weekend. Common to all the four papers is to evenly distribute the work weekends to all the employees.

It is common to find shift patterns that are considered undesirable to assign in the nurse scheduling literature. Examples include having a standalone shift, with pattern "work-off-work", or a single day off, with pattern "off-work-off". Azaiez and Al Sharif (2005), Aickelin and Dowsland (2004), Bard and Purnomo (2005b) and Glass and Knight (2010) all seek to minimize assignments of such patterns by penalizing it whenever they are assigned. Finally, Klyve and Beckmann (2016) include numerous types of shift patterns in their thesis. These include patterns rewarded in the objective function, patterns illegal to assign and patterns mandatory to assign. The thesis is focused on making a scheduling model which produces schedules that are so realistic they can be used as substitutes to manually created schedules, indicating that there may be many more shift patterns which should be taken into account in optimization-based scheduling models than what is common in the related literature.

An employee's work contract specifies how many hours the employee should work. Due to factors such as scheduling rules, shift work and requirements regarding the skill mix on each shift, it can be difficult to make schedules where the employees work the same number of hours every week. To enable a more flexible planning, it is common to allow that employees work more hours than contracted some weeks, and less hours than contracted other weeks (Rönnberg and Larsson, 2010; Klyve and Beckmann, 2016).

All the aspects discussed above are part of the scheduling process at DNIC. Upper limits on both the number of consecutive work shifts and night shifts are found. DNIC also has strict rules regarding work weekends, considering a work weekend the same way as described in the literature, and where each employee has to work every third weekend. Furthermore, at DNIC there are numerous shift patterns considered either good or bad to work, in line with the work by Klyve and Beckmann (2016). Finally, DNIC employs a mix of part-time and full-time employees, and has rules as to how the work hours of each employee may be distributed.

Preferences

Preferences are scheduling rules which are not mandatory in order to make a feasible schedule, but which the employer may choose to follow to satisfy the employees. This includes taking the specific needs and requests of each employee into account, and ensuring that all employees perceive the schedule is as fair.

Within self- and preference scheduling, there are several ways to treat the employees' personal requests. Ensuring that each nurse has a reasonable number of satisfied requests is a commonly used fairness measure (Bard and Purnomo, 2005b; Rönnberg and Larsson, 2010). A simple approach to this type of scheduling is to let the personnel request specific shifts on specific days, and then seek to maximize the fulfillment of these requests (Klyve and Beckmann, 2016; Aickelin and Dowsland, 2004). This is similar to the approach at DNIC. Another method is to incorporate the possibility of grading these requests. Lim and Mobasher (2011) describe a problem where all employees are asked to rank their requests, while Rönnberg and Larsson (2010) introduce a hierarchical structure for prioritizing requests. Such ranking of requests is, however, not currently possible to implement at DNIC.

A different kind of fairness measure is to ensure an even assignment of unpopular shifts such as weekend and night shifts among the nurses. As previously discussed, many papers include constraints meant to evenly distribute the work weekends to all employees. This is also commonly applied to night shifts, e.g. using hard constraints to ensure that each nurse is assigned a minimum number of night shifts (Azaiez and Al Sharif, 2005), or ensuring that all employees working night shifts work approximately the same number of such shifts (Rönnberg and Larsson, 2010). Finally, Klyve and Beckmann (2016) consider an even distribution of unpopular shifts as key, emphasizing the fact that automatic scheduling models assign shifts without any bias, which ensures that all employees are treated fairly and in the same way. A fair distribution of unpopular shifts is also an important consideration in the scheduling problem faced at DNIC.

3.3.3 Solution Methods for the Nurse Scheduling Problem

Several approaches to solving the deterministic nurse scheduling problem exist. Ernst et al. (2004) identify mathematical programming and metaheuristics as the most explored approaches in literature.

Exact Methods

The benefit of using exact methods, such as mathematical programming, is that we are guaranteed to find the optimal solution, given that such a solution exists. The drawback is that it can be very time consuming to find this solution, and that the formulations are sometimes NP-hard. Integer Programming (IP) is one of the simplest ways to formulate the problem, although finding the solution is sometimes harder. One approach is to solve the LP relaxation of the IP and then branch on the variables to find the optimal IP solution. Papers utilizing this approach include Klyve and Beckmann (2016) with their IP problem and Rönnberg and Larsson (2010) with a MIP problem.

To solve the problem more efficiently, the IP problem can be decomposed into a column generation problem (Ernst et al., 2004). Bard and Purnomo (2005b) decompose the nurse scheduling problem into an IP master problem as well as one subproblem for each nurse, which they solve using a heuristic approach.

Metaheuristics

Due to the complex size of the nurse scheduling problem, heuristic approaches with no guarantee of finding optimal solutions should be considered. Ernst et al. (2004) emphasize metaheuristics in particular have a great potential as compared to mathematical programming.

In the related literature, there are several examples metaheuristic approaches to solving the nurse scheduling problem. The most common approaches include genetic algorithms, tabu search, and simulated annealing (Van den Bergh et al., 2012). An example is Aickelin and Dowsland (2004), who develop a genetic algorithm to solve the nurse scheduling problem.

3.4 Uncertainty

Robust optimization, stochastic programming and combining optimization and simulation are three approaches which have been used to address uncertainty in supply and demand in the health care sector. However, the clear majority of publications on personnel scheduling ignores all types of uncertainty, and in particular uncertainty in supply (Van den Bergh et al., 2012). When it comes to the literature on personnel or nurse rescheduling, uncertainty in supply is usually accounted for in some way, while demand is assumed deterministic. Treating supply, demand, or both, as deterministic makes it very difficult to assess the robustness of a schedule.

3.4.1 Uncertainty in Demand

Lim and Mobasher (2011) use robust optimization to consider uncertainty in patient demand. Estimating bounds on uncertain demand parameters using historical data and applying them in a robust optimization approach, they seek to minimize patient dissatisfaction and nurse idle time as one of the objectives in a multiple objective nurse scheduling problem.

Bagheri et al. (2016) address the uncertainty related to demand and stay period of patients, proposing a two-stage stochastic optimization model for nurse scheduling using a distribution of stochastic variables from historical data. The methodology is applied to a scheduling problem for 18 nurses, with a planning period of 31 days. However, stochastic programming has some limitations when the model becomes big. First, it can be very difficult to handle the desired level of detail in a stochastic model, as the method requires a discretization of the possible outcomes. Further, as the number of daily scenarios and the planning horizon increases, the size of the problem explodes. This makes it practically impossible to realistically represent the problem and all its uncertainty.

Both Harper et al. (2010) and Kokangul et al. (2017), who focus on the nurse staffing problem, emphasize the importance of first knowing the best size and skill-mix of nurses, which depends on the patient demand, before the nurse scheduling problem is solved. Both papers use simulation to address uncertainty.

Kokangul et al. (2017) propose a model for finding optimal nurse staffing levels consisting of six steps, and apply the model to a Neonatal Intensive Care Unit (NICU) at a large hospital. The unit divides patients into three levels according to the degree of severity. The nurse-to-patient ratio is different for each level. Patients may be re-categorized into one of the other levels if their condition changes. Both the number of arrivals (both accepted, rejected and transferred), the transfer rates between levels and length of stays (LOS) are random, which makes the number of patients in each level a stochastic process. This is very similar to the situation at DNIC. Because the stochastic nature of the problem makes it difficult to determine the optimal staffing at each level, simulation is used as a decision support tool to evaluate the size of the staff. This thesis differs from this by assuming a fixed, pre-determined staff size. Simulation is only used to evaluate nurse schedules for this fixed staff.

According to Kokangul et al. (2017), in most studies, the patient arrival patterns are considered as Poisson processes. The LOS distribution profiles may vary depending on the patient type, but are typically either exponential, log-normal or Weibull distributions. This fits well with the results obtained at NICU. In the personnel rescheduling literature the uncertainty in demand is also commonly treated as a Poisson distribution in the cases where demand uncertainty is considered (Ingels and Maenhout, 2015, 2017, 2018).

In the preliminary work to this thesis, Løyning and Melby (2017) consider the uncertainty in demand at DNIC. As explained in Section 2.3.3, the patients are divided into five levels, each with its own required nurse-to-patient ratio. The paper proposes a model with dependencies between adjacent levels, and use historical data to calculate the probabilities for a change in one level conditional on the change in the adjacent level.

3.4.2 Uncertainty in Supply

Løyning and Melby (2017) split the absence of nurses into short-term and long-term, and model it using triangular distributions, where the shape of each distribution is based on average duration of short- and long-term absence in Norway. Combined with a model for demand uncertainty, they simulate the realizations of demand and supply and apply the results to a nurse schedule to evaluate under- and overstaffing after uncertainty is accounted for. However, they do not consider the possibilities of reestablishing the schedule when understaffing occurs.

Uncertainty in supply is treated in several ways within nurse rescheduling. Kitada and Morizawa (2010, 2013) randomly generate absences for 24 nurses in a 30-day schedule, where the former consider one-day absences and the latter consider absences with a duration of one to four days. Another approach is to generate nurse absences, or schedule disruptions, in a systematically varied and controlled way (Maenhout and Vanhoucke, 2011, 2013a,b), where the the disruptions are simulated by varying the total number of disruptions over the schedule horizon and the spread of the occurred disruptions. A third method is to assume that the uncertainty in supply can be modelled in an iterative manner using a Bernoulli distribution, where the probability of absence depends on whether the employee was available or absent on the day prior to the current (Ingels and Maenhout, 2015, 2017, 2018).

Uncertainty in supply is also considered in other papers on nurse rescheduling (Moz and Pato, 2003, 2004, 2007; Pato and Moz, 2008; Bard and Purnomo, 2005a; Bäumelt et al., 2016), although it is not clear from the papers which methods were used to generate the disruptions.

3.5 The Personnel Rescheduling Problem

The goal of the personnel rescheduling problem is to reestablish a feasible schedule in which the supply meets demand whenever an imbalance between them has occurred due to the realization of uncertainties in the online operational phase. The problem can be solved using resources internal to the unit studied (Moz and Pato, 2003, 2004, 2007; Pato and Moz, 2008; Kitada and Morizawa, 2010, 2013; Clark and Walker, 2011; Maenhout and Vanhoucke, 2011, 2013a,b; Bäumelt et al., 2016; Ingels and Maenhout, 2015, 2017, 2018) or external employees (Bard and Purnomo, 2005a). Reestablishing the schedule requires that the individual schedules of the employees are changed through actions such as assigning extra shifts or exchanging a work day and an off day, although the available actions varies between the problems studied.

With a few exceptions, who study the general personnel rescheduling problem (Ingels and Maenhout, 2015, 2017, 2018), all papers studied in this section concern the nurse rescheduling problem specifically.

3.5.1 Key Aspects Within Personnel Rescheduling

As for the scheduling problem, several aspects are important to consider when modelling and solving the personnel rescheduling problem. In this section, we discuss several key aspects which are taken into account in the related literature.

Objectives

A key objective within nurse rescheduling is to minimize understaffing. As service must be provided regardless of whether demand can be met or not, demand constraints are commonly modelled as soft constraints, where understaffing is penalized in the objective function (Clark and Walker, 2011; Maenhout and Vanhoucke, 2011, 2013b,a; Ingels and Maenhout, 2015, 2017, 2018). Some papers also minimize overstaffing, either directly by minimizing the number of excess employees assigned each shift (Clark and Walker, 2011; Maenhout and Vanhoucke, 2011, 2013a,b), or indirectly by minimizing costs (Ingels and Maenhout, 2015, 2017, 2018). However, understaffing tends to be heavier penalized than overstaffing.

Another approach towards understaffing is to state that if there are not sufficient available resources to meet demand, finding a feasible solution is outside the model scope, and should rather be dealt with in the decision problem of the hospital administration (Moz and Pato, 2003, 2004, 2007; Pato and Moz, 2008). Bard and Purnomo (2005a) take on this problem in a hospital-wide model, where the objective is to efficiently assign floaters, on-call nurses and agency nurses to the hospital units which are unable to avoid understaffing without the use of employees external to the units.

Nurse preferences when rescheduling consist primarily of retaining the current nurses individual shift assignments as much as possible (Moz and Pato, 2003). It is therefore not surprising that minimizing schedule changes is part of the objective in nearly all the related literature on personnel and nurse rescheduling (Moz and Pato, 2003, 2004, 2007; Pato and Moz, 2008; Kitada and Morizawa, 2010, 2013; Clark and Walker, 2011; Maenhout and Vanhoucke, 2011, 2013a,b; Bäumelt et al., 2016; Ingels and Maenhout, 2015, 2017, 2018).

Also included in some objectives are to account for the employees' preferences towards working specific shifts (Clark and Walker, 2011; Ingels and Maenhout, 2015, 2017) and to ensure a fair or even workload assignment after the rescheduling has been performed (Pato and Moz, 2008; Maenhout and Vanhoucke, 2011, 2013a,b). However, these objectives are generally given less weight than the objectives of minimizing understaffing and schedule changes.

The above mentioned objectives are also found at DNIC, whose most important considerations when it comes to making decisions during the rescheduling phase were discussed in Section 2.3.5. As in most of the related literature, meeting the overall demand is by far the most important goal. It is also important to not make more schedule changes than necessary, as both financial costs and inconvenience to the employees is related to rescheduling. This indicates that the rescheduling problem faced at DNIC is in fact is a multi-objective problem.

Constraints in the Rescheduling Problem

Although the objective and purpose of the scheduling and rescheduling problems differ, many of the same considerations regarding demand, time-related constraints and preferences have to be made in both problems.

When minimizing understaffing in the objective is not a part of the rescheduling problem, meeting demand is found as a hard constraint (Moz and Pato, 2003, 2004, 2007; Pato and Moz, 2008; Kitada and Morizawa, 2010, 2013; Bäumelt et al., 2016), sometimes also in terms of skills (Kitada and Morizawa, 2010, 2013).

The time-related constraints discussed in Section 3.3.2 are commonly found in the rescheduling problem. Working one shift per day, allowing sufficient rest between shifts, limiting the maximum number of consecutive shifts and putting a maximum limit on the total number of work shifts are hard constraints in most cases (Moz and Pato, 2003, 2004, 2007; Pato and Moz, 2008; Clark and Walker, 2011; Kitada and Morizawa, 2010, 2013; Maenhout and Vanhoucke, 2011, 2013a,b; Bäumelt et al., 2016; Ingels and Maenhout, 2015, 2017), although some allow working more than one work shift on the same day through overtime work (Bard and Purnomo, 2005a; Ingels and Maenhout, 2018).

A key difference between the scheduling and rescheduling problem is that absent employees cannot be assigned work shifts when rescheduling. This is commonly treated as a hard constraint, although some papers treat the constraint as soft to ensure feasibility (Maenhout and Vanhoucke, 2011, 2013a,b).

In much of the related literature, hard rules are imposed on several time-related aspects of the problem. This seems to differ from DNIC, where e.g. working more than one shift the same day and working more consecutive shifts than recommended is undesirable, but allowed. With less hard constraints regarding the work conditions in the online operational phase, it could be that it is easier to find a feasible solution in the rescheduling problem faced at DNIC than in some of the related literature.

Uncertainty Realization and Replanning Period

We define the replanning period as the set of consecutive days in which schedule changes are allowed to be made. The duration of the replanning period is commonly related to when it is assumed that disruptions are ascertained. Two main approaches are used in the related literature.

One approach is to assume that all disruptions for the upcoming planning period are known at the start of the period. The entire schedule is then reestablished once, with the first day of the replanning period being the first day of disruptions, and the last day being the final day of the schedule horizon (Moz and Pato, 2003, 2004, 2007; Pato and Moz, 2008; Clark and Walker, 2011; Kitada and Morizawa, 2010, 2013; Maenhout and Vanhoucke, 2011, 2013a; Bäumelt et al., 2016). A variation of this approach, under the same assumption, is performed by Maenhout and Vanhoucke (2013b), who study the optimal length of the rescheduling period under various problem settings.

A second main approach is to use a rolling horizon framework, in which it is assumed that disruptions are ascertained on the day of operation and the rescheduling problem is solved for three shifts at the time (Day, Evening and Night). Depending on how frequently new information is revealed, the problem is solved either once daily (Ingels and Maenhout, 2015, 2017, 2018) or three times daily (Bard and Purnomo, 2005a).

The demand at DNIC usually changes from day to day, and new information regarding absences is revealed daily. As discussed in Section 2.3.5, the scheduling manager can plan for schedule disruptions long time in advance only when the disruption is caused by a long-term absence. However, most disruptions are short-term, and demand is difficult to predict. Therefore, the latter approach seems to be method which most realistically imitates the operations at DNIC.

3.5.2 Rescheduling and Robustness

The related literature is focused on solving the rescheduling problem subject to various objectives intended to meet the preferences of employees and employers while still servicing demand. However, to the best of our knowledge, only a limited set of the papers have discussed scheduling and rescheduling in a robustness perspective.

As discussed in Section 3.2.1, robust schedules are stable and flexible. Proactive and reactive strategies can be used to improve the schedule robustness. Ingels and Maenhout (2015, 2017, 2018) study how different proactive and reactive strategies affect the robustness of personnel schedules using the three-step methodology described in Section 3.2.2.

Ingels and Maenhout (2015) introduce the possibility of proactively assigning reserve duties to the personnel according to five different strategies as to how these reserve duties should be assigned. In the online operational phase, the reserve duties can be transformed into working duties if demand exceeds supply according to two methods. In the fixed reactive method, the duties can only be converted to work duties of the same shift, which gives a good indication of how well the reserve duties are scheduled but offers low flexibility. The adjustable reactive method is more flexible, and allow the reserve duty to be converted to any work shift within the planning period, as long as the time-related constraints are respected.

Scheduling reserve duties gives a trade-off between scheduling costs in terms of wages and cancellation costs and the costs of shortages. Based on this trade-off, Ingels and Maenhout (2015) conclude that the buffer capacity in terms of reserve duties should be a fixed ratio of the minimum staffing requirements. Further, the adjustable reactive method is recommended to improve flexibility. However, the method introduce an additional trade-off between the number of changes to be made and the number of shortages, which should be carefully considered.

In a different approach, Ingels and Maenhout (2017) investigate how to proactively exploit the concept of substitutability to improve the flexibility of a schedule. Employee substitutability exists if an employee can take over the shift assignment of another employee with a particular skill on a particular day.

Ingels and Maenhout (2018) also investigate the use of overtime as a way to improve the robustness, either by proactively scheduling overtime to improve the stability of the baseline schedule, or by allowing to reactively assign overtime work to improve the flexibility. They also discuss the trade-off between the hiring budget and the overtime budget on the tactical level, and how allowing the use of overtime work may reduce the required number of staff as it increases the operational flexibility.

It is clear that the underlying scheduling strategies when making the baseline schedule have a strong impact on the stability and flexibility of schedules; the scheduling at DNIC being no exception. DNIC experience fluctuations in both demand and supply, and it seems very reasonable to study whether proactive strategies such as the ones discussed in this section can improve the stability and flexibility of the schedules created.

3.5.3 Solution Methods for the Personnel Rescheduling Problem

Both exact methods and heuristics are used to solve the rescheduling problem in the related literature.

Exact Methods

Key benefits of using exact methods were discussed in Section 3.3.3. Ingels and Maenhout (2015, 2018, 2017) solve their relatively small rescheduling problems using IP optimization by applying the commercial optimization software Gurobi, and quickly find a solution to their problem in most instances. Bard and Purnomo (2005a) develop a branch-and-price algorithm to find the solution to their IP, which is larger in size, using Dantzig-Wolfe decomposition to divide the problem into a restricted linear Master Problem and one subproblem per nurse. A set covering heuristic is used to find good initial columns. Yet another approach is taken by Moz and Pato (2003, 2004), who solve their IPs using multicommodity flow models.

Heuristics

Heuristic approaches, which do not guarantee finding an optimal solution, can be beneficial to introduce when the problem grows larger. Moz and Pato (2003, 2004) experience a large increase in the solution time for the most complex instances. Therefore, Moz and Pato (2007) propose a genetic heuristic. Pato and Moz (2008) develop the heuristic further, introducing an additional objective, and in turn a Pareto strategy for finding dominating solutions.

Kitada and Morizawa (2010) propose a heuristic tree-search method in two steps. First, they use a recursive search method to find a feasible solution, by generating nodes corresponding to all candidate nurses who can fill up a disrupted shift at each recursive level, and using these nodes to construct an initial search tree. Using the feasible schedule as an incumbent, they backtrack to promising nodes using a depth-first strategy to obtain better feasible schedules. Kitada and Morizawa (2013) divide absences of consecutive days into one subproblem for each day, and apply the algorithm from Kitada and Morizawa (2010), but with slight improvements aimed at more efficiently reestablishing feasibility when the absentee is a high-skilled nurse.

Other examples of heuristic solution methods utilized in the related literature are an evolutionary meta-heuristic operating on a Pareto optimal set of solutions (Maenhout and Vanhoucke, 2011, 2013b), and an artificial immune system which revises and re-optimizes a schedule for a set of heterogeneous nurses (Maenhout and Vanhoucke, 2013a).

3.5.4 Comparison of Papers on Personnel Rescheduling

In Table 3.3 we compare key aspects of this thesis to 12 of the papers on rescheduling that have been discussed in this chapter. The numbers in the top row correspond to the numbering of the papers in Table 3.1.

Table 3.1: Key papers on rescheduling reviewed in this thesis

[1]	This thesis	[8]	Maenhout and Vanhoucke (2013a)
[2]	Bard and Purnomo (2005a)	[9]	Maenhout and Vanhoucke (2013b)
[3]	Pato and Moz (2008)	[10]	Ingels and Maenhout (2015)
[4]	Clark and Walker (2011)	[11]	Ingels and Maenhout (2017)
[5]	Kitada and Morizawa (2010)	[12]	Ingels and Maenhout (2018)
[6]	Kitada and Morizawa (2013)	[13]	Bäumelt et al. (2016)
[7]	Maenhout and Vanhoucke (2011)		

As we perceived some aspects as ambiguously presented in some papers, the reader should be aware that the table only represents how we best interpreted the contents of each paper. Also, in three cases the answer to the question was unclear, indicated by $-$ in the table. The abbreviations used in the table are indicated in Table 3.2.

Table 3.2: Abbreviations used in Table 3.3

H	Heuristic method	N	Nurse
E	Exact method	P	Personnel

As indicated in Table 3.3, while roughly half of the papers base their underlying schedules on a real-life case, none of them utilize historical data in the generation of disruptions. In this thesis, historical data from DNIC is used to calculate the probability distributions on which both the demand and absence simulation models are based. Combined with the fact that the scheduling model in this thesis creates schedules for DNIC using real data as input, we believe that our thesis is among the papers which most realistically evaluates the true performance of nurse or personnel schedules to date.

Very few papers are similar to this thesis in assumptions and methodology. The most similar papers are Ingels and Maenhout (2015, 2017, 2018), which all treat demand and supply as uncertain, use a rolling-horizon model with exact solution methods to perform the rescheduling, and have a specific focus on robustness evaluation. However, the three papers consider personnel scheduling in general, while this thesis concerns a specific nurse scheduling case.

Although we consider simulation combined with a rescheduling model as a good tool in order to evaluate the schedule robustness, few of the papers reviewed do so in practice. Much of the research, specially the papers using heuristic solution methods, seem more focused on studying how different methods cope with solving the rescheduling problem, subject to different objectives and constraints as discussed in Section 3.5.1, than to evaluate the quality of the underlying schedule. Clark and Walker (2011) discuss how two different schedules with corresponding rescheduling strategies perform in terms of fulfilling nurse preferences, but do not mention robustness. Furthermore, Maenhout and Vanhoucke (2013b) study the consequences and outcomes of various nurse rescheduling characteristics and strategies, but not specifically in a robustness context. In addition to this thesis, we therefore only consider Ingels and Maenhout (2015, 2017, 2018) as papers which truly analyze robustness in a structured manner.

Table 3.3: Comparison of this thesis ([1]) and 12 key papers within personnel and nurse rescheduling

Aspect	[1]	[2]	[3]	[4]	[5]	[6]	[7]	[8]	[9]	[10]	[11]	[12]	[13]
Uncertain demand?	Yes	Yes	No	No	No	No	No	No	No	Yes	Yes	Yes	No
Uncertain supply?	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Absence duration	Multi	–	Multi	–	1 day	1-4 days	1 day	1 day	1 day	Multi	Multi	Multi	Multi
Absence duration known?	Partially	–	Yes	Yes	Yes	Yes	Yes	Yes	Yes	No	No	No	Yes
Robustness evaluation?	Yes	No	No	No	No	No	No	No	No	Yes	Yes	Yes	No
Rolling horizon rescheduling?	Yes	Yes	No	No	No	No	No	No	No	Yes	Yes	Yes	No
Schedule based on real-world case?	Yes	Yes	Yes	No	Yes	Yes	No	No	Yes	No	Yes	No	Yes
Real-world data in disruption generation?	Yes	No	No	No	No	No	No	No	No	No	No	No	No
External employees available?	No	Yes	No	No	No	No	No	No	No	No	No	No	No
Skills?	Yes	Yes	No	No	Yes	Yes	No	No	No	No	Yes	No	No
Employees	117	20-200	19/32	20	24	24	30	30	30	10/20	20	10	19/32
Solution method rescheduling	E	E	H	E	H	H	H	H	H	E	E	E	H
Employee group	N	N	N	N	N	N	N	N	N	P	P	P	N

Chapter 4

Problem Description

The overall problem studied in this thesis concerns how to create robust schedules at the Department of Neonatal Intensive Care (DNIC), as well as defining measures to evaluate this robustness. As discussed in Chapter 3, robustness concerns both the degree to which a personnel schedule can absorb disruptions, and its capability to react efficiently to disruptions. We therefore consider the problem as twofold, divided into the offline operational scheduling problem and the online operational rescheduling problem.

In the offline operational scheduling problem, the goal is to make a feasible schedule for the upcoming planning period. In this phase, uncertainty in demand and supply is accounted for using expected values and insights from previous planning periods. The schedule is put into use in the online operational phase, where information about the real demand and absence of employees is revealed daily. Whenever the demand exceeds the capacity, actions can be taken to restore the balance between the two. The uncertainty realization and the ease to which a feasible schedule can be restored may give important managerial insights, which in turn can aid in making more robust schedules in future planning processes. The overall goals of each process, the time horizon of the problems and the way uncertainty affects them are summarized in Figure 4.1.

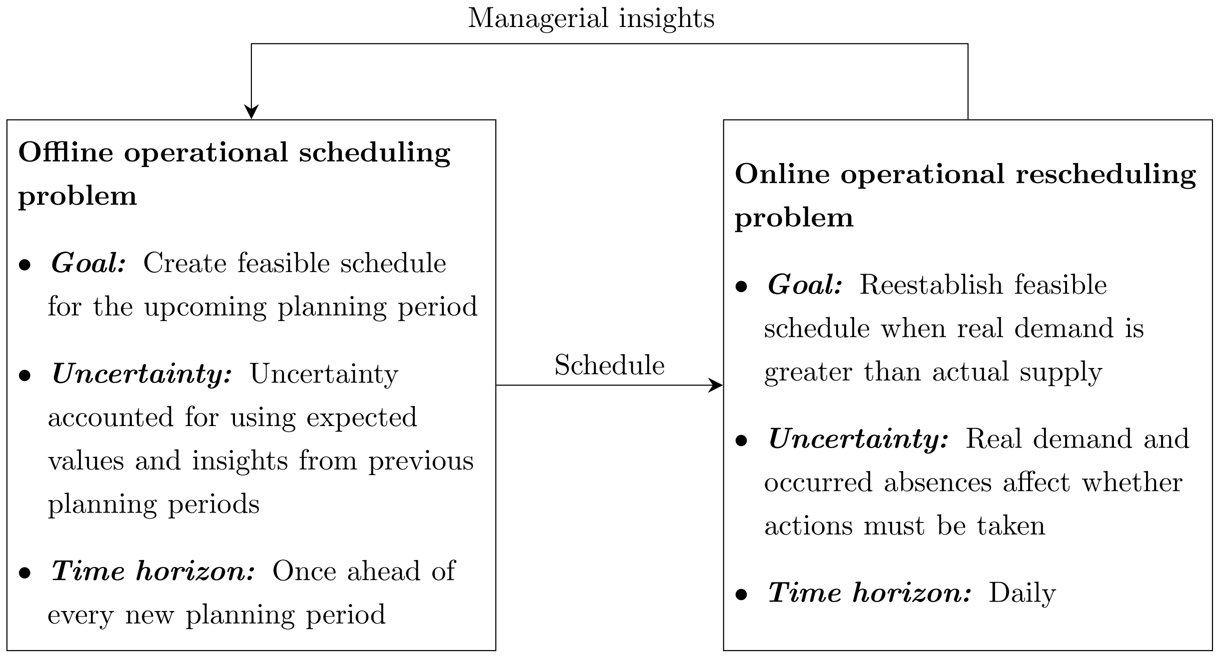

Figure 4.1: The overall robustness problem faced at DNIC

In the remainder of this chapter, we provide a further description of the two problems. We start by elaborating on the offline operational scheduling problem, which is also denoted as the scheduling problem, in Section 4.1. In Section 4.2 we describe the uncertain aspects of both patient demand and employee absence. We present the online operational rescheduling problem, also called the rescheduling problem, in Section 4.3. Finally, in Section 4.4, we introduce several measures of robustness that are intended to provide the managerial insights displayed in Figure 4.1.

4.1 Description of the Scheduling Problem

The scheduling problem involves making a feasible and satisfactory schedule for a group of employees for a given time period. Feasibility means that the demand for employees is satisfied, while respecting relevant rules and regulations in each employee's schedule. Additionally, a satisfactory schedule takes the preferences of the employers and employees into account.

Shifts and requests All employees must be assigned exactly one shift per day in the planning period. This shift can be either a work shift or an off shift. A work shift is either a Day, an Evening or a Night shift, sometimes also denoted by D, E, and N, respectively. Night shifts are assigned to the day when the shifts end as indicated by Figure 4.2. That is, a Night shift on a Monday is the shift that starts Sunday evening and ends Monday morning. Prior to a new planning period, the employees report requests regarding their personal schedules in the upcoming planning period. These should be taken into account.

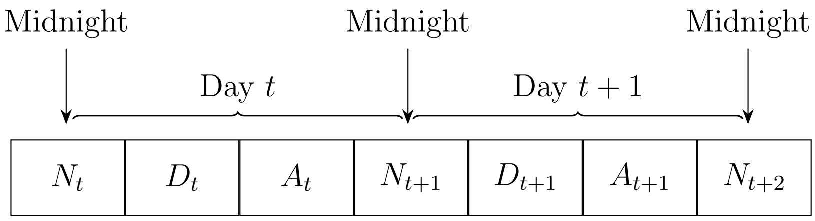

Figure 4.2: Placement of the Night shift in the scheduling problem

Contracted work Each employee has a work contract stating the number of hours he or she should work during the planning period. The employees' work contracts remain the same throughout the planning period. Only small deviations from the contracted work are allowed in the schedule. An employee's workload may vary from week to week, as long as the average weekly workload over the planning period is as stated in the work contract. However, an employee cannot work more than a specified number of hours during each 7-day period.

Demand For each work shift, there is an overall demand for a given number of employees which must be met. Overstaffing up to a specified limit is allowed, while understaffing is not. The employees are divided into hierarchically organized skill categories. These decide which shifts are feasible for which employees. On all shifts, there is a demand for a specified number of employees of each skill. Employees can only cover demand for skills corresponding to its own or a lower skill level. However, the skill-wise demand for employees should be covered by employees with that skill as their main skill to the farthest extent possible. An employee cannot be assigned to cover demand for more than one skill category at each shift. New hires have a training period lasting a specified number of weeks. During these weeks, they are assigned work shifts, but are not counted as one of the employees that contribute to covering the daily demand.

There is also demand for employees who can perform some special tasks, such as working at the outpatient clinic. These tasks are always performed during the Day shift. The shifts meant for performing these tasks are explicitly requested by the employees responsible for them, and have to be assigned whenever requested. Whenever an employee works one of these special shifts, he or she is not counted as one of the employees who contribute to covering the daily demand. New hires cannot take on these special functions.

Required rest Some hard scheduling rules must be followed when making a new schedule. First, there must be a minimum number of hours of rest between each work shift assigned to an employee. Further, each employee must have a protected off day every week. If the employee has a weekend off, this day must be on Sunday. Employees must rest at least a minimum number of consecutive hours when they have their protected off day, resulting in some illegal combinations of work shifts and the protected off shift.

Employees cannot work more than a maximum number of consecutive days, or more than a maximum number of consecutive Night shifts. Further, for each employee, at least a specified minimum of shifts must be an Evening shift or a Night shift. Night shifts have a greater workload than Evening shifts. This rule does not apply to some of the employees who cover the special tasks.

Weekends Weekends are not popular to work, and special scheduling rules are applied here. Each employee has to work every third weekend. During work weekends, they must work both Saturdays and Sundays. In addition, employees working Night shifts during their work weekend must work Night both Saturday, Sunday and Monday.

Shift patterns Some scheduling rules are related to patterns of consecutive shifts. Certain legal shift patterns are considered undesirable by employees and management, and should only be assigned when an employee explicitly requests to work that pattern. An example of such a shift pattern is to work a Night shift between two off shifts. On the other hand, some other shift patterns are considered very beneficial to work, and should be assigned whenever possible. Finally, there are some patterns of successive shifts which should never be assigned, either due to governmental requirements or because the patterns are very unpopular to work.

Objective The objective is to maximize the number of granted requests in the planning period and the number of desirable shift patterns assigned. Cases where employees with more advanced skills than required cover the minimum demand for employees with a certain skill should be penalized.

4.2 Description of the Uncertainty Realization

To model uncertainty means to formulate one or more models that realistically imitate the uncertainty in the demand and supply of employees. The uncertainty in supply is divided into employee absence and availability.

4.2.1 Uncertainty in Demand

The patients are divided into five levels, according to the severity of their conditions. The overall demand for employees is thus calculated based on the number of patients per level. An examination of each patient's condition is performed daily during the Day shift, and patients may change level based on the results. The levels are hierarchically organized according to severity, and thus patients can move to any level depending on whether their conditions improve or get worse. There is a large variance in the duration of the required treatment of each patient, both in terms of the number of days the patient is admitted to a certain level and in terms of the total length of stay.

The number of patients admitted varies from day to day. Each day patients may be admitted, discharged, or change levels.

4.2.2 Uncertainty in Absence

An absent employee is defined as an employee who is unable to work for a given time period. Both working and non-working employees can be absent. We consider both short-term and long-term absences, where a short-term absence is unpredictable in length, but usually shorter than a week, while a long-term absence has a predictable duration.

Both the frequency and the duration of the employee absences are uncertain. Each day, an employee may enter a short- or long-term absence, change from being absent to non-absent.

4.2.3 Uncertainty in Availability

The uncertainty in supply also concerns whether the employees will accept working unplanned shifts or changing their initial schedule, and whether they have requested extra shifts or exchanges on a given day. We therefore define available employees as employees who are non-absent and willing to work.

4.3 Description of the Rescheduling Problem

The rescheduling problem involves the reestablishment of staffing for a previously created schedule when the real demand for employees is higher than the number of non-absent employees in the schedule. This can happen either because the demand is higher than expected, because employees are absent from their scheduled shifts, or both.

Figure 4.3 sums up the inputs, events, and outputs of the rescheduling problem. Each day, the initial schedule from the scheduling problem together with potential changes made in former rescheduling problems are compared to the actual demand and employee availability. If necessary, additional changes to the schedule are made. These changes affect the schedule on the day of the rescheduling, but may also affect the schedule on later days. The process is repeated iteratively.

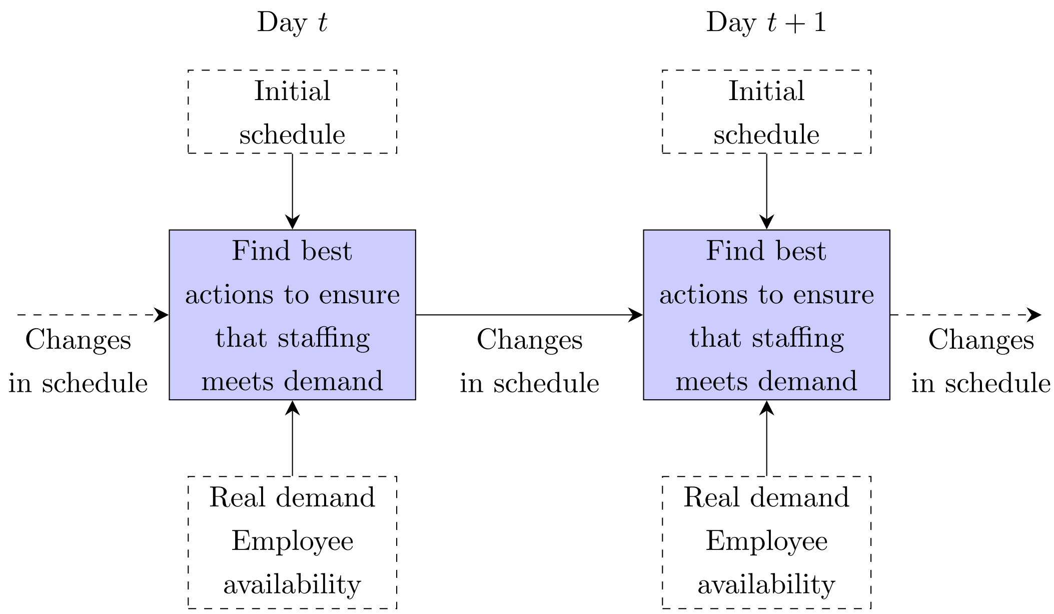

Figure 4.3: Overview of the online operational rescheduling problem at DNIC, where the dashed boxes represent inputs to the daily problems

Replanning period The rescheduling problem is solved daily, taking the real demand and occured absences into account. The set of days in which changes are allowed is called the replanning period. This always includes the current, and may include a set of consecutive days after this. The pre-period is the set of consecutive days prior to the current day, and the post-period is the set of consecutive days after the last day of the replanning period. The shifts worked in the pre-period and the planned shifts in the post-period are important to take into account to ensure optimal decisions.

Shifts There are three daily shifts; Day, Evening and Night. The Night shifts belong to the days when the shifts *start*, as indicated in Figure 4.4. This is as opposed to the scheduling problem in Section 4.1, where the Night shifts falls on the day when the shifts *end*. It is not necessary to distinguish between the mandatory weekly off shift and the remaining off shifts, as DNIC sees these days as equal when solving their real-life rescheduling problem.

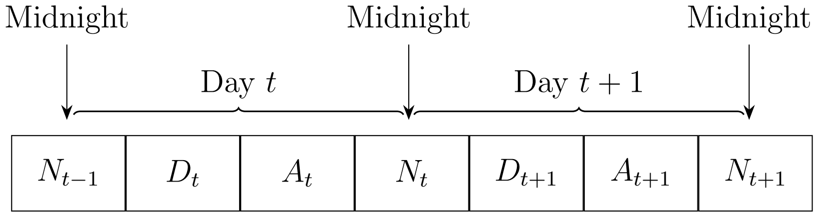

Figure 4.4: Placement of the Night shift in the rescheduling problem

Demand Both the demand for employees of each skill and the overall demand should be met. The actual demand for employees is calculated based on the number of patients per level, multiplied by the expected need for nurses per patient on each level, and both this demand and the department's lower staffing limit should be taken into account. This means that it might be necessary to take actions to ensure sufficient staffing on any of the shifts within the replanning period.

The demand has to be covered by employees at DNIC alone, as there are requirements of having special training in order to properly perform each task. Further, employees with a special shift scheduled may not cover the daily demand, as explained in Section 2.3.1.

Availability Recall from Section 4.2.3 that an available employee is an employee that is non-absent and willing to work. For working employees this means that he or she is available if he or she is non-absent. A non-working employee has to be willing to work on an off day in addition to being non-absent. Only available employees may take on unplanned shifts.

Actions Available, non-working employees can take on extra shifts, meaning that the employees perform a work shift on their off day, without any other changes to their schedule. Employees can give notice that they want to work extra, given shifts ahead of the shifts.

It is possible for employees to work maximum two shifts the same day. Working double shifts should only be allowed on the current day or in the transition to the current day. Both two consecutive shifts without a break in between, and two shifts within the same day, but with a break of one shift in between, are considered double shifts.

Exchanging means that an employee's planned off shift during the replanning period can be exchanged with a planned work shift another day during the replanning period, such that the employee works on the initial off day and not on the initially planned work day. This may happen if that employee has requested an off shift on a specific day during the replanning period.

All available employees with a work shift scheduled sometime during the replanning period can have a shift swapped. Swapping works the same way as an exchange, but incurs a higher cost. Any combination of swaps from one work shift to another within the replanning period is allowed, as well as swapping from one shift to another the same day.

When employees are exchanged or swapped, they still work the same number of shifts as initially planned, only at different times. When swapping or exchanging employees from future work shifts on long notice, a given buffer on top of the minimum or expected real demand must be left. Otherwise, it could be that the absence of one employee or a small rise in demand would lead to problems with understaffing.

When employees are assigned new shifts as a result of the actions performed, the shifts should not be changed again in the future rescheduling problems.

Evaluation factors The evaluation factors for actions regard costs and inconvenience, and should be penalized in the objective function.

Overtime costs in a given week occur whenever an employee works either more hours than a full-time work week that week, or more hours than what they were initially scheduled for, if they were scheduled to work more than a full-time work week due to the averaging of work hours over the planning period. Working extra shifts or double shifts may lead to such overtime pay. When two shifts are exchanged or swapped, no overtime costs occur. The total cost of an exchange is thus equal to the wage cost of the new shift, less the wage cost of the old work shift. If the employees work double shifts or swap shifts, an additional financial compensation is offered per occurrence of the action, which is unrelated to the normal overtime pay.

When changing the initial schedule of an employee, the employee experiences an inconvenience which depends on what kind of action was taken. There is also inconvenience related to working too many consecutive shifts or consecutive nights.

Objective The overall objective is to meet demand. Consequently, unmet demand results in a penalty. Furthermore, it is desirable to not make more schedule changes than necessary, and thus to minimize the number of extra shifts, exchanges, swaps, and double shifts. The penalty related to assigning an extra shift depends on whether the shift was requested. Swaps are penalized differently depending on whether the notice given to the employee whose shift was swapped was provided short or long time in advance of the new shift. Similarly, a double shift where the two shifts follow directly is penalized heavier than a double shift where there is a period of rest between the shifts. Finally, the evaluation factors described in the previous paragraph should also be considered in the objective.

4.4 Description of the Robustness Evaluation

The final step of the problem is to evaluate the performance of the initial schedule in the online operational phase. Insights from this analysis may be used proactively to make future schedules more robust. To gain these insights, different measures of robustness are defined. The measures are divided into stability and flexibility. Recall from Section 3.2 that stability is the degree to which the schedule is able to absorb unexpected events, while flexibility is about how easily the schedule can reestablished when these unexpected events lead to imbalances in demand and supply (Ionescu and Kliewer, 2011).

4.4.1 Robustness Measures

The robustness measures to be used are defined in Table 4.1. The stability measures are used to assess how the initial schedule performs after the uncertainties in demand and supply have become known, but prior to rescheduling, while the flexibility measures assess the robustness after rescheduling actions have been taken.

Table 4.1: Robustness measures, divided into stability and flexibility

Measure	Description	Type
Frequency_pre	On average, how often is DNIC understaffed per shift prior to rescheduling?	Stability
Severity_pre	On average, how severe is the shortage on each understaffed shift prior to rescheduling?	Stability
Overtime	How many overtime hours were worked?	Flexibility
Rest violation	How frequently did the employees work more than the recommended number of consecutive shifts and nights?	Flexibility
Actions required	How many swaps, exchanges, extra shifts and double shifts occurred?	Flexibility
Frequency_post	On average, how often is DNIC understaffed per shift after rescheduling?	Flexibility
Severity_post	On average, how severe is the shortage on each understaffed shift after rescheduling?	Flexibility

4.4.2 Managerial Insights

By testing various proactive and reactive scheduling strategies and assessing them using the stability and flexibility measures defined, managerial insights can be obtained. Specifically, we want information to help answering the questions in Table 4.2. The insights can aid in overcoming the challenges faced in the scheduling process which were described in Section 2.3.4.

Table 4.2: Managerial insights from proactive and reactive strategies

Insight	Question solved	Strategy type
Optimal assignment of surplus work hours	Can the contracted work hours which do not have to contribute to covering the minimum demand be distributed in a more robust way?	Proactive
Optimal placement of off shifts	Can off shifts in the initial schedule be placed in such a way that the violations of rules regarding rest and overtime work in the rescheduling problem are kept at a minimum level?	Proactive
Value of utilizing historical absence information	Does the robustness improve if historical information about absence is taken into account when scheduling?	Proactive
Value of additional staffing during weekends	Does the robustness improve if more weekend work is distributed to some employees in exchange for extra off days?	Proactive
Optimal replanning period	How far in advance is it optimal to assign extra shifts? How does the robustness change when the replanning period changes?	Reactive
Necessity of violating rules	Is it possible to avoid violating the rules regarding rest and work hours?	Reactive

Chapter 5

Scheduling Model

In this chapter, we present the mixed integer mathematical formulation of the scheduling model, as well as several model extensions. The scheduling model is based on the DNIC Scheduling Model by Løyning and Melby (2017). The chapter starts with an explanation of the assumptions and simplifications on which the model is based in Section 5.1. The indices, sets, parameters and variables of the mathematical model are defined in Section 5.2, while the objective function and constraints are presented in Sections 5.3 and 5.4, respectively. To increase the reader's understanding of the model's outputs, we provide a simple example of a feasible schedule in Section 5.5. Finally, we formulate four proactive model extensions in Section 5.6. Compressed versions of the scheduling model and the model extensions are provided in Appendix A.1.

5.1 Assumptions and Simplifications

Because this thesis also concerns the online operational part of the scheduling process, the scope of the scheduling model must be somewhat limited. This implies that not all rules and preferences considered in the real-life scheduling process should be included, and that the model cannot be used to substitute the manual process. Table 5.1 summarizes the underlying assumptions of the model.

First, we assume that no special events such as Christmas, Easter and vacation periods occur. Special scheduling rules apply during these periods, and we do not consider

Table 5.1: The assumptions of the scheduling model

Assumption 1	No holidays or vacation periods
Assumption 2	No personal inclinations
Assumption 3	Sufficient fairness obtained by following departmental guidelines
Assumption 4	Constant number of employees throughout the planning period

implementing these rules as vital. Furthermore, personal inclinations, which are special scheduling rules that only concern a limited set of employees, should be omitted. For example, in real-life many employees older than 50 years do not work Night shifts due to health reasons. However, we do not consider the extra realism added by including this type of personal inclinations for all employees as crucial in order to realistically evaluate the schedule robustness later on.

Some fairness measures are included in the scheduling guidelines at DNIC, such as ensuring a fair assignment of Evening and Night shifts to all employees. We consider implementing these rules as sufficient to obtain a reasonably fair assignment of shifts in the schedules.

The size of the staff and the contracted work of each employee is determined ahead of the scheduling process. Therefore, determining the optimal staffing level should not be a part of the model. Due to the predetermined number of employees, costs are considered sunk, and should not be included in the scheduling model.

5.2 Definitions

In this section, we present the indices, sets, parameters and variables used in the model. We name sets using uppercase calligraphic letters, variables using lowercase letters, and parameters using uppercase letters. Subscripts indicate indices, while superscripts of capital letters specify the meaning of some parameters and sets. Some parameters have over- or underlines to indicate that they represent upper or lower limits, respectively.

5.2.1 Indices

n	employee
s	shift
t	day
k	week
c	skill
p	shift pattern

5.2.2 Sets

$\mathcal{N}$	set of employees
$\mathcal{C}$	set of skills, $\mathcal{C} = \{1, 2, 3, 4\}$, where 1 = Emergency skills, 2 = Intensive Care skills, 3 = Monitoring skills, 4 = Assistant nurse skills
$\mathcal{N}_c$	set of employees with skill c as their highest ranked skill, $\bigcup_{c \in \mathcal{C}} \mathcal{N}_c = \mathcal{N}$
$\mathcal{N}^{GEN}$	generic set of employees, explained further whenever used
$\mathcal{S}$	set of shifts
$\mathcal{S}^W$	set of work shifts, $\mathcal{S}^W = \{D, E, N\}$, $\mathcal{S}^W \subset \mathcal{S}$
$\mathcal{S}^O$	set of off shifts, $\mathcal{S}^O = \{F1, F\}$, $\mathcal{S}^O \subset \mathcal{S}$
$\mathcal{T}$	set of days in the current planning period
$\mathcal{T}^{SUN}$	set of Sundays
$\mathcal{T}^{\mathcal{P}}$	set of days in the current and previous planning period, where $t \leq 0$ indicates days in the previous period and $t = 1$ is the first day in the current period
$\mathcal{K}$	set of weeks in the planning period
$\mathcal{T}_k$	set of days in week k, $\bigcup_{k \in \mathcal{K}} \mathcal{T}_k = \mathcal{T}$
$\mathcal{P}^{DW}$	set of shift patterns occurring during a weekend which are desirable to assign
$\mathcal{P}^{GEN}$	set of shift patterns that are considered undesirable, and should only be assigned if the employee specifically requests it
$\mathcal{P}^{ILL}$	set of shift patterns illegal to assign

5.2.3 Parameters

The parameters used in the model are sectioned by their functionality, and are divided into limit, weighing, indicator and general parameters.

Limit Parameters

The limit parameters set the upper and lower limits for the constraints. D is used for the demand parameters, M denotes parameters which regard the number of consecutive shifts of different types and H is used to denote limits on the number of hours worked.

$\underline{D}_{st}$	minimum number of employees required to cover total demand at shift s on day t
$\overline{D}_{st}$	maximum number of employees allowed to work at shift s on day t
$\underline{D}^{C}_{cst}$	minimum number of employees required to cover demand for skill c at shift s on day t
$\overline{M}^{CW}$	maximum number of consecutive work shifts for each employee
$\overline{M}^{N}$	maximum number of consecutive Night shifts for each employee
$\overline{H}^{7D}_{n}$	maximum number of hours employee n can work during a 7-day period

Weighing Parameters

The weighing parameters are used to reward or penalize certain behaviors in the objective function. All weighing parameters are denoted W.

W^R	reward for assigning a requested shift
W^P	reward for assigning a desirable shift pattern
W^S	penalty if demand for an employee of a particular skill is covered by an employee with a more advanced skill type as their main skill

Indicator Parameters

All indicator parameters are binary, and denoted β.

β^{PA}_{nst}	1 if employee n should have shift s preassigned on day t, 0 otherwise
β^{NA}_{nst}	1 if employee n should never have shift s assigned on day t, 0 otherwise
β^{N}_{nt}	1 if employee n can cover demand on day t, 0 otherwise
$\beta^{F1}_{s_1 s_2}$	1 if there is sufficient time between shifts s_1 and s_2 on days $t-2$ and t, respectively, for an employee to be assigned an 'F1'-day on day $t-1$, 0 otherwise
β^{R}_{nst}	1 if employee n has requested shift s on day t, 0 otherwise

General Parameters

The general parameters are parameters that do not fit into any of the previously defined categories.

H_s	duration of shift s in hours
H^{CW}_n	number of hours employee n should work during the planning period
H^{DEV}	allowed deviation in percent from H^{CW}_n for the number of hours assigned
M^{NW}	employees work every M^{NW} weekend
U	minimum amount in percent of shifts which, for each employee, must be an Evening or a Night shift
B	work load of a Night shift relative to an Evening shift
L_p	duration of shift pattern p in days
L^S_p	start day of shift pattern p relative to Sunday, where $L^S_p = 2$ is a Saturday, $L^S_p = 3$ is a Friday, and so on, where $p \in \mathcal{P}^{DW}$
S_{tp}	shift type on day t in shift pattern p

5.2.4 Variables

When $t \leq 0$, variables x_{nst} represent shift assignments in the previous planning period. This allows for continuity between periods, and is useful in several constraints, like e.g. constraints (5.7), which assign work weekends and have a rolling three-week time horizon. The variables x_{nst} from the previous planning period are given as inputs to the model.

$$x_{nst} = \begin{cases} 1 & \text{if employee } n \text{ is assigned shift } s \text{ on day } t \\ 0 & \text{otherwise} \end{cases}$$

$y_{cst} =$ number of employees in a different skill group than group c who have to rank down to cover demand for that skill on shift s on day t

$$w_{ntp} = \begin{cases} 1 & \text{if employee } n \text{ works desirable shift pattern } p \text{ containing } t\text{, where } t \in \mathcal{T}^{SUN} \\ 0 & \text{otherwise} \end{cases}$$

5.3 Objective Function

$$\max Z = W^R \sum_{n\in\mathcal{N}} \sum_{s\in\mathcal{S}} \sum_{t\in\mathcal{T}} \beta^R_{nst} x_{nst} + W^P \sum_{n\in\mathcal{N}} \sum_{t\in\mathcal{T}^{SUN}} \sum_{p\in\mathcal{P}^{DW}} w_{ntp} - W^S \sum_{c\in\mathcal{C}} \sum_{s\in\mathcal{S}^W} \sum_{t\in\mathcal{T}} y_{cst} \tag{5.1}$$

The objective function (5.1) maximizes the number of personal requests granted throughout the planning period, as well as the number of desirable shift patterns assigned. The final term of the objective function penalizes cases where demand for employees with skill c is covered by employees with more advanced skills than required.

5.4 Constraints

All the constraints used in the model are presented in this section. The constraints are grouped by functionality. For example, constraints related to covering demand are listed in the same subsection.

5.4.1 Covering Demand

$$\sum_{s\in\mathcal{S}} x_{nst} = 1 \qquad n \in \mathcal{N}, t \in \mathcal{T} \tag{5.2}$$

$$\underline{D}_{st} \leq \sum_{n\in\mathcal{N}} \beta^N_{nt} x_{nst} \leq \overline{D}_{st} \qquad s \in \mathcal{S}^W, t \in \mathcal{T} \tag{5.3}$$

$$\sum_{i=1}^{c} \sum_{n \in \mathcal{N}_i} \beta_{nt}^{N} x_{nst} \geq \sum_{i=1}^{c} \underline{D}_{ist}^{C} \qquad c \in \mathcal{C}, s \in \mathcal{S}^W, t \in \mathcal{T} \tag{5.4}$$

$$y_{cst} + \sum_{n \in \mathcal{N}_c} \beta_{nt}^{N} x_{nst} \geq \underline{D}_{cst}^{C} \qquad c \in \mathcal{C}, s \in \mathcal{S}^W, t \in \mathcal{T} \tag{5.5}$$

Constraints (5.2) state that each employee should be assigned exactly one shift per day. Constraints (5.3) make sure that the number of employees assigned work shifts is between the minimum and maximum demand on all shifts, while constraints (5.4) ensure that the minimum demand for employees on each skill level is met on every shift. Constraints (5.4) allow that demand for a skill type is covered by employees with the same or more advanced skills as the skill demanded. Finally, constraints (5.5) keep track of whether demand for particular skills on each shift is covered by employees with that skill as their main skill, or if employees with higher ranked skills will have to step in. In constraints (5.3) to (5.5), parameters β_{nt}^{N} exclude employees that cannot contribute to covering demand on day t, i.e. new hires during their training period and employees performing any of the special tasks that day.

5.4.2 Weekends

$$\sum_{s \in \mathcal{S}^W} (x_{ns(t-1)} - x_{nst}) = 0 \qquad n \in \mathcal{N}, t \in \mathcal{T}^{SUN} \tag{5.6}$$

$$\sum_{s \in \mathcal{S}^W} \sum_{\tau=0}^{M^{NW}-1} x_{ns(t-7\tau)} \leq 1 \qquad n \in \mathcal{N}, t \in \mathcal{T}^{SUN} \tag{5.7}$$

$$2x_{nN(t-1)} - x_{nNt} - x_{nN(t+1)} = 0 \qquad n \in \mathcal{N}, t \in \mathcal{T}^{SUN} \tag{5.8}$$

Constraints (5.6) state that employees must either work both Saturday and Sunday, or have both days off. In constraints (5.7) it is ensured that employees work every M^{NW} weekend. Constraints (5.8) enforce that employees working Night shifts during a weekend, work Night both Saturday, Sunday and Monday.

5.4.3 Work Hours

$$H_n^{CW}(1 - H^{DEV}) \leq \sum_{s \in \mathcal{S}^W} \sum_{t \in \mathcal{T}} H_s x_{nst} \leq H_n^{CW}(1 + H^{DEV}) \qquad n \in \mathcal{N} \tag{5.9}$$

$$\sum_{s \in \mathcal{S}^W} \sum_{\tau = t-6}^{t} H_s x_{ns\tau} \leq \overline{H}_n^{7D} \qquad n \in \mathcal{N}, t \in \mathcal{T} \tag{5.10}$$

Constraints (5.9) ensure that the work load in hours of each employee is kept within the minimum and maximum number of hours, as stated in the work contracts. Constraints (5.10) make sure that employees do not work more than $\overline{H}_n^{7D}$ hours in any given 7-day period.

5.4.4 Required Rest

$$x_{ns_1(t-2)} + x_{n'F1'(t-1)} + \sum_{s_2 \in S | \beta^{F1}_{s_1 s_2} = 0} x_{ns_2 t} \leq 2 \qquad n \in \mathcal{N}, s_1 \in \mathcal{S}, t \in \mathcal{T} \tag{5.11}$$

$$\sum_{t \in \mathcal{T}_k} x_{n'F1't} = 1 \qquad n \in \mathcal{N}, k \in \mathcal{K} \tag{5.12}$$

Constraints (5.11) and (5.12) regard the weekly protected off day, and ensure that the day is assigned once a week and that sufficient rest between the shifts before and after the day is provided.

5.4.5 Shift Patterns

$$\sum_{d=1}^{L_p} x_{nS_{dp}(t - L_p^S + d)} - L_p w_{ntp} \geq 0 \qquad n \in \mathcal{N}, t \in \mathcal{T}^{SUN}, p \in \mathcal{P}^{DW} \tag{5.13}$$

$$\sum_{d=1}^{L_p} x_{nS_{dp}(t - L_p + d)} - \prod_{d=1}^{L_p} \beta^R_{nS_{dp}(t - L_p + d)} \leq L_p - 1 \qquad n \in \mathcal{N}, t \in \mathcal{T}, p \in \mathcal{P}^{GEN} \tag{5.14}$$

$$\sum_{d=1}^{L_p} x_{nS_{dp}(t-L_p+d)} \leq L_p - 1 \qquad n \in \mathcal{N}, t \in \mathcal{T}, p \in \mathcal{P}^{ILL} \tag{5.15}$$

Some shift patterns are very popular, and should be assigned whenever possible. This is indicated by constraints (5.13), which assign a value of 1 to w_{ntp} when such a pattern is assigned. Two examples of desirable patterns, indicating preferable short and long shift patterns during work weekends, are given in Table 5.2.

Table 5.2: Examples of desirable shift patterns containing t, where $t \in \mathcal{T}^{SUN}$

Pattern	$t-3$	$t-2$	$t-1$	t	$t+1$	L_p
$p=1$	-	F1	E	D	-	3
$p=2$	F1	E	D	E	D	5

Constraints (5.14) keep track of whether an employee works an undesirable shift pattern that ends on day t. These shift patterns are generally not assigned, unless an employee explicitly requests to work one of these patterns. Then the product of β^R_{nt} for these days is 1, and the pattern may be scheduled for that employee on the specific days requested. Two examples of such patterns are described in Table 5.3. The first pattern in the table indicates that the employees generally do not fancy working four Night shifts in a row. The second one states that it is not desirable to start working a Night shift on an off day and then go straight to a Day shift, as this results in an irregular circadian rhythm.

Table 5.3: Examples of undesirable shift patterns ending on day t, where $t \in \mathcal{T}$

Pattern	$t-3$	$t-2$	$t-1$	t	L_p
$p=3$	N	N	N	N	4
$p=4$	-	F	N	D	3

Some of the undesirable shift patterns included in constraints (5.14) only occur during weekends, and are defined such that the final shift in the patterns falls on a Sunday. Table 5.4 illustrates two such patterns. Pattern five is undesirable because most employees consider working the Friday Evening shift as equal to working a weekend shift. The sixth pattern indicates that employees prefer not working Evening shifts exclusively during work weekends.

Table 5.4: Examples of undesirable shift patterns ending on day t, where $t \in \mathcal{T}^{SUN}$

Pattern	$t-2$	$t-1$	t	L_p
$p=5$	E	F	-	2
$p=6$	E	E	E	3

Constraints (5.15) state that some patterns of successive shifts should never be assigned. Table 5.5 indicates two such patterns. Pattern number seven indicates that it is illegal to work a Day or Evening shift the day before a Night shift, as this would result in insufficient rest between shifts. Pattern number eight only occurs during weekends, where day t is a Sunday, and states that employees should never work the Friday Night shift if they work Day or Evening shifts that weekend.

Table 5.5: Examples of illegal shift patterns ending on day t, where $t \in \mathcal{T}$

Pattern	$t-2$	$t-1$	t	L_p
$p=7$	-	D/E	N	2
$p=8$	N	D/E	-	2

5.4.6 Other Scheduling Requirements

$$\sum_{s \in \mathcal{S}^W} \sum_{\tau=t-\overline{M}^{CW}}^{t} x_{ns\tau} \leq \overline{M}^{CW} \qquad n \in \mathcal{N}, t \in \mathcal{T} \tag{5.16}$$

$$\sum_{\tau=t-\overline{M}^{N}}^{t} x_{nN\tau} \leq \overline{M}^{N} \qquad n \in \mathcal{N}, t \in \mathcal{T} \tag{5.17}$$

$$\sum_{t \in \mathcal{T}} (x_{nEt} + Bx_{nNt}) \geq U \sum_{s \in \mathcal{S}^W} \sum_{t \in \mathcal{T}} x_{nst} \qquad n \in \mathcal{N}^{GEN} \tag{5.18}$$

Constraints (5.16) and (5.17) ensure that the limits on how many consecutive shifts and consecutive nights, respectively, an employee is allowed to work are respected. Constraints (5.18) state that a minimum of each employee's shifts must be an Evening or a Night shift. The set $\mathcal{N}^{GEN}$ includes all employees which this rule applies for.

5.4.7 Variable Declarations and Fixations

$$x_{nst} \in \{0,1\} \qquad n \in \mathcal{N}, s \in \mathcal{S}, t \in \mathcal{T} \tag{5.19}$$

$$y_{cst} \in \mathbb{N}_0 \qquad c \in \mathcal{C}, s \in \mathcal{S}^W, t \in \mathcal{T} \tag{5.20}$$

$$w_{ntp} \in \{0,1\} \qquad n \in \mathcal{N}, t \in \mathcal{T}^{SUN}, p \in \mathcal{P}^{DW} \tag{5.21}$$

$$x_{nst} = 1 \qquad n \in \mathcal{N}, s \in \mathcal{S}, t \in \mathcal{T}^{\mathcal{P}} \mid \beta^{PA}_{nst} = 1 \tag{5.22}$$

$$x_{nst} = 0 \qquad n \in \mathcal{N}, s \in \mathcal{S}, t \in \mathcal{T} \mid \beta^{NA}_{nst} = 1 \tag{5.23}$$

Constraints (5.19) and (5.21) enforce binary constraints on variables x_{nst} and w_{ntp}, respectively, while constraints (5.20) put the integer variables $y_{sct} \geq 0$. Constraints (5.22) state that all shifts indicated by the parameters β^{PA}_{nst} are mandatory to assign. This includes the shifts scheduled in the previous planning period, and the shifts that cover the special functions that some of the employees have taken on. Similarly, (5.23) make sure that the shifts indicated by the parameters β^{NA}_{nst} are never assigned. This includes the regular off shift on Sundays. During implementation, this constraint can be enforced by never generating these variables.

5.5 Scheduling Example

In Table 5.6 we provide a 9-day example schedule for 9 nurses, where we assume that all nurses have the same skills and that there are no new hires. In our example, the minimum demand $\underline{D}_{st}$ is 1 per shift per day, with an upper limit $\overline{D}_{st}$ of 2 on the weekday Day and Evening shifts, and 1 on the remaining shifts (constraints (5.3)).

Table 5.6: An example schedule for 9 nurses over 9 days, where D, E, N, F and F1 denote Day, Evening, Night, Off and the mandatory weekly off day, respectively

Day Index	Sun 0	Mon 1	Tue 2	Wed 3	Thu 4	Fri 5	Sat 6	Sun 7	Mon 8
Nurse 1	F1	D	D	E	F	F1	E	D	F
Nurse 2	F1	E	E	F	F1	E	D	E	D
Nurse 3	F1	E	D	E	D	F	F	F1	D
Nurse 4	F1	F	E	D	E	E	F	F1	F
Nurse 5	F1	F	F	F	F1	F	N	N	N
Nurse 6	F1	F	N	N	F	D	F	F1	E
Nurse 7	E	D	F	F	N	N	F	F1	F
Nurse 8	D	F	F	D	E	F	F	F1	F
Nurse 9	N	N	F	F	D	D	F	F1	E
#D	1	2	2	2	2	2	1	1	2
#E	1	2	2	2	2	2	1	1	2
#N	1	1	1	1	1	1	1	1	1
#F	6	4	4	4	4	4	6	6	4
$\underline{D}_{st}$	1	1	1	1	1	1	1	1	1

The employees are assigned one shift per day, as enforced by constraints (5.2). Not all employees have the same number of work shifts assigned; either because they have different work contracts, or because of the averaging of work hours over the planning period explained in Section 4.1. Both cases are kept track of by constraints (5.9).

Constraints (5.6) and (5.7) state that employees should work every M^{NW} weekend, where $M^{NW} = 3$ in this example. Nurse 7, 8 and 9 work the first, Nurse 1, 2 and 5 work the second, while the remaining nurses should work the third.

Nurse 1 and 2 work Day or Evening shifts on day 6 and 7, and have been assigned the desirable patterns indicated in Table 5.2. Nurse 5 works Night shifts during her work weekend, and has been assigned the mandatory pattern of three consecutive Night shifts on Saturday, Sunday and Monday, enforced by constraints (5.8).

All employees have been assigned one mandatory off day each week, which is placed on Sunday when the employees have an off weekend (constraints (5.11) and (5.12)).

5.6 Proactive Model Extensions

In this section, we present four extensions to the scheduling model; one for each managerial insight connected to proactive strategies in Table 4.2. The strategies and the related insights are summarized in Table 5.7. The strategies are intended to increase the overall robustness without dedicating too much attention to specific skills and shift types. The only exceptions are the Buffer extension, which also considers skills to some degree, and the Ghost extension, which focuses on Night shifts in particular.

Table 5.7: Proactive scheduling strategies and related managerial insights

Name	Insight	Strategy
Buffer	Optimal assignment of surplus work hours	Ensure even buffer of employees in excess of minimum demand on all shifts in objective function
Ghost	Optimal placement of off shifts	Schedule off shifts such that employees can take on extra Nights shifts without breaking governmental rules and incurring overtime
Absence	Value of utilizing historical absence information	Schedule additional employees on shifts where employees who are vulnerable to absence are assigned
Extra weekends	Value of additional staffing during weekends	Increase staffing during weekend shifts by allowing employees to work extra weekends in exchange for additional off days

All the model extensions are based on the scheduling model already presented in this chapter. We therefore do not present full models in this section, only the elements that deviate from the original model.

5.6.1 Extension 1: Buffer

At DNIC, the nurses have more contracted work hours than what is required to cover the minimum demand during weekdays. Assigning these surplus work hours evenly over all weekday shifts as a buffer may increase the overall robustness. As the problem is very constrained during weekends, the model extension is only applied for weekdays.

We seek to achieve buffers of approximately equal size on each shift in order to obtain a lower variability in the number of employees per shift of the same type. To achieve this, we use a diminishing marginal reward per additional employee in the buffer. Specially ordered sets of type 1 (SOS1), which are sets of binary variables where at most one variable may take on a non-negative value, are used to formulate the diminishing marginal reward.

Definitions

Both a new index, set and several parameters are included in the model extension.

r	index of variable in specially ordered set of type 1
F	maximum buffer size rewarded in objective function
$\mathcal{R}$	set of integers, $\mathcal{R} = \{1, \dots, F\}$
W_r^B	reward of buffer of size r in excess of the minimum demand
W_r^{SB}	reward of buffer of size r in excess of the minimum demand for a certain skill

Specially ordered sets of type 1 (SOS1), each with its corresponding binary variables, are used to keep track of the buffers and corresponding reward.

$$\lambda_{rst}^{B} = \begin{cases} 1 & \text{if there are } r \text{ employees in excess of the minimum demand on shift } s \text{ on day } t \\ 0 & \text{otherwise} \end{cases}$$

$$\lambda_{rcst}^{SB} = \begin{cases} 1 & \text{if there are } r \text{ employees with skill } c \text{ in excess of the minimum demand for that skill on shift } s \text{ on day } t \\ 0 & \text{otherwise} \end{cases}$$

Objective and Constraints

We add the terms in (5.24) to the objective function. The terms reward having buffers of employees in excess of both the overall and skill minimum demand on each shift.

$$z^D = \sum_{r\in\mathcal{R}} \sum_{s\in\mathcal{S}^W} \sum_{t\in\mathcal{T}} \left(W_r^B \lambda_{rst}^B + W_r^{SB} \sum_{c\in\mathcal{C}} \lambda_{rcst}^{SB} \right) \tag{5.24}$$

Constraints (5.26) replace constraints (5.5), while constraints (**??**) and (5.27) to (5.30) are added to the model.

$$\sum_{n\in\mathcal{N}} \beta_{nt}^N x_{nst} - \sum_{r\in\mathcal{R}} r\lambda_{rst}^B \geq \underline{D}_{st} \qquad s\in\mathcal{S}^W, t\in\mathcal{T} \tag{5.25}$$

$$\sum_{n\in\mathcal{N}_c} \beta_{nt}^N x_{nst} + y_{cst} - \sum_{r\in\mathcal{R}} r\lambda_{rcst}^{SB} \geq \underline{D}_{cst}^C \qquad c\in\mathcal{C}, s\in\mathcal{S}^W, t\in\mathcal{T} \tag{5.26}$$

$$\sum_{r\in\mathcal{R}} \lambda_{rst}^B \leq 1 \qquad s\in\mathcal{S}^W, t\in\mathcal{T} \tag{5.27}$$

$$\sum_{r\in\mathcal{R}} \lambda_{rcst}^{SB} \leq 1 \qquad c\in\mathcal{C}, s\in\mathcal{S}^W, t\in\mathcal{T} \tag{5.28}$$

$$\lambda_{rst}^B \in \{0,1\}, SOS1 \qquad r\in\mathcal{R}, s\in\mathcal{S}^W, t\in\mathcal{T} \tag{5.29}$$

$$\lambda_{rcst}^{SB} \in \{0,1\}, SOS1 \qquad r\in\mathcal{R}, c\in\mathcal{C}, s\in\mathcal{S}^W, t\in\mathcal{T} \tag{5.30}$$

The SOS1-variables may take on positive values if a buffer is assigned in constraints (5.25) and (5.26). Constraints (5.27) and (5.28) ensure that at most one SOS1-variable per shift and day may take on a non-negative value, while (5.29) and (5.30) state that the new variables are binary and included in specially ordered sets of type 1.

5.6.2 Extension 2: Ghost

When an employee has a ghost shift of a certain type the employee has an off day, but is able to take on an unplanned shift of that type during the online operational phase without breaking any of the rules regarding consecutive shifts and overtime work. Because Night shifts are the shifts which are the most constrained when it comes to covering demand, we introduce a model extension which assigns ghost shifts of type Night, denoted GN.

Definitions

One set and two parameters are added to the model. Furthermore, some sets and parameters are adapted to include information about the ghost shift.

$\mathcal{S}^G$	set of ghost shifts, $\mathcal{S}^G = \{GN\}$, $\mathcal{S}^G \subset \mathcal{S}$
H	Hours in full-time work week
W^G	Reward for assigning a ghost shift

In constraints (5.15), the patterns including a ghost shift which are illegal to assign are the same patterns as the patterns illegal to assign which include Night shifts. We also extend the set of undesirable patterns used in constraints (5.14) to include patterns including certain combinations of Night shifts, ghost Night shifts and off days.

Parameters $\beta^{F1}_{s_1 s_2}$, which are used to enforce the rules regarding rest around the mandatory off day in constraints (5.11), must be adapted to also include the ghost Night shift. As it is desirable to maintain the rules of rest around the mandatory off day, the same rules regard the ghost shift and the Night shift. The duration of the GN shift is the same as the duration of the Night shift, indicated by parameter H_N.

A binary variable is added to the model to enable the modelling of the ghost shifts. Furthermore, GN is now one of the possible shift types possible to assign with variables x_{nst}. However, in practice, being assigned a ghost shift will appear as the same as being assigned an off shift to the employees.

$$g^H_{nk} = \begin{cases} 1 & \text{if employee } n \text{ is assigned less than } H \text{ hours of work in week } k \\ 0 & \text{otherwise} \end{cases}$$

Objective and Constraints

The term (5.31), which maximizes the number of ghost shifts assigned, is added to the objective function.

$$z^G = W^G \sum_{n \in \mathcal{N}} \sum_{s \in \mathcal{S}^G} \sum_{t \in \mathcal{T}} x_{nst} \tag{5.31}$$

Constraints (5.32) to (5.36) are introduced.

$$\sum_{s \in \mathcal{S}^W \cup \mathcal{S}^G} \sum_{\tau = t - \overline{M}^{CW}}^{t} x_{ns\tau} \leq \overline{M}^{CW} \qquad n \in \mathcal{N}, t \in \mathcal{T} \tag{5.32}$$

$$\sum_{\tau = t - \overline{M}^{N}}^{t} (x_{nN\tau} + x_{n'GN'\tau}) \leq \overline{M}^{N} \qquad n \in \mathcal{N}, t \in \mathcal{T} \tag{5.33}$$

$$\sum_{s \in \mathcal{S}^W \cup \mathcal{S}^G} \sum_{t \in \mathcal{T}_k} H_s x_{nst} \leq H + (\overline{H}_n^{7D} - H)(1 - g_{nk}^H) \qquad n \in \mathcal{N}, k \in \mathcal{K} \tag{5.34}$$

$$0 \leq \sum_{s \in \mathcal{S}^G} \sum_{t \in \mathcal{T}_k} H_s x_{nst} \leq H g_{nk}^H \qquad n \in \mathcal{N}, k \in \mathcal{K} \tag{5.35}$$

$$g_{nk}^H \in \{0, 1\} \qquad n \in \mathcal{N}, k \in \mathcal{K} \tag{5.36}$$

Constraints (5.32) and (5.33) make sure that the rules regarding the maximum number of consecutive shifts and nights are not violated, even if a ghost shift is turned into a work shift. These constraints replace constraints (5.16) and (5.17) in the original model. Constraints (5.34) put an upper limit on the maximum number of hours worked during a 7-day period, and is similar to constraints (5.10). The number of ghost shifts worked without incurring overtime pay is limited by constraints (5.35). Finally, constraints (5.36) enforce the binary constraints on variables g_{nk}^H.

5.6.3 Extension 3: Absence

In Chapter 6, we analyze historical data of employee absence at DNIC. Some of the employees are identified as being more vulnerable to absence than others. For both DNIC and these employees, it can be desirable to add buffers of employees to shifts where they are assigned. This way, the load on the most vulnerable employees is eased, and in turn DNIC is less sensitive towards absence in the online operational phase. During weekends, the shift assignment is too constrained to add buffers. Consequently, we choose to evenly distribute the employees in the high-risk group during the weekend shift as an alternative approach.

Definitions

We add the following sets and parameters to the model.

$\mathcal{N}_t^{HR}$	set of nurses with high risk of being absent on day t
$\mathcal{T}^W$	set of Saturdays and Sundays
C_{nt}	fraction of how much employee n contributes to covering demand on day t
W^{HR}	penalty per additional employee with high risk of being absent that is assigned the same shift during weekends
S^{HR}	number of employees in high-risk group who can be assigned the same shift without getting penalized in the objective function

Recall from Section 5.2.3 that β_{nt}^N indicate whether employee n can contribute to covering demand on day t. Now, this parameter is included in C_{nt} in addition to a factor reducing the contribution from employees in the high-risk group.

The following slack variables are included in the scheduling model.

s_{st}^{HR} = number of nurses with high risk of being absent exceeding S^{HR} that are assigned to shift s on day t

Objective and Constraints

The term (5.37), which minimizes the number of employees with high risk of being absent who are scheduled on the same shift, is appended to the objective function.

$$z^{HR} = -W^{HR} \sum_{s \in \mathcal{S}^W} \sum_{t \in \mathcal{T}} s_{st}^{HR} \tag{5.37}$$

Constraints (5.38) replace (5.3), while constraints (5.39) and (5.40) are added to the model.

$$\underline{D}_{st} \leq \sum_{n \in \mathcal{N}} C_{nt} x_{nst} \leq \overline{D}_{st} \qquad s \in \mathcal{S}^W, t \in \mathcal{T} \tag{5.38}$$

$$\sum_{n \in \mathcal{N}_t^{HR}} x_{nst} - s_{st}^{HR} \leq S^{HR} \qquad s \in \mathcal{S}^W, t \in \mathcal{T}^W \qquad (5.39)$$

$$s_{st}^{HR} \in \mathbb{N}_0 \qquad s \in \mathcal{S}^W, t \in \mathcal{T}^W \qquad (5.40)$$

Constraints (5.38) make sure that the number of employees assigned work shifts is between the minimum and maximum demand on all shifts, while adding buffers to the shifts where employees in the high-risk group are assigned. Variables s_{st}^{HR} take on a value if more than S^{HR} employees in the high-risk group are assigned to the same shift, as ensured by constraints (5.39). Constraints (5.40) declare s_{st}^{HR} as integer variables.

5.6.4 Extension 4: Extra Weekends

In Section 2.3 we stated that weekends are the biggest bottleneck in the planning process due to the strict requirements as to how frequently each employee is allowed to work during weekends. Our hypothesis is that allowing the employees to work more weekends in turn will make the schedules more robust. In order for the employees to accept the policy change, we allow trading extra weekend work with extra off days arbitrarily assigned some other time during the planning period.

The method of trading extra weekend work with extra off days in order to make more robust schedules was also discussed by Klyve and Beckmann (2016), but they did not have the means to test whether the hypothesis is correct.

Definitions

Several parameters are added to the model.

$\overline{M}^{EW}$	maximum number of extra weekends an employee can work
H^{EO}	number of extra off hours gained by working an extra weekend
β_n^{EW}	1 if employee n can work more weekends than normally contracted, 0 otherwise

The binary variables e_{nt} must be included.

$$e_{nt} = \begin{cases} 1 & \text{if employee } n \text{ is works an extra weekend containing Sunday } t \\ 0 & \text{otherwise} \end{cases}$$

Objective and Constraints

The objective function remains unchanged. Constraints (5.43) replace constraints (5.7), while constraints (5.45) and (5.46) replace (5.9).

$$\sum_{t \in \mathcal{T}^{SUN}} e_{nt} \leq \overline{M}^{EW} \beta_n^{EW} \qquad n \in \mathcal{N} \tag{5.41}$$

$$\sum_{s \in \mathcal{S}^W} (x_{ns(t-1)} + x_{nst}) \geq 2\beta_n^{EW} e_{nt} \qquad n \in \mathcal{N}, t \in \mathcal{T}^{SUN} \tag{5.42}$$

$$\sum_{s \in \mathcal{S}^W} \sum_{\tau=0}^{M^{NW}-1} x_{ns(t-7\tau)} - \sum_{\tau=0}^{M^{NW}-1} e_{n(t-7\tau)} \leq 1 \qquad n \in \mathcal{N}, t \in \mathcal{T}^{SUN} \tag{5.43}$$

$$\sum_{\tau=0}^{M^{NW}-1} e_{n(t-7\tau)} \leq \beta_n^{EW} \qquad n \in \mathcal{N}, t \in \mathcal{T}^{SUN} \tag{5.44}$$

$$\sum_{s \in \mathcal{S}^W} \sum_{t \in \mathcal{T}} H_s x_{nst} + \sum_{t \in \mathcal{T}^{SUN}} H^{EO} e_{nt} \geq H_n^{CW}(1 - H^{DEV}) \qquad n \in \mathcal{N} \tag{5.45}$$

$$\sum_{s \in \mathcal{S}^W} \sum_{t \in \mathcal{T}} H_s x_{nst} + \sum_{t \in \mathcal{T}^{SUN}} H^{EO} e_{nt} \leq H_n^{CW}(1 + H^{DEV}) \qquad n \in \mathcal{N} \tag{5.46}$$

$$e_{nt} \in \{0, 1\} \qquad n \in \mathcal{N}, t \in \mathcal{T}^{SUN} \tag{5.47}$$

Constraints (5.41) limit the number of extra work weekends each employee may work, while constraints (5.42) state that employees working an extra weekend have to work both days of the weekend. Constraints (5.43) connect variables x_{nst} and e_{nt} and ensure that employees work every M^{CW} weekend, unless they are assigned extra weekends. Employees working extra weekends should at most be allowed to work two consecutive weekends, which is ensured by constraints (5.44). The employees are provided H^{EO} extra off hours per extra weekend worked, helping them to both exceed the minimum number of work hours and reach the maximum number of work hours quicker, indicated by constraints (5.45) and (5.46). Constraints (5.47) are the binary constraints for variables e_{nt}.

Chapter 6

Data Analysis and Uncertainty Modelling

This chapter contains a description of how the uncertainties in the demand and supply of employees are modelled. We start by explaining the general assumptions on which the uncertainty modelling is based in Section 6.1. We then present the demand and absence uncertainty models in Sections 6.2 and 6.3, respectively. Both models build on historical data provided by the Department of Neonatal Intensive Care (DNIC). We elaborate on the availability of non-absent employees in Section 6.4. Finally, in Section 6.5, we present an example of how uncertainty realization can affect the schedule from Section 5.5.

The result of the chapter is a description of several simulation models used to simulate the demand and supply of employees at DNIC. The output of the models is used as input to the rescheduling model in Chapter 7.

6.1 General Assumptions and Simplifications

The modelling of uncertainty in this thesis is based on the two assumptions in Table 6.1. First, we assume that the employee absences faced at DNIC do not depend on the number of patients and their conditions. Second, the treatment of each patient is assumed independent of the number of employees at work and their weekly workload. Assuming independence between demand and employee absences implies that the uncertainty model can be divided into two components; one for absences and one for patients.

Table 6.1: General assumptions of the uncertainty models

Assumption 1	Frequency and duration of absences independent of demand
Assumption 2	Condition of patients during their stay independent of workload per nurse

Investigating any dependencies between patient demand and employee absences is outside the scope of this thesis, but we do consider this an interesting case for future research. Relevant literature indicate that there is indeed a connection between the two. For example, Rauhala et al. (2007) found a significant connection between patient-associated work overload and increased sickness absence among nurses. It is not unlikely that in practice, there is also a connection between these aspects at DNIC.

6.2 Demand Uncertainty Model

The demand uncertainty model, also denoted the demand model, is used to simulate the real demand at DNIC as accurately as possible, using the historical patient data provided by the department.

6.2.1 Available Patient Data

The data set provided by DNIC contains the daily number of patients at each level in the period August 17th 2014 to October 27th 2016, which gives a total of 803 days.

The data records are aggregated descriptions, with no information regarding individual patients due to privacy concerns. This means that there is no information regarding new admissions and discharges of patients, nor about the patients' transitions between levels. This limits the data analysis in the sense that it is challenging to calculate realistic probabilities of admissions and the duration of a typical patient stay on each level.

6.2.2 Assumptions and Simplifications of the Demand Model

The underlying assumptions and simplifications of the demand uncertainty model are outlined in Table 6.2.

Table 6.2: The assumptions of the demand uncertainty model

Assumption 1	No trends or seasonal variations
Assumption 2	Patients per level measured once per day

First, we assume that there are no trends or seasonal variations in the data set. However, in reality there are considerable seasonal variations in the number of newborns per month in Norway, with a peak in June and July and a low in December (Andersen, 2018). This is likely to affect the number of patients admitted to DNIC as well, indicating that there are indeed seasonal variations in the number of patients. This is in line with the impression of the scheduling manager at DNIC, who has stated that seasonal variations are not uncommon. However, as the main scope of the thesis is to study robust schedules in general, but not on a level of detail corresponding to e.g. different staffing strategies per month of the year, we consider the assumption a reasonable simplification.

The second assumption is that the number of patients change once daily. This is a simplification, as in reality, patients may be admitted to DNIC at any time of the day.

6.2.3 Objectives of the Demand Model

Fulfilling the objectives in Table 6.3, we aim to accurately model the daily demand at DNIC.

Table 6.3: The objectives of the demand uncertainty model

Objective 1	Simulate number of patients on each level
Objective 2	Daily variation consistent with historical variation
Objective 3	Capture dependencies between levels
Objective 4	Keep number of patients within upper capacity limit

First, the model should simulate the number of patients and their corresponding conditions, which the demand for nurses at DNIC dependents upon. Because the demand is not dependent on the individual patients themselves, we consider modelling the number of patients at each level each day as an equally relevant possibility as modelling individual patients and their stay at each level.

When simulating the aggregated number of patients, the model should meet some requirements regarding the correct day-to-day behavior, keeping the daily variation consistent with the historical variation.

Finally, the model should capture any dependencies between levels, as well as reflect that DNIC has an upper capacity limit.

6.2.4 General Demand Model

We introduce a first order Markov model for each hospital bed at DNIC. In the model, we assume that there is a given number of independent hospital beds. The state of each bed can be considered as the condition of the patient that occupies that bed. The state is thus between 1 and 5 if the bed is occupied, corresponding to the five levels of severity explained in Section 2.3. An unoccupied bed is represented by state 0.

The transition process is illustrated in Figure 6.1. Transitions between all states are allowed. A transition from unoccupied to either one of the other states is considered an admission, while a transition the other way around is considered a discharge.

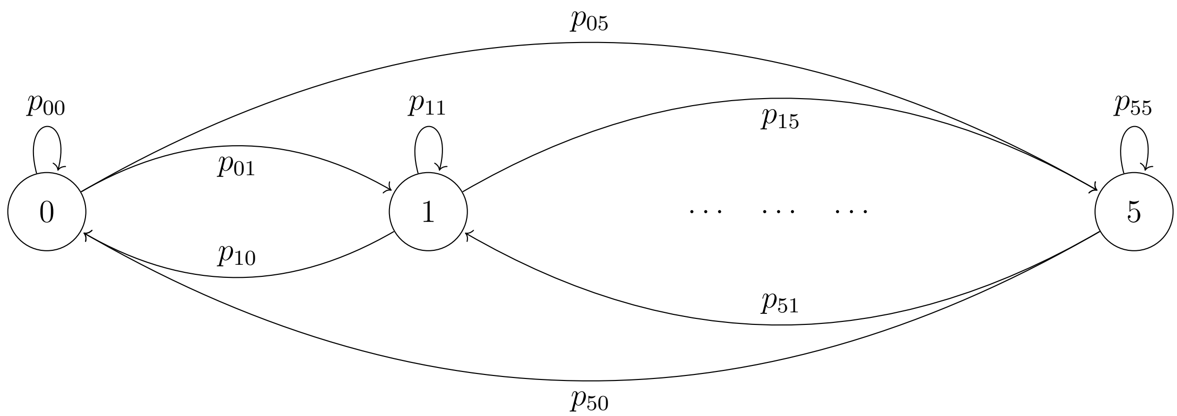

Figure 6.1: The transitions between states 0-5 for each hospital bed, where the probabilities of going from one state to the next is indicated on each arc

A benefit of the model displayed in Figure 6.1 is that it captures the dependencies between all levels. The model also smoothly includes an upper limit on the total number of patients admitted by not including more hospital beds than the upper capacity limit at DNIC.

The model is a slight simplification of the real-world situation, as the number of patients with severe conditions that DNIC can take in may be dependent on the patient mix already admitted, in addition to some beds being reserved specific severity levels.

6.2.5 Demand Transition Probability Distributions

Calculating transition probabilities when individuals transition over a time period is not a trivial task when the only data available is an aggregated description, such as in the patient data set. Jones (2005) suggests a quadratic programming approach for estimating transition probabilities when individual transitions are unknown.

Using this approach, we introduce a mathematical model, where we define the patient proportion at a given level as the number of patients at that level divided by the total number of hospital beds. This patient proportion should, at a given level and time increment, equal the sum of patient proportions at all levels in the previous time period multiplied by their respective transition probability of going to the given level. However, the relationship between the patient proportions from one day to the next is not exact, and an error term must be added as an approximation. The mathematical description of the relationship between the previous and current time period is given by Equation (6.1), where U_{tj} is an input to the model representing the patient proportion at level j on day t, p_{ij} is the probability of a transition from level i to level j, and e_{tj} is the error term.

$$U_{tj} = \sum_{i \in \mathcal{I}} U_{(t-1)j} p_{ij} + e_{tj} \tag{6.1}$$

Equation (6.1) can be reformulated and substituted into the objective function (6.2), which minimizes the sum of squared errors for all time periods and levels. Consequently, we seek to find the transition probabilities that make the relationship in (6.1) as exact as possible. The objective function is subject to constraints (6.3) and (6.4), where (6.3) make sure that the transition probabilities from a given state and to all states must sum to one, while (6.4) ensure that all transition probabilities are positive.

$$\min \sum_{t \in \mathcal{T}} \sum_{j \in \mathcal{I}} \left(U_{tj} - \sum_{i \in \mathcal{I}} U_{(t-1)i} p_{ij} \right)^2 \tag{6.2}$$

$$\sum_{j \in \mathcal{I}} p_{ij} = 1 \qquad i \in \mathcal{I} \tag{6.3}$$

$$p_{ij} \geq 0 \qquad i, j \in \mathcal{I} \tag{6.4}$$

A weakness of this method is that it is very sensitive to outliers and irregular behavior in the data records. If there for some reason were many irregularly high or low numbers of patients at a level during this time period, the probabilities would be affected by this. However, as there are approximately 800 days with data points in the data set, we believe that there are sufficient records to keep potential outliers from having a too large effect on the results.

6.2.6 Demand Simulation Algorithm

Algorithm 6.1 displays the demand simulation procedure. The number of hospital beds in the system is given by $|Beds|$. For each iteration, the state of each bed in the next iteration is updated based on a draw from the cumulative distribution of each bed's current state. The cumulative distributions are based on the the transition probabilities calculated in Section 6.2.5.

A warm up period is necessary to obtain a realistic distribution of states prior to the actual simulation. For simplicity, the warm up period is not explicitly described in the algorithm, but rather incorporated in the initialization on line 2.

Algorithm 6.1 Algorithm for the demand uncertainty model

procedure DEMANDSIMULATION
  $\mathbf{b_0} \leftarrow$ INITIALIZESTATE$(1 \dots |Beds|)$
  $t \leftarrow 0$
  while $t <$ TimeLimit **do**
    for all b in Beds **do**
      $r \leftarrow$ DRAWRANDOMNUMBER
      $b^b_{t+1} \leftarrow$ CUMULATIVEDEMANDLOOKUP(b^b_t, r)
    end for
    $t \leftarrow t + 1$
  end while
end procedure

6.3 Absence Uncertainty Model

The overall objective of the absence uncertainty model, sometimes denoted the absence model, is to obtain a model which realistically simulates the employee absences occurring at DNIC.

6.3.1 Available Employee Data

The data set provided by DNIC contains detailed anonymous records of planned work shifts and employee absences for each day in the period January 1st 2015 to December 31st 2017. The date of each employee's first day at work, and, if they quit the job, their resignation date, has also been provided.

Each row in the data set contains information regarding a specific employee's schedule for a given day, and contains the employee number, date and shift code (Day, Evening, or Night). There is also information regarding whether the shift was planned in the initial schedule, if it was an extra shift, or whether the shift has been swapped. When a shift was planned in advance, special codes indicate whether the employee was absent and the cause of this absence. This means that we can tell whether an absence was short- or long-term, as described in Section 2.3.5, based on this code. Examples of such codes are given in Table 6.4.

Table 6.4: Examples of absence codes

Code	Cause	Duration
SE	Self-certification of absence	Short
SB	Home with sick child	Short
SS	Sick leave	Long
OS	Maternity leave	Long

For each day, the employee data set only contains information about employees who had a work shift scheduled that day. The missing entries about the remaining employees entail that we do not know whether they were absent or not.

Data Cleansing

A total of 265 employees has worked one or more shifts at DNIC over the duration of the data set. The duration of employment for each employee varies from small temporary positions to continuous employment for the entire period. A significant portion of the employees in the data set had worked either only extra shifts, or worked short-time temporary positions during e.g. the summer holidays. With few and sometimes irregular records, it is difficult to obtain a picture of these employee's real probabilities of being absent. For instance, the employees who only worked extra shifts naturally did not have any records of absence in the data set, as they had no initially planned shifts to be absent from. Therefore, these employees were removed from the data set prior to the further analysis. After this removal, 165 employees were left in the set.

Missing Values

There are frequent missing values in the data set due to the observations being dependent on the schedule. The dependency of the schedule affects the number of observations because all employees must have at least one day off each week, no matter their number of weekly contracted hours, in addition to most of the employees in the data set working part-time positions.

The missing observations present a challenge in the data analysis, as they may introduce bias that can lead to false conclusions. In our case, the missing data complicates the analysis of both the frequency and the duration of absence. For example, assume an employee working a 40% position always shows up at her scheduled work shifts. The employee will then appear to always be non-absent when looking at the recorded data. However, working 40% amounts to an average of two work shifts per week, leaving a significant amount of unobserved values. Thus, it is not unlikely that the employee might have been sick on several occasions, without this being recorded. Also, if any absence had been recorded, it would have been impossible to know the true duration of the absence if it occurred right before or after an off day.

According to Lachin (2016), a common statistical approach when analyzing missing observations in longitudinal data, in which the dependent variable is measured at several points in time for each subject, is to use last observation carried forward (LOCF). Using LOCF, the missing value is imputed by the value recorded prior to the missing value. The data set consisting of both observed and imputed data is then analyzed as if the original data set was complete. However, LOCF may introduce significant bias in the data set, and therefore the method should not be used (Lachin, 2016).

Although LOCF is not recommended, there is one particular case where the method is a reasonable imputation approach. Long-term absence is relatively predictable in nature, and we consider it reasonable to assume that missing values occurring between two entries of the same long-term absence type can be imputed using LOCF. Imputing this long-term absence, we are left with 63521 observations in the data set, which is 45.4% of all possible entries. Obtaining an observation rate close to 100% would not be realistic even if all employees had a full-time position. As working full-time amounts to approximately five shifts a week, the maximum likely observation rate per employee is roughly $\frac{5}{7} \approx 70\%$.

Imputing the remaining missing values would imply making some assumptions about what the values of these entries are likely to be. These assumptions could further bias the resulting distributions, and then again the simulation results. It is far from certain that these results would be more accurate than the results derived from distributions based on the incomplete data set with missing values. We regard applying advanced statistical methods for imputation as outside the scope of this thesis. Because of this, we consider the data set after performing the pre-processing actions discussed above as sufficiently complete.

Figure 6.2 illustrates how the number of observations recorded in the data set is limited due to missing values on off days. The final size of the data set is determined by the number of observations where the value on the day prior to the observation was also recorded. This is because we are interested in how the employees transition between being available and being absent, and not just the frequency of each event, which is an important property in the model described in Section 6.3.4. Thus, the number of observed transitions is reduced with the number of transitions from an off day to a work day.

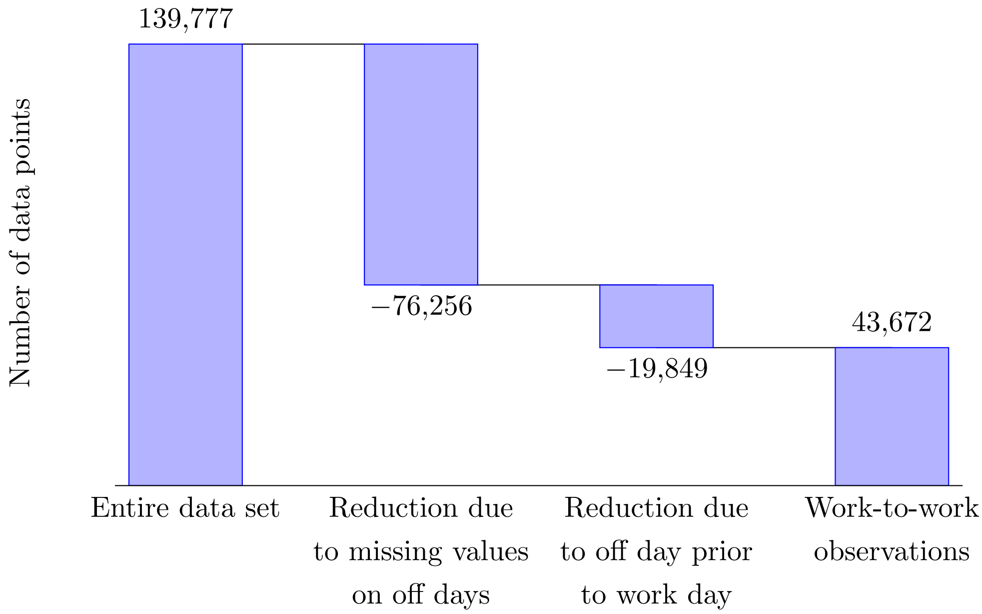

Figure 6.2: Key numbers for the employee data set

6.3.2 Assumptions and Simplifications of the Absence Model

The absence uncertainty model rests on three assumptions, as outlined in Table 6.5.

First, we assume that there are no trends in the data set, as these are very difficult to realistically identify due to the missing observations. We also assume that there are no seasonal variations for the same reason. The exception is seasonal variations within a

Table 6.5: The assumptions of the absence uncertainty model

Assumption 1	No trends in data set
Assumption 2	No seasonal variations, except for days of the week
Assumption 3	Employees are independent of each other
Assumption 4	Same probability of absence on work days and off days

week, which we further investigate. This means that we assume that the probabilities of absences are the same each month, but that they may vary from Monday to Sunday. In reality, it is likely that there are seasonal variations in the absences. For example, in Norway there are usually considerably more self-certifications of absence and sick leave certificates used during the winter than during the summer (SSB, 2018). Thus, these assumptions are likely to be simplifications of the actual situation.

The third assumption is that that the absence of each employee is independent of the absence occurring to the other employees. This way, it is not more or less likely that an employee is absent when other employees are absent. The assumption also implies that we can simulate the absence of each employee independent of the simulations for the other employees. Finally, because we have no records on off days, we assume that the employees have the same probabilities of being absent on both work days and off days.

6.3.3 Objectives of the Absence Model

To accurately simulate the real-life situation at DNIC, we introduce the objectives in Table 6.6 for the absence uncertainty model.

Table 6.6: The objectives of the absence uncertainty model

Objective 1	Simulate absence for individual employees
Objective 2	Simulate absence on both work and off days
Objective 3	Distinguish between short- and long-term absence

The first objective is to simulate whether particular employees are absent. It is essential to know which employees to remove from the schedule, as the employees are not a homogeneous group. The alternative would be to simulate the number of absent employees each day, but this would leave it impossible to tell which employees were absent. The absence model should therefore maintain information about the state of each employee. However, the transition probabilities between the states for each employee do not necessarily have to be unique. The resolution of the transition probability distributions between states is discussed in Section 6.3.5.

The second objective is to be able to know which employees can take on unplanned shifts. Information about the absence state on both work and off days is therefore required.

The third objective is to distinguish between short-term and long-term absence, suggesting that there should be be two absence states for each employee. It should also be possible to transition between being short- and long-term absent. In real-life this happens e.g. if a self-certification of absence is used during the first few days of absence (code SE in Table 6.4), before a sick-leave certificate from a doctor is provided.

6.3.4 General Absence Model

By incorporating the three objectives mentioned in Section 6.3.3, we obtain the absence uncertainty model. As illustrated in Figure 6.3, individual employees can transition between the states non-absent, short-term absent, and long-term absent, denoted a_N, a_S and a_L, respectively, on both off days and work days. All possible transitions between these states are allowed, each with its own transition probability.

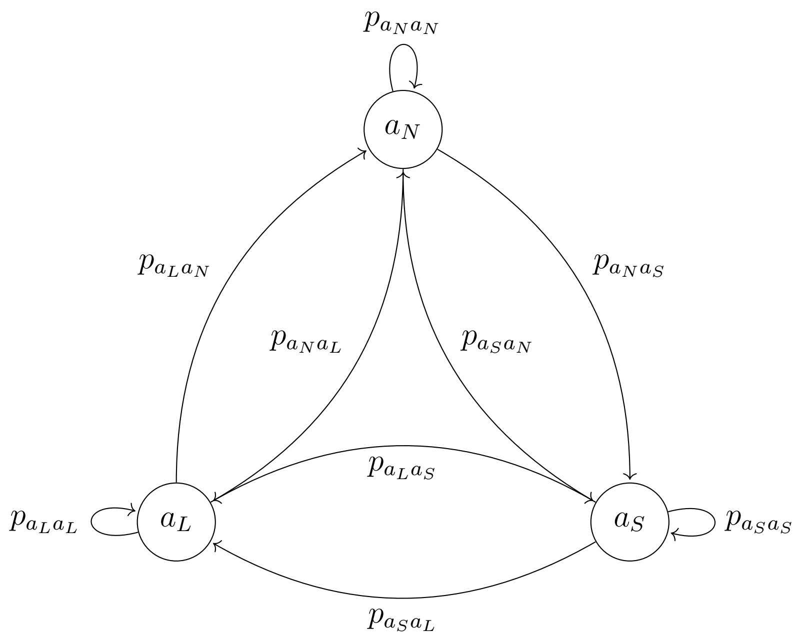

Figure 6.3: Transitions between the states non-absent (a_N), short-term absent (a_S) and long-term absent (a_L), as well as their respective transition probabilities

6.3.5 Absence Transition Probability Distributions

The transition probability distributions should be calculated using the available data. Several possible resolutions for the distributions exist, both concerning the level of aggregation of the employees and the aggregation of time. In general, the greatest benefit of calculating detailed distributions is that they then may be more realistic and capture possible special cases better. However, the more detailed the distribution, the higher are the requirements to the quality of the underlying data set. If the data set does not contain sufficiently many records, or if there is missing data, distributions with much detail can end up being biased.

Time Resolution

The time resolution of the transition probability distributions may vary from full aggregation, with the same probability of absences occurring at all times, through an aggregation of some days, such as one distribution for weekends and one for weekdays, to a resolution of one distribution per day of the week. There is also the possibility of having one distribution per shift type for all alternatives. As we assume that there are no seasonal variations in the data set, it does not make sense to consider e.g. different distributions per month.

We use hypothesis testing to evaluate which resolution is the best fit. Using a two-tailed, two-sample t-test, two data sets can be compared to conclude whether they have significantly different means. In this type of test, the null hypothesis is that the means of the two samples are equal, while the alternative hypothesis is that the means are different (Snedecor and William, 1989). For a more advanced analysis, methods such as neural nets, support vector machines and boosting can be used (Efron and Hastie, 2016).

We first conduct separate t-tests for short-term and long-term absences. Our observations are the rate of short-term absence and long-term absence for each day of the week. The rate of a short-term absence for a particular weekday is defined as the number of transitions into the short-term absent state, divided by the total number of observed transitions to that day. The rate of long-term absence is defined similarly.

A t-test for each combination of the days of the week can then be carried out. To avoid bias, the samples consist of the observations for employees with more than a certain number of shifts scheduled on each day of the week. The samples are unpaired and we assume unequal variances. We conduct all t-tests with a significance level equal to 5 %. The resulting p-values from all tests performed are provided in Appendix C.2.

As displayed in Table 6.7, for short-term absences there are no significant differences for Monday through Friday, or for Saturday and Sunday. As indicated, the null hypothesis is rejected for all combinations of a weekday and weekend day, with the only exception being Friday and Sunday. We also examined the differences between the various shift types on weekdays and weekends, with the result that the null hypothesis was accepted for all combinations.

Based on the results from the t-tests, we consider weekdays and weekends a sufficient distinction for short-term absence, where the weekdays are the days Monday through Friday, and the weekend days are Saturday and Sunday. Ideally, the null hypothesis stating that the mean of Friday is equal to the mean of Sunday should have been rejected. For simplicity, we include Friday in the weekday set as that seems to be the best fit.

Table 6.7: Results of a two-sample t-test for all combinations of days of the week for the short-term absence rate, where crosses or check marks indicate that the null hypothesis is rejected or accepted, respectively

	Tuesday	Wednesday	Thursday	Friday	Saturday	Sunday
Monday	✓	✓	✓	✓	✗	✗
Tuesday		✓	✓	✓	✗	✗
Wednesday			✓	✓	✗	✗
Thursday				✓	✗	✗
Friday					✗	✓
Saturday						✓

For long-term absences, there were no significant differences between any days of the week, nor for any combinations of shift types. We therefore assume that the same distribution of transitioning to long-term absence can be used for all days and shifts.

Employee Aggregation

The aggregation of employees may vary from no aggregation, with one unique probability distribution for each employee, to complete aggregation, where the same distribution applies to all employees. It is also possible to split the employees into groups, with different probabilities of being absent in each group.

In theory, one unique probability distribution for each employee can be generated with the data provided. This has the benefit of giving a direct connection between the employee in the schedule and the its historically observed absence. However, the probability of absence for a given employee may in practice end up being biased if the distribution is based on that employee's historical absence alone. For example, if an employee does not have any long-term absence registered in the historical data, the resulting distribution will not allow any long-term absence either.

A way to overcome this problem is to aggregate the employees into different groups depending on their historical risk of being absent. This method connects the historical data for each employee to that employee, while also allowing previously unobserved events such as long-term absences to happen. We therefore consider this method as less biased than the method with no aggregation.

Using the selected time resolution, we split the employees into one group with high risk of being absent and one group with low risk for both weekdays and weekends. The high-risk group is the set of employees with a historical absence rate larger than one standard deviation from the mean. The low-risk group contains the remaining employees. An employee does not have to be in the same risk group for both weekends and weekdays; it is possible to be in e.g. the high-risk group during weekends and low-risk during weekdays.

In Table 6.8, the number of observations remaining when the data set is split this way is indicated. Long-term transitions are excluded from the table, as we previously found that there is an equal possibility of entering a long-term absence on all days. The no aggregation alternative is included for comparison with the alternative of aggregating the employees into groups. Clearly, the former alternative presents a challenge with few observations per employee.

Table 6.8: Comparison of no aggregation and high- and low-risk aggregation when the data set is divided into weekdays and weekends

		Time resolution	
Aggregation	**Measure**	**Weekdays**	**Weekends**
None	Employees in each distribution	1	1
	Avg. observations per employee	138	47
	Std. observations per employee	90	26
High- and low-risk group	Employees in high-risk	28	27
	Employees in low-risk	137	138
	Observations in high-risk	3403	1170
	Observations in low-risk	19505	6568

Calculating the Transition Probability Distributions

As the transitions of individuals are recorded in the available data, calculating transition probabilities can be achieved using Equation (6.5), where we count the frequency of transitions from one state to another, and divide it by the total number of transitions from that particular state. In Equation (6.5), p_{ij}^{gw} denotes the probability that an employee in risk group g and day group w is in state j in the current time period, given that he or she was in state i in the previous time period. This probability is calculated using the number of times the employees in the same risk group for a given day of the week transitioned from state i to state j. This frequency is denoted n_{ij}^{ed}, where e is a given employee and the d index is included to make sure to count the transitions *to* a day in either the set of weekdays or set of weekend days.

$$p_{ij}^{gw} = \frac{\sum_{e \in \mathcal{E}_{gw}} \sum_{d \in \mathcal{D}_w} n_{ij}^{ed}}{\sum_{j \in \mathcal{J}} \sum_{e \in \mathcal{E}_{gw}} \sum_{d \in \mathcal{D}_w} n_{ij}^{ed}} \qquad i, j \in \mathcal{J}, g \in \mathcal{G}, w \in \mathcal{W} \tag{6.5}$$

Here, $\mathcal{J}$ is the set of states, and $\mathcal{G}$ is the set containing the risk groups, $\mathcal{W}$ is the set containing the day groups. Finally, $\mathcal{E}_{gw}$ is the set of employees in risk group g for day group w.

Recall Figure 6.3, which describes the possible transitions between the three states in the simulation model. With the selected resolution for time and employee aggregation, we are left with four such systems; weekday high- and low-risk, and weekend high- and low-risk. The probabilities p_{ij}^{gw} can then be substituted onto the corresponding arc for each system.

6.3.6 Absence Simulation Algorithm

The procedure for the absence simulation is displayed in Algorithm 6.2. An employee's state in the next iteration is dependent on its state in the current, while the transition probability depends on whether the upcoming day is a weekday or weekend day and the employee's corresponding risk group. Based on this information, the next state can be found using a look-up from the corresponding cumulative distribution of the transition probabilities calculated in Section 6.3.5. For simplicity, the warm up period is incorporated in the initialization on line 2, just as in Algorithm 6.1.

Algorithm 6.2 Algorithm for the absence uncertainty model

procedure AbsenceSimulation
  $\mathbf{a_0} \leftarrow$ InitializeState$(1 \ldots |N|)$
  $t \leftarrow 0$
  while $t <$ TimeLimit **do**
    for all n in Nurses **do**
      $r \leftarrow$ DrawRandomNumber
      $w \leftarrow$ CheckDayOfWeek$(t+1)$
      $g \leftarrow$ CheckRiskGroup(n, w)
      $a^n_{t+1} \leftarrow$ CumulativeAbsenceLookup(a^n_t, g, r)
      $t \leftarrow t+1$
    end for
  end while
end procedure

6.4 Availability of Non-Absent Employees

There are uncertainties related to which employees will accept working unplanned shifts. As employees may have other plans on their off days, or simply do not feel the need to work more shifts than contracted, the employees might accept an offer to work an unplanned shift with a certain probability.

It is also stochastic which employees have requested working extra shifts or exchanges on specific days. Using an appropriate probability distribution, we can for each shift and day draw how many requests for extra shifts have occurred. As a simplification, we assume an equal probability that each employee with an off day scheduled on a given day will request an extra shift, and draw which employees have requested the shifts from this pool of employees. In a similar manner, we can draw whether an exchange has been requested each day, and then draw which employee requested it from the group of employees with a work shift scheduled.

There is not sufficient data available at DNIC to enable a mathematical calculation of the probabilities above. However, reasonable values may be estimated based on conversations with the scheduling manager. Monte Carlo simulation can then be used to tell whether an employee will accept or not, and which schedule changes have been requested.

6.5 Uncertainty Example

Table 6.9 is based on the example schedule in Section 5.5. Assume that we have made it until Wednesday without any disruptions, and with demand equal to the expected demand of 1 per shift. On Wednesday, it turns out that the real demand is two, and that this is predicted to also be the demand tomorrow. Furthermore, Nurse 4 calls in sick, and is absent from her scheduled Day shift. Nurse 8 hands in a sickness note of seven days and is therefore absent the next week.

With higher demand than expected and two absent nurses, the department will be understaffed on the Day and Night shift today, and on the Evening and Night shift tomorrow, unless actions are taken. Different ways to restore the balance between demand and supply are elaborated on in Chapter 7.

Table 6.9: An example schedule, where the demand on Wednesday and Thursday is higher than expected, Nurse 4 is absent on Wednesday, and Nurse 8 hands in a sickness note of 7 days

Day	Sun	Mon	Tue	Wed	Thu	Fri	Sat	Sun	Mon
Index	-3	-2	-1	0	1	2	3	4	5
Nurse 1	F1	D	D	E	F	F1	E	D	F
Nurse 2	F1	E	E	F	F1	E	D	E	D
Nurse 3	F1	E	D	E	D	F	F	F1	D
Nurse 4	F1	F	E	~~D~~	E	E	F	F1	F
Nurse 5	F1	F	F	F	F1	F	N	N	N
Nurse 6	F1	F	N	N	F	D	F	F1	E
Nurse 7	E	D	F	F	N	N	F	F1	F
Nurse 8	D	F	F	~~D~~	~~E~~	~~F~~	~~F~~	~~F1~~	~~F~~
Nurse 9	N	N	F	F	D	D	F	F1	E
#D	1	1	2	1	2	2	1	1	2
#E	1	2	2	2	1	2	1	1	2
#N	1	1	1	1	1	1	1	1	1
Real demand	1	1	1	2	2	1	1	1	1

Chapter 7

Rescheduling Model

The mixed integer rescheduling model takes a schedule from the scheduling model as well as results from the demand, absence, and availability uncertainty models as input. Based on information about the real supply and demand seen at the Department of Neonatal Intensive Care (DNIC), actions can be taken to ensure that the actual staffing meets the actual demand. The rescheduling model is a rolling horizon model that is called daily. This way, changes made to the initial schedule today is an input to the rescheduling model tomorrow.

This chapter is structured in the same way as Chapter 5. We start by presenting the underlying assumptions and simplifications of the rescheduling model in Section 7.1, before the indices, sets, parameters and variables of the model are defined in Section 7.2. The objective function and constraints are mathematically formulated and explained in Sections 7.3 and 7.4, respectively. We thereby provide an example of how a schedule can be reestablished when the real demand exceeds the supply in Section 7.5, using the same example as in Sections 5.5 and 6.5. We finally provide a model extension in Section 7.6. The rescheduling model and the model extension, is also provided in a compressed form in Appendix A.2.

7.1 Assumptions and Simplifications

The assumptions and simplifications made in Table 7.1 lay the boundaries of the rescheduling model.

Table 7.1: The assumptions of the rescheduling model

Assumption 1	Staff is notified about employee absences and real demand for the whole day at the start of the Day shift
Assumption 2	No knowledge of activity at Maternity Ward
Assumption 3	The full duration of long-term absences are known
Assumption 4	When a short-term absence occurs, it is known whether the absent employee will be absent tomorrow
Assumption 5	Constant number of employees throughout the planning period
Assumption 6	New hires work as regular nurses in Monitoring Unit
Assumption 7	Employees performing special functions are not replaced if they are absent
Assumption 8	Employees always work the full shift when they have a work shift

First, we assume that DNIC is notified about employee absences and real demand at the start of the Day shift. The real demand becomes known for the upcoming day, including the Day, Evening and Night shift. We also assume that there is no communication between the Maternity Ward and DNIC, such that the real demand in the remainder of the current day as well as the expected real demand in the upcoming days is based on the observed demand at the start of the Day shift today alone.

There are two absence types; short-term and long-term. A short-term absence occurs abruptly, and typically lasts for a few days. Although the duration is uncertain, the employees can often tell whether they will be able to work the next day. We therefore assume that DNIC is notified about whether a short-term absence will last until tomorrow or not. For long-term absences, the first day happens unexpectedly, while the rest of the period can be somewhat planned for. We therefore assume that the length of a long-term absence is uncovered on the day when the employee enters a long-term absence.

At DNIC, common approaches to dealing with long-term absences are to hire new employees, or to provide training to selected employees to allow them to cover more demand types. However, in the rescheduling model we assume that the size and skill-mix of the staff is constant throughout the planning period, making such actions impossible.

In the schedule from Chapter 5, special conditions exist for new hires and employees performing special functions on certain days. However, these conditions are assumed to be different in the rescheduling model. First, new hires are assumed to work as regular nurses who contribute to covering demand in the rescheduling model, as this is how they are normally seen in the real-life rescheduling problem at DNIC. Second, employees responsible for the special tasks are not to be replaced if they are absent. This is because we want to analyze how the general demand is met, and considering these special shifts makes the model and analysis unnecessarily complicated.

The final assumption is that employees always work a full shift when they have a work shift scheduled; overtime work of a few hours before or after another work shift is not allowed. Although working such partially double shifts is allowed in the real-life rescheduling problem, we believe that allowing only full double shifts is a reasonable simplification because demand is assumed constant throughout the day, such that e.g. working a half work shift would only make sense if another employee worked the other half.

7.2 Definitions

As in Chapter 5, the section regarding definitions contains the indices, sets, parameters and variables used in the model. Sets are named using uppercase calligraphic letters, variables using lowercase letters, and parameters using uppercase letters. Subscripts indicate indices, while superscripts of capital letters specify the meaning of some parameters and sets. Some parameters have over- or underlines to indicate that they represent upper or lower limits, respectively.

7.2.1 Indices

n	employee
c	skill
s	shift
t	day
k	week
q	double shift type

7.2.2 Sets

Sets denoted $\mathcal{T}$ indicates time, $\mathcal{S}$ is used for sets of shifts, $\mathcal{N}$ is for employees, $\mathcal{C}$ indicates skills, while $\mathcal{Q}$ denotes double shifts.

$\mathcal{T}^{PRE}$	set of days in pre-period
$\mathcal{T}^{R}$	set of days in the replanning period
$\mathcal{T}^{POST}$	set of days in post-period
$\mathcal{T}^{ALL}$	set of days, $\mathcal{T}^{ALL} = \{\mathcal{T}^{PRE} \cup \mathcal{T}^{R} \cup \mathcal{T}^{POST}\}$
$\mathcal{T}^{SN}$	set of days which trigger a short notice ahead of swapping a shift
$\mathcal{T}^{LN}$	set of days which trigger a long notice ahead of swapping a shift
$\mathcal{K}$	set of weeks containing the days $\mathcal{T}^{R} \cup \mathcal{T}^{POST}$
$\mathcal{T}_k$	set of days in week k
$\mathcal{S}$	set of shifts
$\mathcal{S}^W$	set of work shifts, $\mathcal{S}^W = \{D, E, N\}$, $\mathcal{S}^W \subset \mathcal{S}$
$\mathcal{S}^O$	set of off shifts, $\mathcal{S}^O = \{F\}$, $\mathcal{S}^O \subset \mathcal{S}$
$\mathcal{N}$	set of employees
$\mathcal{N}_t^W$	set of available employees who are assigned a work shift on day t, where $t \in \{\mathcal{T}^{R} \cup \mathcal{T}^{POST}\}$
$\mathcal{N}_t^O$	set of available employees who are assigned to an off shift on day t, where $t \in \{\mathcal{T}^{R} \cup \mathcal{T}^{POST}\}$
$\mathcal{N}_t^A$	set of available employees on day t, $\mathcal{N}_t^W \bigcup \mathcal{N}_t^O = \mathcal{N}_t^A$
$\mathcal{C}$	set of skills, $\mathcal{C} = \{1, 2, 3, 4\}$, where 1 = Emergency skills, 2 = Intensive Care skills, 3 = Monitoring skills, 4 = Assistant nurse skills
$\mathcal{N}_c$	set of employees with skill c as their highest ranked skill, $\bigcup_{c \in \mathcal{C}} \mathcal{N}_c = \mathcal{N}$
$\mathcal{Q}$	set of double shift types

In the sets indicating time, the pre-period is the set of days prior to the current day. The replanning period always includes the current day, and possibly also a set of consecutive days after the current day, in which changes in the schedule are allowed to be made. The post-period is a set of consecutive days after the replanning period. Furthermore, the replanning period consists of $\{\mathcal{T}^{SN} \cup \mathcal{T}^{LN}\}$.

The three time dependent sets of nurses, $\mathcal{N}_t^A, \mathcal{N}_t^W$ and $\mathcal{N}_t^O$ are exploited in the formulation of the constraints to include or exclude nurses from the various rescheduling possibilities modelled. Recall from Section 4.3 that the set of available nurses includes all nurses that are non-absent and willing to work an extra shift. The set of work nurses is then then the available nurses with a work shift scheduled, either from the initial schedule or as a result of a change made in a previous rescheduling problem. The set of off nurses includes the available employees with an off shift in the initial schedule, and the employees who have been assigned an off shift due to a previously made change.

7.2.3 Parameters

The parameters in this subsection are sectioned by their functionality. For example, all parameters used to weigh the terms in the objective function are described together.

General Parameters

The parameters concerning the duration of shifts and the work hours of each employee are denoted H. A capital D is used to indicate demand, while M denotes limits on consecutive shifts.

H	number of hours in a full-time work week
H_s	duration of shift s in hours
H_{nk}^{MAX}	number of work hours employee n can work in week k without incurring overtime pay
$\overline{M}^{CW}$	maximum number of consecutive work shifts without penalty per employee
$\overline{M}^{N}$	maximum number of consecutive Night shifts without penalty per employee
$\underline{M}^{B}$	minimum size of buffer assigned each shift in order for swaps and exchanges from the shift to be allowed
$\underline{D}_{st}^{RE}$	minimum online operational demand for employees for shift s on day t
$\underline{D}_{cst}^{RE}$	minimum online operational demand for employees with c as their highest skill for shift s on day t
D_{st}^{SIM}	real demand for employees for shift s on day t
D_{cst}^{SIM}	real demand for employees with c as their highest skill for shift s on day t

Indicator Parameters

The indicator parameters are binary. All except the parameters indicating initially scheduled work shifts are denoted α.

X_{nst}	1 if employee n was initially scheduled to work shift s on day t, 0 otherwise
α_{nst}^{EX}	1 if employee n has requested to work an extra shift s on day t, 0 otherwise
α_{nt}^{OFF}	1 if employee n has requested to exchange a work day on day t, 0 otherwise
α_{nst}^{PRE}	1 if employee n should be preassigned shift s on day t, 0 otherwise
α_{nt}^{CW}	1 if employee n previously has been assigned a number of consecutive work shifts ending on day t that exceeds $\overline{M}^{CW}$, 0 otherwise
α_{nt}^{N}	1 if employee n previously has been assigned a number of consecutive Night shifts ending on day t that exceeds $\overline{M}^{N}$, 0 otherwise

Parameters α_{nst}^{PRE} indicate shifts which have to be assigned, including the shifts the employees worked in the pre-period and changes made to the schedule of today or the upcoming days in previous iterations of the rescheduling model. The preassigned shifts in the replanning period can be both work shifts and off shifts, meaning that e.g. an employee whose work shift has been changed to an off shift on a certain day, cannot be reassigned to a work shift that day.

α_{nt}^{CW} and α_{nt}^{N} are inputs from the previous rescheduling problem. These parameters are included to ensure that only the consecutive shifts assigned in the current rescheduling problem are penalized in the objective function.

Weighing Parameters

The weighing parameters are used to weigh the terms in the objective function, and are all denoted W with a superscript.

W^{REX}	penalty per assigned extra shift requested by an employee
W^{NEX}	penalty per assigned extra shift not requested by an employee
W^{EXC}	penalty per exchanged shift
W^{SN}	penalty per shift swapped on short notice

W^{LN}	penalty per shift swapped on long notice
$W^{DB,(q)}$	penalty per double shift of type q worked
W^{CW}	penalty per consecutive shift worked that exceeds the governmental maximum limit
W^{N}	penalty per consecutive Night shift worked that exceeds the governmental maximum limit
W_t^{D}	penalty per employee in shortage of covering overall demand on day t
W_t^{C}	penalty per employee in shortage of covering skill specific demand on day t
W^{CO}	penalty per hour of overtime worked

The parameters for extra shifts, exchanges, swaps and double shifts penalize each occurrence of the corresponding action. The parameters for working too many consecutive shifts and nights penalize each day in excess of the corresponding maximum limit. The overtime parameter penalizes each hour of overtime assigned as a result of an action taken today.

7.2.4 Variables

The variables are split into decision variables, slack variables and indicator variables. Each variable type is presented in its own section.

Decision variables

When $t < 0$, variables x'_{nst} represent shift assignments in the pre-period. The variables indicating actions, $u_{nt_1t_2}$, v_{nst} and $z_{nt}^{(q)}$, are generally only defined for the days in the replanning period.

$$x'_{nst} = \begin{cases} 1 & \text{if employee } n \text{ is assigned shift } s \text{ on day } t \\ 0 & \text{otherwise} \end{cases}$$

$$u_{nt_1t_2} = \begin{cases} 1 & \text{if employee } n \text{ was initially scheduled to a work shift on day } t_2\text{, but swapped} \\ & \text{or exchanged shifts to get a work shift on day } t_1 \text{ and an off day on day } t_2 \\ 0 & \text{otherwise} \end{cases}$$

$$v_{nst} = \begin{cases} 1 & \text{if employee } n \text{ takes on an extra shift of type } s \text{ on day } t \\ 0 & \text{otherwise} \end{cases}$$

$$z_{nt}^{(1)} = \begin{cases} 1 & \text{if employee } n \text{ works two consecutive shifts without a break in between, where the second shift occurs on day } t \\ 0 & \text{otherwise} \end{cases}$$

$$z_{nt}^{(2)} = \begin{cases} 1 & \text{if employee } n \text{ works two consecutive shifts with a break in between, but without getting sufficient rest according to rules and regulations, where the second shift occurs on day } t \\ 0 & \text{otherwise} \end{cases}$$

Slack Variables

The slack variables are useful to indicate cases when scheduling rules and preferences are violated due to actions taken when the rescheduling problem is solved.

s_{cst}^{C} = unsatisfied demand for skill c for shift s on day t

s_{st}^{D} = unsatisfied demand for employees for shift s on day t

$$s_{nt}^{CW} = \begin{cases} 1 & \text{if employee } n \text{ is assigned a pattern of consecutive work shifts exceeding } \overline{M}^{CW} \text{ that ends on day } t\text{, incurred by changes in the work schedule which were approved on the current day} \\ 0 & \text{otherwise} \end{cases}$$

$$s_{nt}^{N} = \begin{cases} 1 & \text{if employee } n \text{ is assigned a pattern of consecutive Night shifts exceeding } \overline{M}^{N} \text{ that ends on day } t\text{, incurred by changes in the work schedule which were approved on the current day} \\ 0 & \text{otherwise} \end{cases}$$

s_{nk}^{H} = overtime worked by employee n in week k incurred by changes in the work schedule which were approved on the current day

Indicator Variables

The indicator variables in this model keep track of whether assigning a swap or exchange on a given day is allowed or not.

$$d_{st} = \begin{cases} 1 & \text{if there is sufficient staff assigned shift } s \text{ on day } t \text{ to allow a swap or exchange} \\ & \text{from the shift} \\ 0 & \text{otherwise} \end{cases}$$

7.3 Objective Function

To ensure a neater presentation of the objective function, the variables are aggregated into weighted variables, as indicated by the terms (7.1) to (7.7).

$$w_t^{DEM} = \sum_{s \in \mathcal{S}^W} \left(W^D s_{st}^D + \sum_{c \in \mathcal{C}} W^C s_{cst}^C \right) \qquad t \in \mathcal{T}^R \tag{7.1}$$

$$w_t^{GOV} = \sum_{n \in \mathcal{N}} \left(W^{CW} s_{nt}^{CW} + W^N s_{nt}^N \right) \qquad t \in \mathcal{T}^R \tag{7.2}$$

$$w_t^{EX} = \sum_{n \in \mathcal{N}} \sum_{s \in \mathcal{S}^W} \left(W^{REX} \alpha_{nst}^{EX} v_{nst} + W^{NEX} (1 - \alpha_{nst}^{EX}) v_{nst} \right) \qquad t \in \mathcal{T}^R \tag{7.3}$$

$$w_t^{DB} = \sum_{q \in \mathcal{Q}} W^{DB,(q)} \sum_{n \in \mathcal{N}} z_{nt}^{(q)} \qquad t \in \mathcal{T}^R \tag{7.4}$$

$$w_t^{EXC} = W^{EXC} \sum_{n \in \mathcal{N}} \sum_{t_2 \in \mathcal{T}^R \setminus \{t\}} \alpha_{nt_2}^{OFF} u_{ntt_2} \qquad t \in \mathcal{T}^R \tag{7.5}$$

$$\begin{aligned} w^{SWAP} = {} & W^{SN} \sum_{n \in \mathcal{N}} \sum_{t_1 \in \mathcal{T}^{SN}} \sum_{t_2 \in \mathcal{T}^R \setminus \{t_1\}} (u_{nt_1t_2} + u_{nt_2t_1})(1 - \alpha_{t_2}^{OFF}) \\ & + W^{LN} \sum_{n \in \mathcal{N}} \sum_{t_1 \in \mathcal{T}^{LN}} \sum_{t_2 \in \mathcal{T}^{LN} \setminus \{t_1\}} u_{nt_1t_2} (1 - \alpha_{t_2}^{OFF}) \\ & + W^{SN} \sum_{n \in \mathcal{N}} \sum_{t_1 \in \mathcal{T}^{SN}} u_{nt_1t_1} + W^{LN} \sum_{n \in \mathcal{N}} \sum_{t_1 \in \mathcal{T}^{LN}} u_{nt_1t_1} \end{aligned} \tag{7.6}$$

$$w_k^{OVER} = W^{CO} \sum_{n \in \mathcal{N}} s_{nk}^H \qquad k \in \mathcal{K} \tag{7.7}$$

(7.1) are the weighted sums of not meeting total demand and skill specific demand, while

(7.2) are the aggregations of the penalties for working too many consecutive shifts and consecutive nights. (7.3) ensure that working extra shifts is penalized differently, depending on whether the employee requested the shift or not. The aggregated penalties for working double shifts are given by (7.4). (7.5) are the weighted penalties of performing exchanges, while (7.6) gives the weighted sum of assigning a swap, where the penalty depends on whether the swap was assigned on short or long notice. The two last terms in the swap formulation are necessary to avoid that the model interprets same-day swaps as exchanges. Recall that an exchange cannot happen within the same day. Finally, (7.7) keep track of the overtime incurred per week.

Because the variables in (7.1) to (7.7) are already weighted, they are not given any weight in the objective function (7.8).

$$\min Z = \sum_{t \in \mathcal{T}^R} \left(w_t^{DEM} + w_t^{GOV} + w_t^{EX} + w_t^{DB} + w_t^{EXC} \right) + w^{SWAP} + \sum_{k \in \mathcal{K}} w_k^{OVER} \quad (7.8)$$

The objective function (7.8) minimizes understaffing, violations of governmental laws and regulations, the number of schedule changes in terms of extra shifts, swaps, exchanges and double shifts and overtime work.

7.4 Constraints

As in the scheduling problem, the constraints are sectioned by their functionality. For example, all constraints specifically regarding double shifts are formulated in Section 7.4.3.

7.4.1 Covering Real Demand

$$\sum_{n \in \mathcal{N}_t^A} x'_{nst} + s_{st}^D - \underline{M}^B d_{st} \geq \max\{\underline{D}_{st}^{RE}, D_{st}^{SIM}\} \qquad s \in \mathcal{S}^W, t \in \mathcal{T}^R \quad (7.9)$$

$$\sum_{i=1}^{c} \sum_{n \in \mathcal{N}_i \cap \mathcal{N}_t^A} x'_{nst} + s_{cst}^C \geq \max\{\underline{D}_{cst}^{RE}, D_{cst}^{SIM}\} \qquad c \in \mathcal{C}, s \in \mathcal{S}^W, t \in \mathcal{T}^R \quad (7.10)$$

Constraints (7.9) state that the maximum of real overall demand and minimum demand should be covered, and make sure that the variables indicating that swaps are allowed only take on a non-negative value when the staff buffer is larger than $\underline{M}^B$. Constraints (7.10) make sure that the maximum of real demand and minimum demand for employees of each skill is covered.

7.4.2 Technical Constraints for Actions

$$x'_{nst_1} - v_{nst_1} - \sum_{t_2 \in \mathcal{T}^R} u_{nt_1t_2} - \sum_{q \in \mathcal{Q}} z^{(q)}_{nt_1} \leq X_{nst_1} + \alpha^{PRE}_{nst_1} \quad t_1 \in \mathcal{T}^R, n \in \mathcal{N}^A_{t_1}, s \in \mathcal{S}^W \tag{7.11}$$

$$\sum_{s \in \mathcal{S}} x'_{nst} - \sum_{q \in \mathcal{Q}} z^{(q)}_{nt} \leq 1 \quad t \in \mathcal{T}^R, n \in \mathcal{N}^A_t \tag{7.12}$$

$$\sum_{t_1 \in \mathcal{T}^R} u_{nt_10} + \sum_{q \in \mathcal{Q}} z^{(q)}_{n0} \leq 1 \quad n \in \mathcal{N}^A_0 \tag{7.13}$$

All decision variables are connected by constraints (7.11), which also put a limit to the number of cases where variables x'_{nst} may take on a positive value. Constraints (7.12) ensure that each employee is assigned exactly one shift per day, unless the employee is assigned a double shift. Constraints (7.13) make sure that an employee whose shift is swapped or exchanged from the current day, is not assigned a double shift on the two remaining shifts which the employee was not initially assigned this day.

7.4.3 Double Shift

$$x'_{nN(t-1)} + x'_{nDt} - z^{(1)}_{nt} \leq 1 \quad t \in \{0 \dots |\mathcal{T}^R|\}, n \in \mathcal{N}^A_t \tag{7.14}$$

$$x'_{nDt} + x'_{nEt} - z^{(1)}_{nt} \leq 1 \quad t \in \mathcal{T}^R, n \in \mathcal{N}^A_t \tag{7.15}$$

$$x'_{nEt} + x'_{nNt} - z^{(1)}_{nt} \leq 1 \quad t \in \mathcal{T}^R, n \in \mathcal{N}^A_t \tag{7.16}$$

$$x'_{nDt} + x'_{nNt} - z^{(2)}_{nt} \leq 1 \quad t \in \mathcal{T}^R, n \in \mathcal{N}^A_t \tag{7.17}$$

$$x'_{nN(t-1)} + x'_{nEt} - z^{(2)}_{nt} \leq 1 \quad t \in \{0 \dots |\mathcal{T}^R|\}, n \in \mathcal{N}^A_t \tag{7.18}$$

$$\sum_{q \in \mathcal{Q}} z^{(q)}_{nt} \leq 1 \quad t \in \mathcal{T}^R, n \in \mathcal{N}^A_t \tag{7.19}$$

Constraints (7.14) to (7.18) indicate whether double shifts are assigned. Working one double shift of each type on the same day is not allowed, as ensured by constraints (7.19).

Constraints (7.14) and (7.18) are defined for one day more than the duration of the replanning period. This is to ensure that changes made to the schedule on the final day of the period do not result in double shifts in the transition from the replanning period to the post-period.

7.4.4 Swap and Exchange

$$\sum_{t_1 \in \mathcal{T}^R} u_{nt_1t_2} + \sum_{s \in \mathcal{S}^W} \left(X_{nst_2} x'_{nst_2} + \alpha^{PRE}_{nst_2} \right) = 1 \qquad t_2 \in \mathcal{T}^R, n \in \mathcal{N}^W_{t_2} \tag{7.20}$$

$$\sum_{t_2 \in \mathcal{T}^R} u_{nt_1t_2} \leq 1 \qquad t_1 \in \mathcal{T}^R, n \in \mathcal{N}^O_{t_1} \tag{7.21}$$

$$\sum_{t_1 \in \mathcal{T}^R} u_{nt_1t_2} \leq 1 \qquad t_2 \in \mathcal{T}^R, n \in \mathcal{N}^W_{t_2} \tag{7.22}$$

$$\sum_{t_1 \in \mathcal{T}^R} u_{nt_1t_2} - \sum_{s \in \mathcal{S}^W} d_{st_2} \left(X_{nst_2} + \alpha^{PRE}_{nst_2} \right) \leq 0 \qquad t_2 \in \mathcal{T}^{LN}, n \in \mathcal{N}^W_{t_2} \tag{7.23}$$

Recall that employees are swapped or exchanged *to* a shift on day t_1, *from* a shift on day t_2. It is possible to be swapped between shifts on the same day, while the days have to be different for exchanges to happen.

(7.20) ensure that employees who have a work shift scheduled on day t, either work the shift or swap or exchange the shift. It is not necessary to multiply $\alpha^{PRE}_{nst_2}$ by x'_{nst_2}, as the variables are locked to the values of the indicator parameters in constraints (7.32). Constraints (7.21) make sure that an employee can maximum be swapped or exchanged to each day once. Ensuring that an employee can maximum be swapped or exchanged once from each work shift is done by constraints (7.22). Finally, constraints (7.23) ensure that swaps are only allowed to happen on shifts with a sufficiently large buffer.

7.4.5 Consecutive Work

$$\sum_{\tau=t-\overline{M}^{CW}}^{t} \left(\sum_{s\in\mathcal{S}^W\setminus\{N\}} x'_{ns\tau} + x'_{nN(\tau-1)} \right) - s^{CW}_{nt} \leq \overline{M}^{CW} + \alpha^{CW}_{nt} \quad n\in\mathcal{N}, t\in\mathcal{T}^R\cup\mathcal{T}^{POST} \tag{7.24}$$

$$\sum_{\tau=t-\overline{M}^{N}}^{t} x'_{nN\tau} - s^{N}_{nt} \leq \overline{M}^{N} + \alpha^{N}_{nt} \qquad n\in\mathcal{N}, t\in\mathcal{T}^R\cup\mathcal{T}^{POST} \tag{7.25}$$

$$\sum_{s\in\mathcal{S}^W}\sum_{t\in\mathcal{T}_k} H_s x'_{nst} - s^{H}_{nk} \leq H^{MAX}_{nk} \qquad n\in\mathcal{N}, k\in\mathcal{K} \tag{7.26}$$

Constraints (7.24) and (7.25) indicate whether an employee ends up working more consecutive shifts or consecutive nights than what is recommended by governmental regulations. In (7.24) we distinguish between the Night shift and the other work shifts as a consequence of the placement of the Night shift in the rescheduling problem compared to the scheduling problem. The overtime assigned to each employee is kept track of by constraints (7.26).

7.4.6 Variable Declarations and Fixations

$$x'_{nst} \in \{0,1\} \qquad t\in\mathcal{T}^{ALL}, n\in\mathcal{N}^A_t, s\in\mathcal{S} \tag{7.27}$$

$$u_{nt_1t_2} \in \{0,1\} \qquad t_1,t_2\in\mathcal{T}^R, n\in\mathcal{N}^A_{t_1}\cap\mathcal{N}^W_{t_2} \tag{7.28}$$

$$v_{nst} \in \{0,1\} \qquad t\in\mathcal{T}^R, n\in\mathcal{N}^O_t, s\in\mathcal{S}^W \tag{7.29}$$

$$z^{(q)}_{n0} \in \{0,1\} \qquad n\in\mathcal{N}^W_0, q\in\mathcal{Q} \tag{7.30}$$

$$z^{(q)}_{nt} = 0 \qquad t\in\{1\ldots|\mathcal{T}^R|\}, n\in\mathcal{N}^A_t, q\in\mathcal{Q} \tag{7.31}$$

$$x'_{nst} = 1 \qquad n\in\mathcal{N}, s\in\mathcal{S}, t\in\mathcal{T}^{ALL} \mid \alpha^{PRE}_{nst} = 1 \tag{7.32}$$

$$x'_{nst} = X_{nst} \qquad t\in\mathcal{T}^R, n\in\mathcal{N}\setminus\mathcal{N}^A_t, s\in\mathcal{S} \mid \alpha^{PRE}_{nst} = 0 \tag{7.33}$$

$$s^{C}_{cst} \in \mathbb{N}_0 \qquad c\in\mathcal{C}, s\in\mathcal{S}^W, t\in\mathcal{T}^R \tag{7.34}$$

$$s^{D}_{st} \in \mathbb{N}_0 \qquad s\in\mathcal{S}^W, t\in\mathcal{T}^R \tag{7.35}$$

$$s^{CW}_{nt} \in \{0,1\} \qquad n\in\mathcal{N}, t\in\mathcal{T}^R\cup\mathcal{T}^{POST} \tag{7.36}$$

$$s^N_{nt} \in \{0,1\} \qquad n \in \mathcal{N}, t \in \mathcal{T}^R \cup \mathcal{T}^{POST} \tag{7.37}$$

$$s^H_{nk} \geq 0 \qquad n \in \mathcal{N}, k \in \mathcal{K} \tag{7.38}$$

$$d_{st} \in \{0,1\} \qquad s \in \mathcal{S}^W, t \in \mathcal{T}^{LN} \tag{7.39}$$

Constraints (7.27) to (7.30) state that all decision variables are binary. Variables $u_{nt_1t_2}$ only exist if employee n was initially assigned a work shift on day t_2 and is available on day t_1. Similarly, variables v_{nst} only exist for available employees that initially were assigned an off shift on day t. The variables indicating double shifts may only take on a value of 1 on the current day, as employees are only allowed to be assigned double shifts on the current day. This is ensured by constraints (7.30) and (7.31).

Constraints (7.32) lock variables x'_{nst} to 1 for the shifts the employees were actually assigned in the pre-period. The constraints also ensure that changes previously made to the schedule of today or the upcoming days actually are enforced. Unavailable employees are assigned their initial shifts by constraints (7.33) to allow the constraints regarding consecutive shifts and overtime to work correctly, although they do not count in the set of employees who cover demand.

The slack variables s^C_{cst} and s^D_{st} should be non-negative integers, ensured by constraints (7.34) to (7.35). Furthermore, s^{CW}_{nt} and s^N_{nt} are binary variables, ensured by constraints (7.36) and (7.37). (7.38) are the non-negativity constraints for the slack variables keeping track of the overtime hours worked. Finally, constraints (7.39) make sure that the variables indicating whether swaps are allowed from a certain shift are binary.

7.5 Rescheduling Example

Recall the example schedule for 9 nurses for a 9-day period in Table 5.6 of Chapter 5, which was subject to uncertainty in Table 6.9 of Chapter 6. We now illustrate the various rescheduling actions using the same schedule, where D, E, N and F denotes a Day, Evening, Night and Off shift, respectively. In Table 7.2, the current day is Wednesday, and the duration of the replanning period is three days. An important distinction from Table 5.6 is that the Nights shifts are now scheduled on the days when the shifts start as opposed to when the shifts end. Also, the F1 day is now denoted as F because the two off day types are treated equally by the scheduling manager when the real-life rescheduling problem is solved. Both these distinctions were explained in Section 4.3.

The examples from all three chapters are assembled in Appendix B.1 to B.3.

Table 7.2: An example schedule, where the current day is denoted by index 0 and the replanning period consists of three days. Nurse 4 is absent today, Nurse 8 is long-term absent, and the real demand today and tomorrow is higher than expected

	Pre-period			Replanning period			Post-period		
Day	Sun	Mon	Tue	Wed	Thu	Fri	Sat	Sun	Mon
Index	-3	-2	-1	0	1	2	3	4	5
Nurse 1	F	D	D	E	F	F	E	D	F
Nurse 2	F	E	E	F	F	E	D	E	D
Nurse 3	F	E	D	E	D	F	F	F	D
Nurse 4	F	F	E	~~D~~	E	E	F	F	N
Nurse 5	F	F	F	F	F	N	N	N	F
Nurse 6	F	N	N	F	F	D	F	F	E
Nurse 7	E	D	F	N	N	F	F	F	F
Nurse 8	D	F	F	~~D~~	~~E~~	~~F~~	~~F~~	~~F~~	~~F~~
Nurse 9	N	F	F	F	D	D	F	F	E
#D	1	2	2	0	2	2	1	1	2
#E	1	2	2	2	1	2	1	1	2
#N	1	1	1	1	1	1	1	1	1
#F	6	4	4	6	4	4	6	6	4
Real demand	1	1	1	2	2	1	1	1	1

In the scheduling problem, the *expected* demands $\underline{D}_{st}$ for all shifts are one. Now assume that the *real* demand on the current day is two, and that this is expected to also be the real demand tomorrow. Furthermore, assume that Nurse 4 is absent from today's scheduled Day shift, and Nurse 8 is absent from the Day shift today and the Evening shift tomorrow. With the current schedule, the department will understaffed on the Night shift both today and tomorrow, as well as on the Day shift today and the Evening shift tomorrow. In the process of reestablishing the schedule, the scheduling manager can change the schedule of Wednesday, Thursday and Friday, as these are the days in the replanning period.

One feasible solution to the rescheduling problem on Wednesday is indicated in Table 7.3. In this example solution, we assume that all non-absent employees are available. Whether this solution would be the optimal one depends on the values of the weights in the objective function. The purpose of Table 7.3 is simply to illustrate how all the actions work rather than how to find the optimal solution. All the swaps could have been exchanges, depending on whether they had been requested by the employees or not.

As there are constraints and preferences governing how several of the actions are allowed to be performed, it varies which unplanned shifts each employee may take on. For example, it would make little sense to swap the work shifts of Nurse 3 on day 0 and 1, as these shifts contribute to covering the real demand. However, it could be that the best action regarding tomorrow's Night shift was to wait with assigning it until tomorrow, depending on how unmet future demand is penalized in the objective function. This way, Nurse 3 would be allowed to work double shifts tomorrow, covering both Day and Night shift.

Table 7.3: Example of one possible assignment of actions when the real demand on day 0 and 1 is two, Nurse 4 is absent today and Nurse 8 has a long-term absence. A change to an employee's initial shift is illustrated by a slash cancellation and highlighted by blue

Day Index	Tue -1	Wed 0	Thu 1	Fri 2	Action	Variable
Nurse 1	D	$\cancel{\text{E}}$ E + N	F	F	Double shift of type 1	$z_{10}^{(1)} = 1$
Nurse 2	E	F	F	E	No actions are taken	
Nurse 3	D	E	D	F	No actions are taken	
Nurse 4	E	$\xcancel{\text{D}}$	E	E	No actions are possible	
Nurse 5	F	F	$\cancel{\text{F}}$ N	N	Extra shift Thursday	$v_{5N1} = 1$
Nurse 6	N	F	$\cancel{\text{F}}$ E	$\cancel{\text{D}}$ F	Day shift Friday swapped	$u_{612} = 1$
Nurse 7	F	$\cancel{\text{N}}$ D + N	N	F	Double shift of type 2	$z_{70}^{(2)} = 1$
Nurse 8	F	$\xcancel{\text{D}}$	$\xcancel{\text{E}}$	$\xcancel{\text{F}}$	No actions are possible	
Nurse 9	F	$\cancel{\text{F}}$ D	D	D	Extra shift today	$v_{9D0} = 1$
#D	2	2	2	1		
#E	2	2	2	2		
#N	1	2	2	1		
#F	4	3	3	5		
D_{st}^{SIM}	1	2	2	1		

7.6 Reactive Model Extensions

In Table 4.2, we presented two reactive managerial insights which we should seek to obtain. The first insight was to find the optimal duration of the replanning period, while the second was to assess whether a feasible schedule can be reestablished during the online operational phase without violating the rules regarding rest. The first insight can be obtained by varying the size of the set of days in the replanning period, $|\mathcal{T}^R|$. The second insight requires a model extension.

7.6.1 Extension 1: Stricter Rescheduling

Not getting sufficient rest may negatively affect the health of the employees and their performance at work. In this strategy, we therefore prohibit double shifts and violations of the rules regarding consecutive work shifts and Night shifts. Although the managers at DNIC seek to follow the rules as best they can, it is most important to take proper care of the patients. Sometimes, they therefore have no choice other than to let these events occur in order to meet demand. In this extension, we want to find out if this is strictly necessary.

Definitions

Variables $z_{nt}^{(q)}$ are removed from the problem. The same are the slack variables s_{nt}^{CW} and s_{nt}^{N}. Their corresponding weighing parameters, $W^{DB,(q)}$, W^{CW} and W^{N}, as well of the set of double shift types, $\mathcal{Q}$, can also be removed.

Objective and Constraints

When double shifts are illegal, several terms and constraints can be removed from the scheduling model. First, we need not include the terms penalizing double shifts and violation of rules regarding consecutive shifts, w_t^{DB} and w_t^{GOV}, in the objective function. Furthermore, constraints (7.15) to (7.17), (7.19), (7.30), (7.31), (7.36) and (7.37), are all removed from the model.

Constraints (7.40) to (7.43) replace constraints (7.11) to (7.14), respectively, while (7.44) replace constraints (7.18). Constraints (7.45) and (7.46) regard consecutive shifts and nights, and replace constraints (7.24) and (7.25). The constraints have the same function as the respective constraints they replace.

$$x'_{nst_1} - v_{nst_1} - \sum_{t_2 \in \mathcal{T}^R} u_{nt_1t_2} \leq X_{nst_1} + \alpha^{PRE}_{nst_1} \quad t_1 \in \mathcal{T}^R, n \in \mathcal{N}^A_{t_1}, s \in \mathcal{S}^W \tag{7.40}$$

$$\sum_{s \in \mathcal{S}} x'_{nst} = 1 \quad t \in \mathcal{T}^{\mathcal{R}}, n \in \mathcal{N}^A_t \tag{7.41}$$

$$\sum_{t_1 \in \mathcal{T}^R} u_{nt_10} \leq 1 \quad n \in \mathcal{N}^A_0 \tag{7.42}$$

$$x'_{nN(t-1)} + x'_{nDt} \leq 1 \quad n \in \mathcal{N}^A_t, t \in \{0 \ldots |\mathcal{T}^R|\} \tag{7.43}$$

$$x'_{nN(t-1)} + x'_{nEt} \leq 1 \quad n \in \mathcal{N}^A_t, t \in \{0 \ldots |\mathcal{T}^R|\} \tag{7.44}$$

$$\sum_{\tau = t - \overline{M}^{CW}}^{t} \left(\sum_{s \in \mathcal{S}^W \setminus \{N\}} x'_{ns\tau} + x'_{nN(\tau-1)} \right) \leq \overline{M}^{CW} \quad n \in \mathcal{N}, t \in \mathcal{T}^R \cup \mathcal{T}^{POST} \tag{7.45}$$

$$\sum_{\tau = t - \overline{M}^N}^{t} x'_{nN\tau} \leq \overline{M}^N \quad n \in \mathcal{N}, t \in \mathcal{T}^R \cup \mathcal{T}^{POST} \tag{7.46}$$

Finally, although several elements can be stripped from the rescheduling model when double shifts and working too many consecutive shifts and nights are illegal, a simple way to implement the changes are to simply put variables $z^{(q)}_{nt}$, s^{CW}_{nt} and s^N_{nt} equal to zero for all cases.

Chapter 8

Case Study

Both the scheduling model and the rescheduling model are implemented using the programming language Mosel, run through the commercial optimization software FICO® Xpress Optimisation Suite 8.3. All input parameters are stored in Microsoft Excel and .txt-files. The simulation script is written in Python 2.7.

The scheduling model can be run as a standalone model, giving a nurse schedule as output. The rescheduling model requires input from the simulation models as well as an initial schedule to work. Furthermore, the rescheduling model is iteratively called using a Python script, once for each day in the planning period. The output from the rescheduling model is written to a .txt-file, which again is used as input to the model in the next iteration.

All instances are solved on computers with Intel i7-6700 CPU and 32GB RAM, running on a Windows 10 Education 64-bit Operating System.

An overview of all case instances is provided in Section 8.1. Sections 8.2 and 8.4 contain descriptions of the data used in all scheduling and rescheduling instances, respectively. In Section 8.3, we present the results from the simulation models. Finally, a brief technical analysis of the results of running all instances is included in Section 8.5.

8.1 Case Instances

Table 8.1 contains a summary of the proactive and reactive scheduling extensions mathematically modelled in Chapters 5 and 7, respectively. Each of the extensions correspond to a case instance. The data included in the case instances are presented in the remainder of this chapter.

Table 8.1: Definitions of case instances and their connection to model extensions and managerial insights

Name	**Model Extension**	**Insight**
`s0`	None	Robustness of today's scheduling strategy
`s1`	Buffer	Optimal assignment of surplus work hours
`s2`	Ghost	Optimal placement of off shifts
`s3`	Absence	Value of utilizing historical absence information
`s4`	Extra weekends	Value of additional staffing during weekends
`r0`	None	Today's rescheduling strategy
`r1`	Strict rescheduling	Necessity of violating rules

8.2 Data in the Scheduling Instances

In this section, we elaborate on the parameter values of the different instances the scheduling model. Unless otherwise stated, the overlapping parameters in the model extensions and `s0` have the same values.

The nurse schedules generated in this thesis are created with the purpose of identifying robust characteristics of schedules, and are not to be used in real-life. In Section 5.1 we assume that the general robustness of schedules is best tested in an environment when no holidays and special events occur. For that reason, and because we use real personal shift requests from the Department of Neonatal Intensive Care (DNIC) as input to the model, the longest possible time period we can create a nurse schedule for using real data is from September 4th 2017 to December 17th 2017.

8.2.1 S0: Base Case

The base case scheduling instance is best described as the data that represents the real-life situation at DNIC as closely as possible.

General settings To provide an overview of the size and complexity of the scheduling problem solved, key data regarding the problem settings and important real-life parameter values are presented in Table 8.2.

Table 8.2: Key settings in the scheduling problem

Setting	Value	Description
$\lvert\mathcal{T}\rvert$	105	Days in planning period
$\lvert\mathcal{T}^P\rvert_{t\leq 0}$	7	Days from previous planning period used as input
$\lvert\mathcal{N}\rvert$	117	Number of employees
H^{DEV}	5%	Allowed deviation between the hours in the work contract and the hours actually assigned for the entire planning period
U	50%	Minimum amount of shifts which must be Evening or Night
B	1.5	Work load of a Night shift relative to an Evening shift
M^{NW}	3	Frequency of work weekends
$\overline{M}^{CW}$	7	Maximum number of consecutive shifts
$\overline{M}^{N}$	4	Maximum number of consecutive Night shifts
$\sum \beta^R_{nst}$	4940	Total number of requests
$\lvert\mathcal{P}^{DW}\rvert$	2	Number of desirable shift patterns
$\lvert\mathcal{P}^{GEN}\rvert$	18	Number of undesirable shift patterns
$\lvert\mathcal{P}^{ILL}\rvert$	2	Number of illegal shift patterns

Some employees had filled in few or zero personal requests in their personal schedules, meaning that the scheduling model barely rewards the assignment of shifts to these employees in the objective function. Therefore, ideally $\sum \beta^R_{nst}$ should have been higher. If all employees had filled in the number of requests corresponding to their work contract, there would be roughly 8300 requests for shifts in total.

Some deviations from the actual contracted work had to be allowed to keep the model implementation from being too tight. Because the scheduling manager tries to keep the deviations to a minimum, we have worked to keep the maximum allowed deviation, H^{DEV}, as small as possible.

Demand and employees The parameter values for the overall demand are listed in Table 8.3. All demand limits are based on conversations with the scheduling manager, and reflect the demand requirements used in the real-life scheduling problem. Not indicated in Table 8.3, the demand on Night shifts on Mondays is the same as the demand on Night shifts during weekends, due to the requirement that the same employees should work Night both Saturday, Sunday and Monday.

Table 8.3: Minimum and maximum demand per shift and day of the week

Parameter	Weekdays		Weekends		Description
	Min	Max	Min	Max	
$\underline{D}_{Dt}$ / $\overline{D}_{Dt}$	16	30	13	14	Demand on Day shift
$\underline{D}_{Et}$ / $\overline{D}_{Et}$	14	20	13	14	Demand on Evening shift
$\underline{D}_{Nt}$ / $\overline{D}_{Nt}$	14	20	13	14	Demand on Night shift

Table 8.4 indicates the minimum demand for employees per skill category and shift. The minimum demand for assistant nurses is always 0, as assistant nurses do not possess all skills necessary to cover the minimum demand, consistent with Section 2.3.2. As opposed to the total minimum demand, the limits for skill specific demand per shift are the same on all days of the week. The total minimum demand at DNIC is higher than the sum of minimum demands for the specific skill types. We consider the skills of the last employees required to fulfill the total minimum demand per shift as arbitrary.

Table 8.4: Minimum demand per skill category per shift

Parameter	Day	Evening	Night	Description
$\underline{D}_{1st}$	3	2	2	Demand for Emergency skills
$\underline{D}_{2st}$	4	4	4	Demand for Intensive Care skills
$\underline{D}_{3st}$	3	2	2	Demand for Monitoring skills
$\underline{D}_{4st}$	0	0	0	Demand for Assistant skills

The number of nurses of each skill employed at DNIC is indicated in Table 8.5. As indicated in Table 8.3, a minimum of 39 employees must work each weekend. As employees only work every third weekend, at least 117 employees must be employed. To be able to cover the weekend demand, while simultaneously staying within the department's budget constraints, mostly part-time workers are employed, which was also discussed in Section 2.3.2. The average contracted work is approximately 73%. Employees with full-time work contracts should work 35.5 hours per week on average.

Table 8.5: Number of employees by skill category

Skill set	Value	Description
$\lvert\mathcal{N}_1\rvert$	36	Emergency skills
$\lvert\mathcal{N}_2\rvert$	39	Intensive Care skills
$\lvert\mathcal{N}_3\rvert$	37	Monitoring skills
$\lvert\mathcal{N}_4\rvert$	5	Assistant skills

Weights in the objective function The values of the three weighting parameters in the scheduling objective function are given in Table 8.6. We consider assigning requested shifts as the most important objective, and set the corresponding reward, W^R, to a value of 1 due to simplicity.

Assigning desirable patterns is also considered important, but not to the extent that it should compromise the assignment of requested shifts as would be the case if the reward for a desirable pattern exceeded the request. We therefore set the value of the corresponding parameter, W^P, a bit lower than W^R. Finally, it is our impression that having requests satisfied and working good shift patterns is significantly more important to the health and satisfaction of the employees than having to sometimes rank down in skill. We reflect by setting the corresponding penalty, W^S, to the lowest value of the three parameters.

Table 8.6: Values of the weighing parameters in `s0`

Parameter	Value	Description
W^R	1	Request reward
W^P	0.8	Desirable pattern reward
W^S	0.3	Penalty of ranking down in skill

8.2.2 S1: Buffer

The final values of the weighing parameters in `s1` are displayed in Table 8.7. The weights W_r^B are based on the buffer size F, and the weight for the first employee, W_1^B. Using a linear reduction, W_r^B are calculated using Equation (8.1). Parameters W_r^{SB} are calculated similarly.

$$W_r^B = \sum_{i=1}^{r} \left(W_0^B - \frac{W_0^B}{F}(i-1)\right) \qquad r \in \{1, \ldots, F\} \tag{8.1}$$

It is essential to keep $W^S > W_0^{SB}$. Otherwise, the model finds it optimal to assign severe understaffing in terms of skills on some shifts in exchange for a big buffer of employees with the same skills on other shifts, as the buffer is rewarded higher than the understaffing is penalized. This would naturally be an undesirable result.

Table 8.7: Key settings and weighing parameters in `s1`

Parameter	Value	Description
F	5	Maximum size of buffer rewarded in objective function
W^P	0.5	Desirable pattern reward
W^S	2	Penalty of ranking down in skill
W_1^B	1.5	Reward for first employee in excess of minimum demand for overall staffing
W_1^{SB}	1.5	Reward for first employee in excess of minimum demand in terms of skills

In addition to rewarding the assignment of surplus employees up to a certain limit in the objective function, we reduce the upper demand limit on each shift in `s1` for the weekday shifts. The maximum demand for Day, Evening and Night shifts are $\overline{D}_{nDt} = 23$, $\overline{D}_{nEt} = 19$ and $\overline{D}_{nNt} = 19$, respectively.

8.2.3 S2: Ghost

Two parameters are introduced in `s2`; each with its corresponding value indicated in Table 8.8. The ghost reward was set after preliminary testing, where we tried to set $W^G > W^R$ in order to enhance the importance of this particular strategy as compared to the other objectives. However, this significantly increased the time required to find a solution with a reasonably low optimality gap.

Table 8.8: Key settings in `s2`

Setting	Value	Description
H	35.5	Duration of full-time work week
W^G	0.5	Reward per ghost shift assigned

8.2.4 S3: Absence

Table 8.9 contains the values of the new settings and parameters introduced in `s3`. The size of the high-risk group amounts to 21 for both weekdays and weekends, but with a different group of employees in each set. The employees were identified in the data analysis in Section 6.3.5.

The contribution to demand from the high-risk group should be smaller than the contribution from the rest of the employees. The most robust choice would be to put the contribution equal to zero, but this would result in an infeasible problem instance. To ensure both a feasible solution and buffers on each weekday shift with high-risk employees, we set the contribution equal to 0.5.

Approximately seven high-risk employees must be assigned work shifts each weekend. We do not want to penalize each shift they are assigned to, suggesting that S^{HR} should be higher than 1. To evenly distribute the high-risk employees, we consider 2 an appropriate limit. Further, the penalty of assigning an employee in the high-risk group to a work shift exceeding S^{HR} must exceed the reward for fulfilling a request. Otherwise, the request would be dominant and we would not achieve an even distribution.

Table 8.9: Key settings in `s3`

Setting	Value	Description
$\|\mathcal{N}_t^{HR}\|$	21	Size of set with high risk of being absent, $t \in \mathcal{T} \setminus \mathcal{T}^W$
$\|\mathcal{N}_t^{HR}\|$	21	Size of set with high risk of being absent where $t \in \mathcal{T}^W$
C_{nt}	0.5	Contribution to demand from employee n in the high-risk group, $n \in \mathcal{N}_t^{HR}$, $t \in \mathcal{T} \setminus \mathcal{T}^W$
C_{nt}	1	Contribution to demand from employee n in the high-risk group, $n \in \mathcal{N}_t^{HR}$, $t \in \mathcal{T}^W$
S^{HR}	2	Number of employees in high-risk group who can be assigned to the same shift without getting penalized during weekends
W^{HR}	1.2	Penalty per additional employee with high risk of being absent

8.2.5 S4: Extra Weekends

The key settings for model extension `s4` are displayed in Table 8.10. The minimum demand is increased by 1 to 14 employees for all weekend shifts. As working employees work both Saturday and Sunday, and 3 shifts must be covered each day, at least $3 \times 15 = 45$ extra weekends have to be assigned for the 15 week long planning period. With maximum $\overline{M}^{EW} = 2$ extra weekends per employee, minimum $\lceil \frac{45}{2} \rceil = 23$ employees must accept working extra weekends.

The nurses working extra weekends are arbitrarily chosen based on a set of selection criteria. First, the employees should have a contracted work of at least 75 % to be sure that they have sufficient contracted hours for numerous parameter settings to work. Second, to isolate the effect of being in a high-risk group, explained in `s3`, all the selected employees should be in the low-risk group. Third, equally many employees of each skill should be selected. In the end, 24 nurses are selected.

Assuming a duration of 8 hours per work shift, the number of off hours gained per extra work weekend, H^{EO}, is set based on the number of extra off days gained multiplied by this shift duration. It is likely that the employees will welcome the trade better the more extra off days are offered in exchange. However, offering too many extra off days could result in problems with covering demand during weekdays instead. As a compromise, we set the number of off days to 3, or 24 hours.

Table 8.10: Key settings in `s4`

Setting	Value	Description
$\underline{D}_{st}$	14	Minimum demand for all work shifts s when t is a Saturday or Sunday
$\overline{M}^{EW}$	2	Maximum number of extra weekends per employee
$\sum \beta_n^{EW}$	24	Number of employees who can work extra weekends
H^{EO}	24	Extra off hours gained per extra weekend worked

8.3 Setup of the Uncertainty Models

All uncertainties are realized using simulation. The simulation results for a given day, which are an important input to the rescheduling model, consists of one realization of demand, one of absence, and one of availability. The demand realization includes the number of patients at each of the five levels. Whether any of the employees at DNIC have entered or recovered from either a short-term or long-term absence is realized through the absence simulation. Finally, the availability realization is dependent on the realization of absence and represents whether a non-absent employee is willing to work an extra shift on an off day. Any requests for either a work shift or an off shift are also a part of the simulation of availability.

8.3.1 Technical Settings

To obtain realistic and comparable results, some technical settings are set. First, we consider the generation of random numbers in our simulations. To be able to reproduce our results we manually set the seed for the random number generator for each simulation run.

Another important setting is the number of simulation runs. We run tests of 50, 100, 200, and 500 simulations to observe how the stability of the output is affected. Figure 8.1 shows how the variance in the total number of patients is reduced when going from 50 to

500 simulation runs. There is a small decrease in variance when increasing the number of simulation runs from 50 to 200. However, the marginal gain per added simulation levels out when increasing the number of runs further. Consequently, we put the number of simulation runs to 200. The number of runs for the absence and availability simulation models is set to 200 using the same approach.

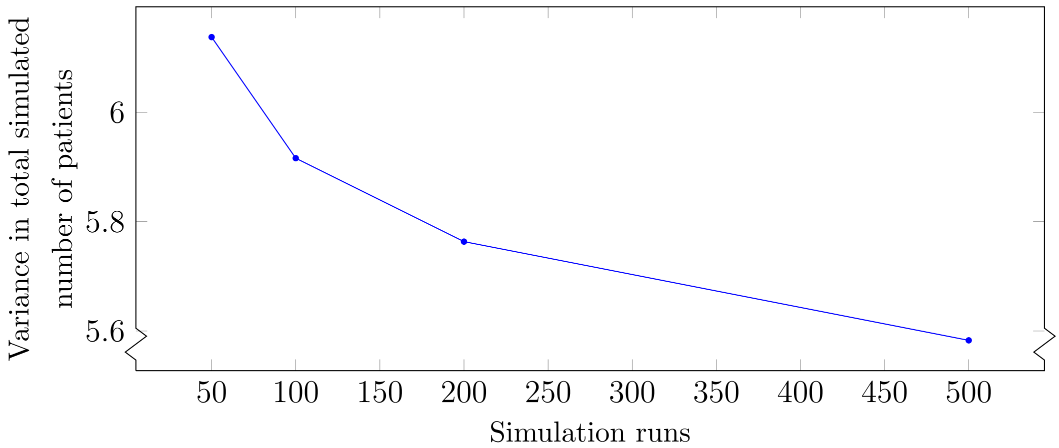

Figure 8.1: Variance in total number of patients for various simulation runs

The final setting for the simulation models is the duration of the warm-up period. This period must be sufficiently long for the results to stabilize. We set this duration to 364 days for the demand and absence simulation models. As the availability simulation is a static Monte Carlo simulation, warm-up is not relevant for this model.

8.3.2 Demand Simulation

The results of doing the calculations proposed in Section 6.2.5 are the transition probabilities in Table 8.11. Each row in the table indicates the probabilities that the state of a hospital bed at the given level transitions to each of the other states, where state 0 is an empty bed, and state 1-5 means that a patient on the corresponding level occupies the bed.

Table 8.11: The probabilities of transitioning between patient levels

		To level					
		0	**1**	**2**	**3**	**4**	**5**
From level	**0**	0.91	0.02	0.02	0.04	0.02	0.00
	1	0.19	0.50	0.29	0.01	0.00	0.01
	2	0.06	0.03	0.88	0.03	0.00	0.00
	3	0.00	0.02	0.12	0.80	0.06	0.01
	4	0.03	0.00	0.01	0.05	0.90	0.01
	5	0.23	0.00	0.00	0.00	0.08	0.69

To validate the results of the demand uncertainty model, the simulated values are compared to the historical ones. In Figure 8.2 the average number of patients per level is plotted for both the simulation results and the historical values. Error bars displaying the spread one standard deviation from the mean are included to illustrate how the variation of the simulation results is consistent with the historical variation. The difference between the plots are minimal, suggesting that the daily variations between the patient levels are realistic.

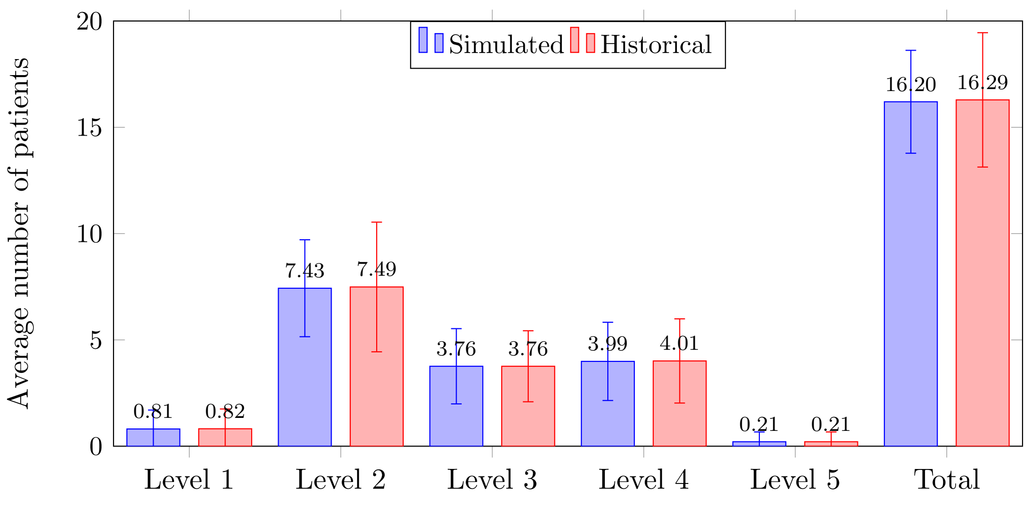

Figure 8.2: Comparison of the average number of patients for simulation results and historical values for 200 simulation runs

Figure 8.3 is included to give an impression on the daily variation in demand per patient.

The figure shows that the simulation model produces an average need for care per patient that fluctuates in a similar manner as a random historical sample.

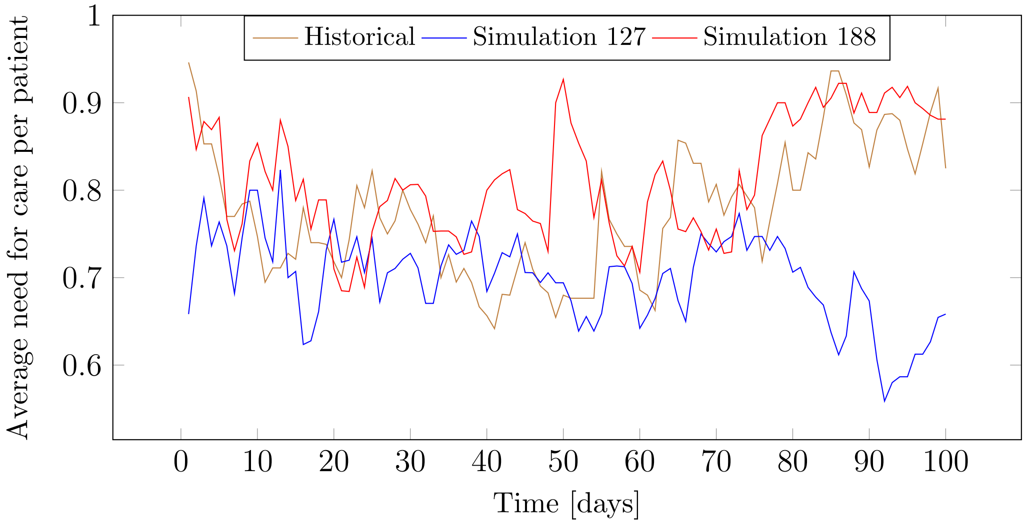

Figure 8.3: Two arbitrary simulations compared to the historical average need for care per patient during the period from January 1st 2016 through April 9th 2016

The results from the demand simulation, which are the number of patients on each level on a given day, are used as inputs to all the upcoming rescheduling instances in Section 8.4.

8.3.3 Absence Simulation

The cores of the absence simulation are the four transition probability matrices described in Section 6.3.5. Two of them are included in Table 8.12, while the remaining are presented in Appendix C.1. Notice how the probability of remaining non-absent is much higher while the probability of remaining short-term absent is lower for the low-risk group in Table 8.12a than for the high-risk group in Table 8.12b. The probabilities of transitioning to a long-term absence state are practically the same in both tables, as they should be.

Table 8.12: Probabilities for transitioning between the states non-absent (a_N), short-term absent (a_S), and long-term absent (a_L)

(a) Low-risk group on weekdays

		To		
		a_N	a_S	a_L
From	a_N	0.97	0.03	0.00
	a_S	0.43	0.56	0.01
	a_L	0.05	0.00	0.95

(b) High-risk group on weekdays

		To		
		a_N	a_S	a_L
From	a_N	0.89	0.10	0.01
	a_S	0.35	0.64	0.01
	a_L	0.03	0.01	0.96

Figure 8.4 shows the comparison between the average absence rate on all days for the simulation and historical values. The short-term absence rate is divided into weekdays and weekends to reflect that the time resolution is important for this type of absence, as concluded in Section 6.3.5. The figure illustrates that we overestimate the rate of both short-term and long-term absences, and in turn the total rate of absence. There are several possible explanations for this. First, we use historical data from work days to generate probability distributions used to simulate absences on both off days and work days. If we assume that employees more often are classified as sick when they are scheduled to work than when they enjoy a day off, we would in fact overestimate the probability of being absent on off days. This could be the case for example if an employee showed up sick at work and as a result was sent home. If this employee had a day off instead, it might be that he or she would not be classified as sick. Using this line of thought, it could be that our absence rates are not that far off on work days, and that the difference from the historical values can be explained by the overestimates on off days.

Another explanation is that the proportions of high-risk employees are higher in our scheduling instances than in the data set, which can result in higher absence rates. In the data set from the data analysis in Section 6.3, 27 and 28 out of 165 employees are classified in the high-risk group on weekdays and weekends, respectively. This amount to proportions of 17.0 % and 16.4 %. In our scheduling instances, 21 out of 117 employees, a proportion of 17.9 %, are put in the high-risk group on both weekdays and weekends.

The spread is generally lower in the simulation results than for the historical values. The most likely explanation for this is the aggregation of the employees in the two groups as proposed in Section 6.3.5. When using only two different groups of probability distributions to model absences that in the historical data are unique for each employee, it follows that the variations in the simulation results are lower than the historical ones.

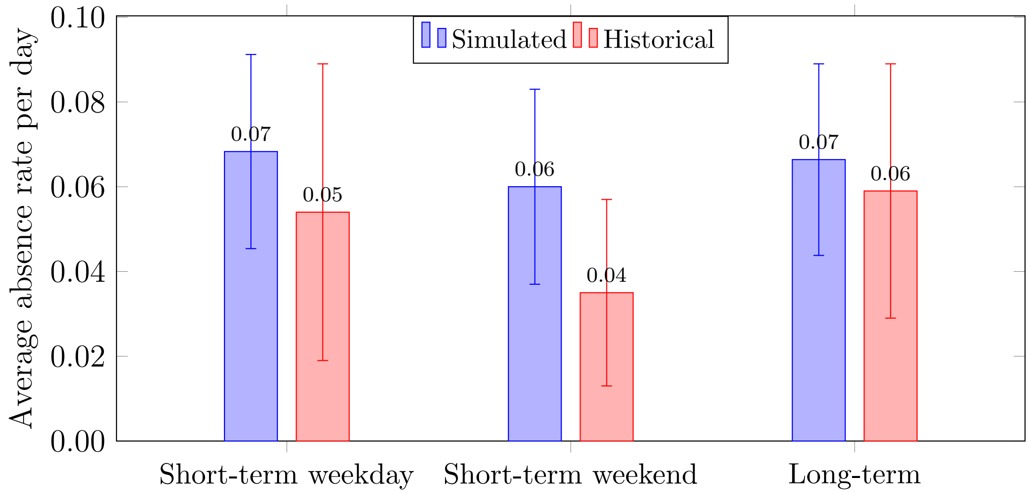

Figure 8.4: Comparison of the rate of employee absences each day for simulation results and historical values

The results from the absence simulations, which indicate whether each employee is non-absent, short-term absent, or long-term absent for all days in the planning period, are used as inputs for all the upcoming rescheduling instances in Section 8.4.

8.3.4 Availability Simulation

As explained in Section 6.4, there is not sufficient data available at DNIC to mathematically calculate the values of the simulated parameters in the rescheduling model. Therefore, all values are set based on conversations with the scheduling manager at DNIC.

Uncertainty in acceptance of extra shifts According to the scheduling manager, approximately one fourth of all employees never accept extra shifts. During weekdays, a large proportion of the remaining employees accept them, while it is more difficult to call in extra staff during weekends. Another important dimension is how many days in advance the employees are asked to accept the extra shifts. It is approximately twice as difficult to call in someone on short notice as on long notice. For simplicity, the probabilities, which are presented in Table 8.13, are assumed equal for all employees.

Table 8.13: Probabilities for accepting an extra shift depending on how far in advance the notice is given

	Short notice	**Long notice**
Extra shift weekday	0.30	0.60
Extra shift weekend	0.10	0.20

Uncertainty in requests for extra shifts For weekdays, $2-3$ requests per shift is the most common, with a variation between 0 and 6. For weekdays, 1 request per shift is the most common, with a variation of 1 in both directions. The final probabilities used are displayed in Figure 8.5.

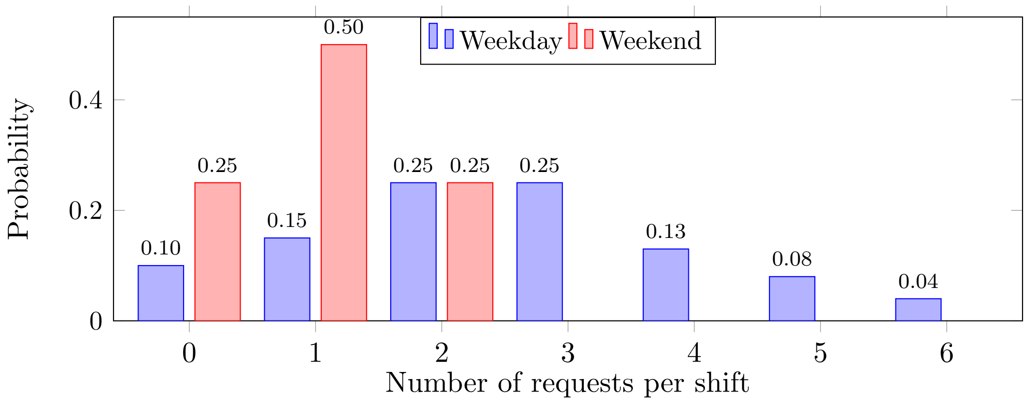

Figure 8.5: Probability distribution for the number of requests for extra shifts

Uncertainty in requests for exchanges The probability of an employee having requested an exchange is set in a similar manner. The scheduling manager estimates that exchanges happen a couple of times per month, but not as often as once per week. There is no distinction between weekdays and weekends. Based on this, we assume that the probability that any employee with an off day scheduled has requested an exchange each day is 10%.

The results from the availability simulations are used as inputs for all the upcoming rescheduling instances in Section 8.4.

8.4 Data in the Rescheduling Instances

In order to run an instance of the rescheduling model, an initial schedule is required as input. The size of the problem therefore depends on the size of the initial schedule. Using the scheduling instances from Section 8.2, the number of employees is 117, and the maximum number of days considered is 105. Due to requirements regarding the duration of the pre- and post-period, the number of days rescheduled is 77.

The rescheduling base case, `r0` is the reactive strategy which is most similar to how the real-life rescheduling at DNIC is performed. The same parameter values and settings are used in `r0` and `r1`. The difference is that some variables and parameters are removed from the model in `r1`, as defined in Section 7.6.

General settings The rescheduling problem is solved daily, with a rolling time horizon and several simulated values. Therefore, the elements and sizes of many of the sets may vary from day to day. The settings that remain constant in all problem instances are summarized in Table 8.14.

Table 8.14: Key settings in the rescheduling problem

Setting	Value	Description
$\lvert\mathcal{T}^{PRE}\rvert$	7	Days in pre-period
$\lvert\mathcal{T}^{POST}\rvert$	8	Days in post-period
H	35.5	Work hours per week for full-time employees
$\underline{M}^B$	2	Size of buffer of employees required to allow swaps or exchanges

Parameters H_{nk}^{MAX}, stating how many extra hours employee n can work week k without incurring overtime pay, depend on both the number of hours initially assigned and the actions taken. If the employee is assigned less than H work hours in week k, the employee can work extra or double shifts until the sum of weekly hours reaches H without receiving any overtime pay. New actions taken must therefore be included in the calculation of H_{nk}^{MAX}. If the employee initially is assigned more work hours than this, overtime will be payed for all additional shifts that week.

Replanning Period The set of days in which changes to the schedule are allowed to be made is denoted as the replanning period, $\mathcal{T}^R$. According to the scheduling manager at DNIC the length of this period varies in real-life, although the the most common length is four days. The length of the replanning period affects the sets of days triggering a swap on short and long notice, $\mathcal{T}^{SN}$ and $\mathcal{T}^{LN}$, respectively, as indicated in Table 8.15. Swaps on long notice are required to happen at least two days in advance, such that the set is empty when the replanning period contains just the current and the next day.

Table 8.15: Length of replanning period and its impact on $\mathcal{T}^{SN}$ and $\mathcal{T}^{LN}$

$\lvert\mathcal{T}^R\rvert$	**1**	**2**	**3**	**4**	**5**	**6**	**7**
$\mathcal{T}^{SN}$	$\{0\}$	$\{0,1\}$	$\{0,1\}$	$\{0,1\}$	$\{0,1\}$	$\{0,1\}$	$\{0,1\}$
$\mathcal{T}^{LN}$	$\{\}$	$\{\}$	$\{2\}$	$\{2,3\}$	$\{2\dots4\}$	$\{2\dots5\}$	$\{2\dots6\}$

Minimum Demand The minimum demand in the online operational problem is displayed in Table 8.16. The limits are lower than in the offline operational problem for all shifts, indicating that the department always adds a buffer of employees on the shifts in the initial schedules.

Table 8.16: Minimum demand limits for online operational staffing

Parameter	**Weekday**	**Weekend**	**Description**
$\underline{D}_{Dt}^{RE}$	14	12	Demand on the Day shift
$\underline{D}_{Et}^{RE}$	12	12	Demand on the Evening shift
$\underline{D}_{Nt}^{RE}$	12	12	Demand on the Night shift

Table 8.17 indicates the minimum demand per skill in the rescheduling problem. As the table shows, the limits do not distinguish between days of the week or shift types. There must always be at least five employees with respirator skills present to take care of the patients with the most severe conditions, where at least one of them must have Emergency skills. Recall from Section 2.3.2 that employees with Intensive Care skills or higher knows how to use respirators.

Table 8.17: Minimum demand per skill category per shift in the rescheduling problem

Parameter	Value	Description
$\underline{D}_{1st}^{RE}$	1	Demand for Emergency skills
$\underline{D}_{2st}^{RE}$	4	Demand for Intensive Care skills
$\underline{D}_{3st}^{RE}$	2	Demand for Monitoring skills
$\underline{D}_{4st}^{RE}$	0	Demand for Assistant skills

Real demand The real demand on the current day is based on the number of patients per level simulated by the model in Section 8.3.2. Using these values the number of patients in the upcoming days is estimated based on the expected values of the number of patients, which are obtained using the transition probabilities in Table 8.11.

D_{st}^{SIM} and D_{cst}^{SIM} for the replanning period are then calculated using the values for the estimated number of nurses required to treat one patient at each level displayed in Table 8.18. Additionally, there is always demand for one coordinator and one employee with Emergency skills on all shifts. An example of how to calculate demand with a given patient mix is included in Appendix B.4.

Table 8.18: Estimated need for nurses per patient per level (Halsteinli, 2017) and the skills required to treat the patients at each level, where skill 2 is Intensive Care and 3 is Monitoring

Level	1	2	3	4	5
Nurses per patient	0.3	0.4	0.7	1.0	1.5
Minimum skill	3	3	3	2	2

Availability The set of available employees are based on outputs from the absence and availability simulation models. The set of working nurses on day t is the set of non-absent nurses with a work shift scheduled that day, either in the initial schedule or as a result of a schedule change made in a previous replanning problem. The set of off nurses on day t consists of the non-absent nurses with an off shift scheduled who will accept a work shift that day, with the probabilities given in Table 8.13. The set of available nurses on day t is the union of these sets.

Weights in objective function The values of the weights in the objective function of the rescheduling problem are indicated in Table 8.19. The weights are divided into three types; actions, rules and demand. Each weight is composed of two parts; a cost penalty and an inconvenience penalty.

Table 8.19: Values of the penalties in the objective function of the rescheduling problem, where each weight consists of a cost term and an inconvenience term

Weight	Value	Cost	Inconvenience	Type	Penalty description
W^{EXC}	1	0	1	Action	Exchange
W^{REX}	5	4	1	Action	Extra, requested shift
W^{NEX}	7	4	3	Action	Extra, non-requested shift
W^{LN}	3	2	1	Action	Swap, long notice
W^{SN}	7	4	3	Action	Swap, short notice
$W^{DB,(2)}$	14	8	6	Action	Double shift with break
$W^{DB,(1)}$	17	7	10	Action	Double shift without break
W^{CO}	1.5	0.5	1	Rules	One hour of overtime work
W^{CW}	7	0	7	Rules	Too many consecutive shifts
W^{N}	12	0	12	Rules	Too many consecutive nights
W_0^C	100	0	100	Demand	Not meeting skill demand
W_0^D	100	0	100	Demand	Not meeting overall demand

Assuming a shift duration of 8 hours ($\frac{24}{3} = 8$), the values of the cost terms reflect the number of extra, unplanned hours worked, plus, for some of the actions, the number of hours which are connected to an additional financial compensation. The sum of these values are multiplied by the cost term of one hour of overtime, W^{CO}. W^{CO} reflects the overtime incurred by the number of hours actually worked, and not the additional financial compensation incurred by some of the actions. The total compensation for each type of action was elaborated on in Section 2.3.5. The size of the additional financial compensations are, in a simplified version, based on governmental rules and regulations as well as agreements with the trade union (Unio, 2018).

Inconvenience concerns how the action or violation of rules and preferences affects the satisfaction and health of the employees and patients. The inconvenience terms are set based on conversations with the scheduling manager. The least inconvenient actions are the actions which do not add any additional shifts to the employees' schedules, but only move the existing shifts to another time. Therefore, swaps and exchanges are given low inconvenience scores. However, a swap on short notice demands that the employees are flexible with their short-term plans, which is inconvenient, and this is reflected in a higher inconvenience term than swaps on long notice.

Working extra shifts is generally not considered too inconvenient if it does not lead to violations of the rules regarding sufficient rest between shifts. The inconvenience is lowest when the employee has requested the extra shift. Working double shifts violates the rules regarding sufficient rest between shifts, where working two consecutive shifts without a break in between (type 1) is the worst violation. These weights are therefore given a high inconvenience value relative to the other weights for actions.

Working more consecutive shifts, and specially nights, than recommended by the governmental regulations may affect the health and work performance of the employees negatively. We therefore put the corresponding inconvenience values relatively high. The terms aid in making sure that each employee is not assigned unreasonably many schedule changes. The same effect can be incurred by working overtime, but depends on the accumulated number of overtime hours. Therefore, the inconvenience per isolated hour is considered being relatively low.

Not meeting demand is the most severe violation of the department's preferences and guidelines regarding rescheduling, as this could put the patients' health at risk. We therefore assign high penalties to the parameters penalizing understaffing in the objective function to avoid the corresponding slack variables taking on positive values whenever possible.

Parameters W_t^D and W_t^C are calculated using Equations (8.2). The weights gradually decrease by 10 % for each day we move into the replanning period. This is to reflect that the demand becomes more uncertain the longer into the future we see, and in turn that the model should not take unreasonable actions to cover potentially unmet demand too many days in advance.

$$W_t = 0.90 \times W_{(t-1)} = 0.90^t \times W_0 \qquad t \in \mathcal{T}^R \backslash \{0\} \tag{8.2}$$

8.5 Technical Analysis

In this section we elaborate on some key statistics of the scheduling instances, as well as the run time and optimality gap of the scheduling and rescheduling instances. The thorough robustness analysis of all instances is left for Chapter 9.

8.5.1 Testing the Scheduling Instances

Table 8.20 contains key statistics, run time and optimality gap of instances `s0` to `s4`. To obtain comparable results, measures were made to ensure that the number of shifts scheduled did not deviate significantly between instance `s0` and instances `s1` to `s3`. The number of shifts assigned in instance `s4` is naturally lower, as the employees working extra weekends in this instance have to work less hours than in the other instances.

The number of desirable patterns assigned are expressed as a percentage of the total number of possible shift patterns assigned in Table 8.20. It could be that more desirable patterns had been assigned if this objective was given a higher reward relative to the reward for granting requested shifts.

As requiring employees to rank down in skills on a shift is heavier penalized in `s1` than in the other instances, it makes sense that the number of employees ranking down is lowest in this instance. The fact that there are fewer employees ranking down in `s4` than `s0` indicates that at least some of the occurrences in the base case happen during weekends, when the problem is the most constrained. This could also be why employees ranking down were not eliminated in `s1`, as nothing was done to improve the robustness during weekends in this instance.

There was a large difference in the time required to produce a solution with an optimality gap lower than 1 % for the five instances. As seen in Table 8.20, neither `s1` nor `s2` had found a solution with a gap lower than 10 % within the first three hours. These instances were therefore run for a full day, and even then `s2` did not get below 1 %. We believe that the main reason for `s1` being slower than the other instances is that the reward for assigning the two first buffer employees is higher than the reward for assigning requested shifts. When the reward for assigning requested shifts is the highest weight, it seems easy to the model to find a satisfying solution as there is a clear preference as to which employees should work which shift. However, other weights dominate the request reward in `s1`. This makes the search space much larger, especially since we made no measures to speed up the search through e.g. symmetry breaking constraints. In the case of `s2`, our best guess is that the model took long to find a solution because there is a large amount of equally good placements of the ghost shifts, making the initial search space huge.

Table 8.20: Key statistics, run time and optimality gap for instances `s0`-`s4`

Measure	`s0`	`s1`	`s2`	`s3`	`s4`
Requests granted [%]	88.46	86.03	87.67	88.21	87.81
Work shifts	5550	5575	5490	5574	5459
Desirable patterns [%]	62.9	49.7	62.1	60.2	60.7
Rank-downs in skills	31	11	35	30	22
First solution under 10 % gap [s]	76	10286	12792	140	2655
Run time [s]	3600	86400	86400	3600	3600
Optimality gap [%]	0.27	0.55	1.56	0.31	0.37

8.5.2 Testing the Rescheduling Instances

In this section the run time and problem size of the rescheduling instances `r0` and `r1` when the replanning period varies between one and seven days is analyzed. The upper limit of seven days is chosen because it is very difficult to predict the real demand further into the future. `s0` is used as input schedule for all instances, meaning that 14 instances are analyzed in total. With 200 simulation runs and 77 days in the rescheduling period, the rescheduling model must be called 15400 times to perform the full rescheduling process on one schedule. Whether actions are taken during each call depends on the balance between demand and supply. Consequently, it could be that the model chooses to do nothing on a given day. All 15400 problems for all 14 instances were solved to optimality.

The average number of variables and constraints for each instance is indicated in Figure 8.6. The problem size is clearly dependent of the duration of the replanning period, which is expected as a longer replanning period corresponds to more possible actions. The size of `r1` is naturally smaller than the size of `r0`, as less variables and constraints are needed when some actions are prohibited.

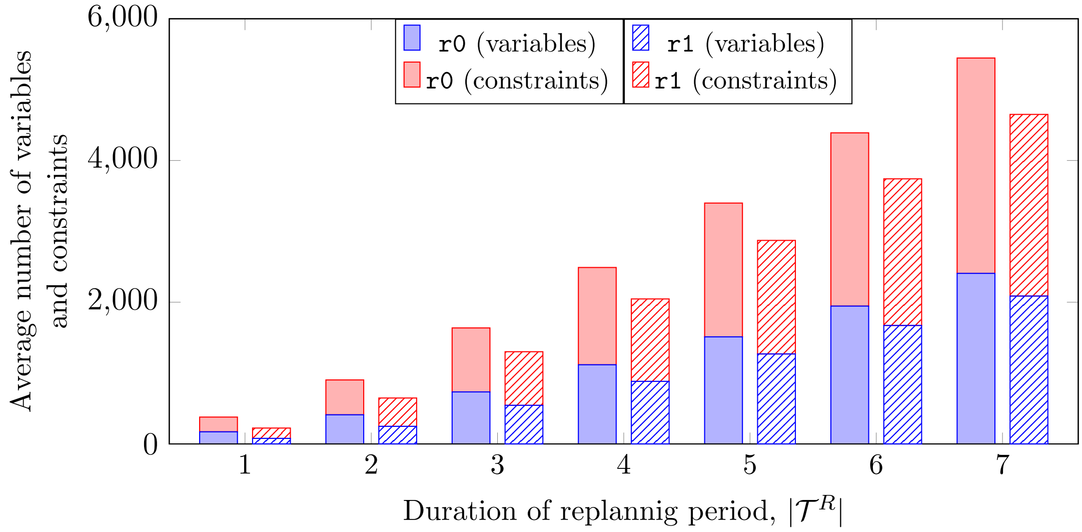

Figure 8.6: Average number of variables and constraints after presolve per iteration over the entire planning period and 200 simulations

In Figure 8.7, the average run time for all rescheduling instances are plotted against the duration of the replanning period. As displayed, the run time for each daily problem is very low. The fact that the problem size of `r1` is smaller than `r0` is reflected in a lower run time. There is an increase in both run time and standard deviation for longer replanning periods, which is consistent with the increasing problem size in Figure 8.6.

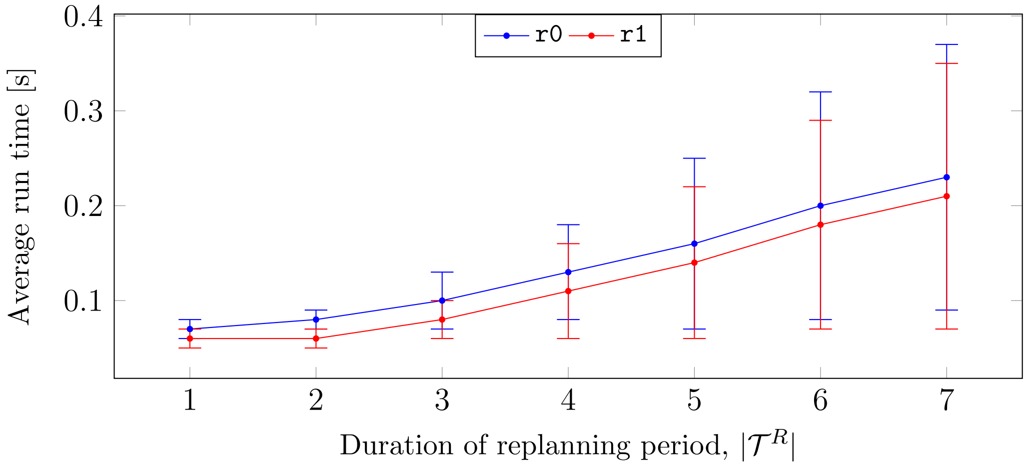

Figure 8.7: Average run time per iteration over the entire planning period and 200 simulations in addition to error bars displaying the spread one standard deviation from the mean

Chapter 9

Computational Study

In this chapter, we combine the scheduling and rescheduling instances from the case study and evaluate the results using the robustness measures from Section 4.4. We start by presenting the instances to be tested in Section 9.1, before the quantitative robustness analysis is performed in Section 9.2.

9.1 Test Phases

The test instances used in the computational study are explained in Table 9.1. The type of scheduling instance, the type of rescheduling instance and the length of the replanning period give the name of each test instance. For example, a test instance consisting of the ghost scheduling instance, `s2`, and the base case rescheduling instance, `r0`, where the duration of the replanning period is three days, is named `s2_r0_3`.

The testing is split into four phases, where each phase has a specific goal. During the first phase, we test the base case instance produced by the scheduling model on `r0` with a replanning period varying from one to seven days. The goal is to identify the replanning period that results in the best actions during the online operational phase. The optimal duration of replanning period is then set in stone for the remaining test phases. In phase 2, all scheduling instances from Section 8.2 are tested and compared to the base case schedule with the purpose of identifying the most promising combinations of proactive strategies. These strategies are then combined to generate new and possibly even more robust schedules, which are tested in phase 3. The goal of the third phase is to reveal the most robust strategy. During the fourth and final test phase, we learn how the flexibility in the online operational phase is affected if no governmental rules can be violated.

We use the settings for the demand, absence, and availability simulations from Section 8.3 for all instances.

Table 9.1: The four test phases of the computational study, their respective goals as well as the test instances included in each phase

Test phase	**Goal**	**Name**	**Proactive strategy**	**Reactive strategy**	**Inst-ances**
1	Find optimal duration of replanning period, t^*	`s0_r0_1-7`	None	Vary length of replanning period	7
2	Identify three most promising combinations of proactive strategies, `s5a-c`	`s1_r0_t`*	Even buffer of surplus work hours	Use t^*	1
		`s2_r0_t`*	Introduce ghost shift	Use t^*	1
		`s3_r0_t`*	Less contribution to demand from high-risk group	Use t^*	1
		`s4_r0_t`*	Schedule extra weekends for some employees	Use t^*	1
3	Find most robust combination of proactive strategies, `s5`*	`s5a-c_r0_t`*	Three most promising combinations of `s1-s4`	Use t^*	3
4	Learn about necessity of violating rules	`s0_r1_t`*	None	No governmental violations	1
		`s5`*`_r1_t`*	Best combination of strategies	No governmental violations	1

9.2 Quantitative Robustness Analysis

The four phases in Table 9.1 are analyzed in this section. For each phase, we assess the stability and flexibility of the corresponding instances, before we make a conclusion regarding the goal of the phase. When the analysis of the fourth phase is completed, we leave some final remarks on the validity of the results obtained.

9.2.1 Phase 1

The goal of phase 1 is to determine the optimal duration of the replanning period by evaluating the results of running instances `s0_r0_1-7`. With the replanning period being a part of the reactive strategy, it only makes sense to consider the flexibility measures, as the same input schedule results in identical stability for all replanning periods.

The flexibility of the instances in phase 1 are displayed in Table 9.2. The first part of the instance names, `s0_r0`, a left out of the table for simplicity. There was no understaffing after rescheduling for any of the instances, and the measures of frequency and severity of understaffing after rescheduling are therefore also omitted. In the table, the actions are not multiplied by their corresponding weighing parameters except for in the objective.

Table 9.2: Flexibility of instances `s0_r0_1-7`, where the lowest score on each measure is highlighted by green and the highest by red. Each measure is the total number of occurrences of the corresponding action during the full planning period of 77 days, averaged over 200 simulations

	Duration of replanning period						
Measure	**1**	**2**	**3**	**4**	**5**	**6**	**7**
Overtime hours	114.86	102.81	80.55	77.18	75.89	77.01	74.90
Consecutive shifts	13.48	6.25	4.18	3.55	2.54	1.45	0.74
Consecutive nights	0.13	0.17	0.12	0.06	0.04	0.02	0.02
Extra shifts	140.06	168.69	155.61	136.53	118.77	106.57	100.09
Exchanges	0.00	0.01	0.16	0.54	1.23	1.64	2.08
Swaps	14.83	21.34	50.98	78.37	103.74	120.50	130.91
Double shifts	26.44	13.14	11.03	11.01	10.80	10.88	10.79
All actions	181.33	203.17	217.79	226.45	234.54	239.58	243.86
Objective	1716.65	1659.23	1628.56	1618.83	1582.90	1556.28	1538.38

The longest replanning period has the lowest objective function value, while the shortest replanning period has the worst. The main reason for the poor performance when only considering the current day is that more consecutive shifts, consecutive nights, double shifts and overtime work is assigned. This is consistent with the fact that it is more difficult to call in extra staff on short notice, and that assigning double shifts to some employees may be left as the only alternative. Both working double shifts and too many consecutive shifts and nights are considered severe violations of governmental rules, and are heavily penalized in the objective function.

When increasing the duration of the replanning period from one day to seven, the overtime and number of consecutive and double shifts decrease significantly. Also, the number of swaps, which is one of the preferred actions at DNIC, reflected by the relatively low corresponding assignment penalties, increases at the expense of fewer of the more penalized extra shifts. This is not a surprising result; the more days considered when replanning, the more possibilities for swaps become available.

There are no significant differences in the number of overtime hours, consecutive shifts and nights and double shifts when increasing the replanning period from three to seven days. While more actions are taken when the replanning period increases, the actions are less expensive in terms of costs and inconvenience, resulting in an overall lower objective function value. Although more actions are taken when the replanning period is seven days, recall that a swap does not mean that any additional shifts are worked, just that a planned work shift is changed to another time. Therefore, the actual number of *additional* shifts worked is actually lowest with the longest replanning period, indicating a better utilization of the resources available.

As the number of violations of the rules intended to ensure sufficient rest are approximately the same when the replanning period is between three and seven days, we do not consider the current practice of replanning up to four days ahead as a very bad practice either. In addition, in practice it can be difficult to take a week of shifts into account when performing the rescheduling manually, as there is simply too much information to process within the short time during which it is preferred that the problem is solved.

We conclude phase 1 by stating that the optimal replanning period for automatic rescheduling, t^*, is seven days. However, it seems sufficient to consider only four days when the rescheduling is performed manually, as the differences are not striking.

9.2.2 Phase 2

In the second phase, we put $t^* = 7$ and run the test instances `s0_r0_7` - `s4_r0_7`. The difference between these instances is the underlying scheduling model. For simplicity, we will mainly refer to the instances by the name of their respective scheduling instance for the remainder of the phase.

Stability

In Figure 9.1 we split the frequency of understaffing for the various instances into weekday and weekend shortage. The frequency is the number of understaffed shifts divided by all shifts. There is a large difference in the frequency of understaffing during weekdays and weekends. This is a result of the staffing level during weekdays usually exceeding the staffing during weekends, while the demand fluctuates independently of which day it is.

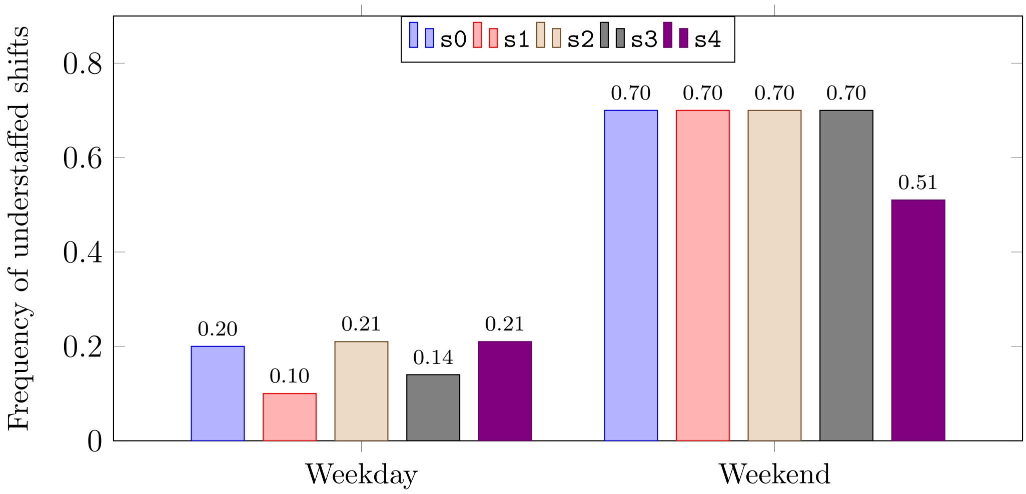

Figure 9.1: Frequency of understaffing for instances `s0-s4_r0_7`

Three instances positively stand out in Figure 9.1; `s1`, `s3` and `s4`. First, the proactive strategies modelled in `s1` and `s3` seem to succeed in making the schedules more stable during weekdays, as the frequency of understaffing compared to `s0` is reduced from 20 % to 10 % and 14 %, respectively. `s1` performs better than `s3` in terms of understaffing, suggesting that it is more efficient to evenly distribute the buffer on all shifts than to only buffer up on shifts where employees vulnerable to absence are scheduled.

During weekends, none of the three first strategies makes a difference, confirming that the staffing problem is too constrained during weekends. However, the fourth strategy, which aims at reducing the weekend understaffing, succeeds in just that. Compared to `s0`, the weekend understaffing is reduced from 70 % to 51 %, while increasing the understaffing during weekdays by only 1 percentage point. The latter point is interesting, as far less work shifts are assigned in `s4` than in `s0` due to the policy of offering extra off days to the employees who work extra weekends.

Instance `s2` does not improve in terms of stability, which is quite as expected, as the optimal placements of off shifts solely are intended to increase the flexibility during the online operational phase.

There are no major differences between the test instances when it comes to the severity of understaffing, which is the number of employees in shortage on the shifts with insufficient staffing. During weekdays, `s1` and `s3` have the lowest average severity with a value of 1.9 employees, while the remaining instances have approximately 2.0 employees in shortage. For weekends, the average severity is 2.1 for `s4` and 2.3 for the other instances. Thus, the instances with the lowest frequency of understaffing for weekdays or weekends are also the instances with the lowest severity during the corresponding time period. This makes sense, as ensuring a certain buffer of employees on all shifts means that the variance in the number of employees scheduled should be reduced, with less shifts that are very overstaffed or very understaffed as compared to the case when surplus work hours are more freely assigned. This implies a much better utilization of resources, as the number of work shifts scheduled in instance `s0` to `s3` are approximately the same, and `s4` proving good stability although 91 less shifts are scheduled than in the base case instance.

Flexibility

The flexibility of the test instances is compared in Table 9.3. The time resolution describes on which day an action was taken, or on which day a pattern of too many consecutive shifts occurred as a result of actions taken, and *not* which day that actually triggered the action. This is particularly important to keep in mind when evaluating the results for the weekday resolution, as many of the actions taken during weekdays are consequences of expected understaffing during the weekend. As in phase 1, there was no understaffing after rescheduling for any instances.

Table 9.3: Flexibility of instances `s0-s4_r0_7`, where the lowest score on each measure is highlighted by green and the highest by red. Each measure is the total number of occurrences of the corresponding action during the full planning period of 77 days, averaged over 200 simulations

Time resolution	**Measure**	`s0_r0_7`	`s1_r0_7`	`s2_r0_7`	`s3_r0_7`	`s4_r0_7`
Weekday	Overtime hours	22.91	22.23	21.90	23.56	17.81
	Consecutive shifts	0.36	0.18	0.47	0.34	0.24
	Consecutive nights	0.01	0.03	0.02	0.02	0.03
	Extra shifts	62.92	46.52	70.54	52.41	63.17
	Swaps	85.98	71.15	84.78	75.57	61.05
	Exchanges	1.28	1.01	1.30	1.12	1.31
	Double shifts	0.22	0.12	0.19	0.17	0.38
	All actions	150.41	118.79	156.82	129.27	125.91
	Objective	820.03	633.21	851.18	701.61	712.86
Weekend	Overtime hours	51.99	57.14	48.78	54.52	33.81
	Consecutive shifts	0.38	0.51	0.62	0.59	0.24
	Consecutive nights	0.01	0.00	0.01	0.00	0.00
	Extra shifts	37.17	32.33	39.03	34.17	34.42
	Swaps	44.93	41.73	43.20	41.83	27.83
	Exchanges	0.79	0.74	0.70	0.76	0.59
	Double shifts	10.57	11.04	11.25	10.86	7.09
	All actions	93.45	85.83	94.18	87.61	69.93
	Objective	718.35	694.04	721.19	700.88	539.85
Overall	All actions	243.86	204.62	250.99	216.88	195.83
	Objective	1538.38	1327.24	1572.37	1402.49	1252.72

For `s1`, about 40 less actions are required than in the base case instance, most likely due to the schedule being more stable during weekdays for `s1` than `s0`. This is also reflected in an objective function value that is significantly lower than for `s0`. The overtime occurred during weekends is slightly increased compared to the base case instance, which is interesting since the number of extra shifts is much lower. One possible explanation is that more double shifts are assigned during weekends in `s1`.

`s2` generally performs worse than the base case instance. However, it results in less overtime with more actions, with overtime being one of the aspects `s2` aims at improving. There is no increased performance in terms of consecutive nights, although the reason for this might be that none of the test instances has a problem with avoiding too many consecutive Night shifts either. Also, recall from Section 8.5 that 60 less shifts were scheduled with `s2` than `s0`, which could be the reason why there are more extra shifts and less swaps assigned in the former.

Instance `s3` triggers fewer actions than `s0`, which most likely is a result of being understaffed less frequently. Still, it results in slightly more overtime work than the base case instance, suggesting that the proactive strategy on which `s3` is based is not very flexible.

`s4` performs best in terms of almost all measures, indicating that solving the weekend bottleneck is the key to obtaining more robust schedules at DNIC. The weekday results of `s4` give an impression of increased performance during weekdays as well. Recall that this is likely to be a consequence of a need for making fewer actions during the weekdays as a result of expecting to be understaffed with a lower frequency and severity during the weekends.

Another interesting result is that much less overtime is assigned in `s4` than in the other instances. This could be due to the assumption made in Section 6.3.1, stating that the number of employees who accept to work extra shifts during weekends is only 10 %. When there are this few employees to choose from, it is not surprising if overtime work is difficult to avoid during weekends in the other instances. Another explanation is that it in general is easier to find employees who can work extra shifts without overtime in `s4` as there are far less work shifts scheduled in this instance. For the same reason, less swaps are assigned in `s4` than in the other instances.

Conclusion

When comparing all results we see that no instances strictly Pareto dominate any other instances, indicating that in the end the evaluation of which strategies are the best are also dependent on human judgment. However, in the case of instances discussed in this section, `s4` is superior to the other instances in most of the flexibility measures. Although it does not improve the stability during weekdays compared to `s0`, the weekend stability was improved from 70 % to 51 %. Also, `s4` had the best improvement in objective function value, with a 19 % reduction compared to the base case instance. We therefore consider the strategy of allowing to trade extra weekend work for extra off shifts as the most promising single standing proactive strategy.

Although `s4` performs the best, phase 2 has revealed that several of the other strategies lead to schedules more robust than the base case instance. While `s1` is the most stable instance during weekdays, `s4` outperforms the other instances during weekends. It seems promising to exploit the benefits of both of these to see if it is possible to increase the robustness any further.

Additionally, `s3` is understaffed less frequently than `s0` during weekdays, while still assigning more overtime work. It seems that `s3` is a stable, but not flexible schedule. A combination of `s3` and `s2` could be beneficial, as the latter seems like it could be more flexible than `s0` if equally many shifts are scheduled.

Finally, we combine all the instances which had a clear improvement in robustness compared to the base case instance; `s1`, `s3`, and `s4`. Thus, the three new instances to test in phase 3 are the following:

$$
\begin{aligned}
\texttt{s5a} &= \texttt{s1} + \texttt{s4} \\
\texttt{s5b} &= \texttt{s2} + \texttt{s3} \\
\texttt{s5c} &= \texttt{s1} + \texttt{s3} + \texttt{s4}
\end{aligned}
$$

9.2.3 Phase 3

During phase 3, the goal is to identify the strategy or strategy combination leading to the most robust schedules. The candidates are `s4`, which was the overall most robust single standing strategy, and instances `s5a-c`, which were defined in the end of phase 2.

A presentation of the key model settings used for instances s5a-c, as well as a technical analysis, is provided in Appendix D.

Stability

The frequency of understaffing for instance s4 and s5a-c are displayed in Figure 9.2, which is designed equally as Figure 9.1 in phase 2.

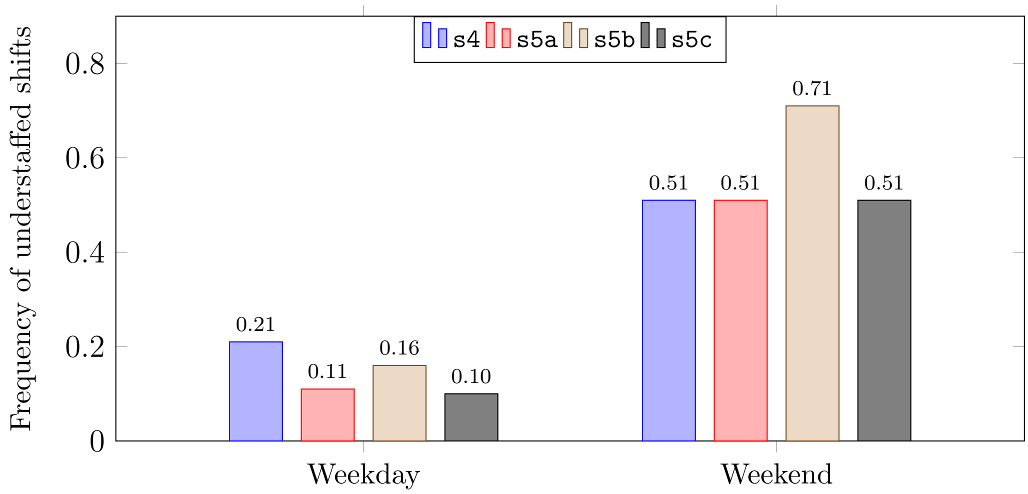

Figure 9.2: Frequency of understaffing for instances s4_r0_7 and s5a-c_r0_7

s5a and s5c manage to keep the understaffing as low as s1 during weekdays and s4 during weekends, successfully improving the stability significantly on both weekdays and weekends compared to the base case. s5b is more often understaffed compared to s3 and less frequently in shortage compared to s2, while still outperforming the base case instance by 4 percentage points.

s5a and s5c slightly improve the severity of understaffing during weekdays compared to s1, with an average severity of 1.8 employees in shortage when a shift is understaffed for the former instances and 1.9 for s1. During weekends, the two instances have the same severity as s4, namely 2.1. The severity of s5b and s3 is the same all week, with 1.9 during weekdays and 2.3 during weekends.

Flexibility

Table 9.4 contains the values of the flexibility measures for the instances of phase 3. As in phase 1 and 2, there was no understaffing after rescheduling for any instances.

Table 9.4: Flexibility of instances `s4_r0_7` and `s5a-c_r0_7`, where the lowest score on each measure is highlighted by green and the highest by red. Each measure is the total number of occurrences of the corresponding action during the full planning period of 77 days, averaged over 200 simulations

Time resolution	**Measure**	`s4_r0_7`	`s5a_r0_7`	`s5b_r0_7`	`s5c_r0_7`
Weekday	Overtime hours	17.81	14.79	22.40	14.48
	Consecutive shifts	0.24	0.17	0.49	0.16
	Consecutive nights	0.03	0.01	0.05	0.03
	Extra shifts	63.17	42.90	61.78	40.46
	Swaps	61.05	48.20	79.24	46.00
	Exchanges	1.31	1.00	1.19	1.04
	Double shifts	0.38	0.13	0.12	0.10
	All actions	125.91	92.22	142.32	87.59
	Objective	712.86	506.61	770.14	480.46
Weekend	Overtime hours	33.81	35.12	56.73	34.29
	Consecutive shifts	0.24	0.31	0.67	0.22
	Consecutive nights	0.00	0.01	0.00	0.01
	Extra shifts	34.42	28.66	38.22	27.37
	Swaps	27.83	24.44	40.93	23.80
	Exchanges	0.59	0.53	0.71	0.55
	Double shifts	7.09	7.22	11.09	7.11
	All actions	69.93	60.85	90.95	58.83
	Objective	539.85	496.14	721.81	481.11
Overall	All actions	195.83	153.07	233.26	146.42
	Objective	1252.72	1002.75	1491.94	961.56

`s5a` generally performs better than `s4`, with the improvement during weekdays being the most noticeable. The number of actions is considerably reduced, and a larger proportion of the actions are swaps compared to extra shifts. As stated in Section 9.2.1, a swap is a more efficient action than an extra shift. This proposes that `s5a` is more flexible than `s4`.

`s5b` ends up somewhere in between the results of `s2` and `s3`, which obliviously is worse than `s4`. Compared to `s3`, the overtime assigned during weekdays is the only measure that is improved. The number of actions required is increased, while the overtime is approximately the same. While it is possible to interpret this as a somewhat more efficient action allocation, the overall performance decreased.

`s5c` performs slightly better than and very similar to `s5a`, meaning that the instance performs well during both weekdays and weekends. There are no particular measures in which `s5c` outperforms `s5a`; the performance is generally slightly better.

Conclusion

It is clear that `s1` and `s4` have great synergies, as both `s5a` and `s5c` significantly increases the performance. More staffing during weekends, combined with a more efficient use of the surplus work hours on weekdays, results in schedules that are more robust during the whole week. `s3` boosts the performance of `s1` and `s4` even further, with a reduction in objective function value compared to `s0` of 37 %. This makes `s5c` the most robust combination of strategies, giving that $\texttt{s5}^* = \texttt{s5c} = \texttt{s1} + \texttt{s3} + \texttt{s4}$.

`s5b` does not succeed in reducing the overall overtime such as intended. `s3` performs better than `s5b` both in terms of stability and flexibility. The attempt of making the instances `s2` and `s3` complement each other clearly failed.

9.2.4 Phase 4

In phase 4, we change the reactive strategy, and use strategy `r1` to perform rescheduling of instance `s0` and $\texttt{s5}^*$ with a replanning period of $t^* = 7$. Recall that with this reactive strategy, it is not allowed to assign double shifts or too many consecutive shifts or consecutive nights. For the same reasons as in phase 1, it only makes sense to evaluate the results in terms of flexibility, which is conducted in Table 9.5.

Table 9.5: Flexibility of instances `s0_r1_7` and `s5*_r1_7`. The _post measures are averaged over the full planning period and 200 simulations, while the remaining measures are the total number of occurrences of the corresponding action during the full planning period of 77 days, averaged over 200 simulations

Time resolution	Measure	s0_r1_7	s5*_r1_7
Weekday	Overtime hours	24.04	18.83
	Extra shifts	62.67	39.97
	Swaps	87.92	49.94
	Exchanges	1.25	0.99
	Frequency_post	0.00	0.00
	Severity_post	0.00	0.00
	All actions	151.84	90.90
	Objective	848.78	564.85
Weekend	Overtime hours	43.56	32.71
	Extra shifts	38.10	27.27
	Swaps	44.45	25.70
	Exchanges	0.79	0.62
	Frequency_post	0.08	0.05
	Severity_post	1.77	1.67
	All actions	83.33	53.58
	Objective	1463.08	921.19
Overall	All actions	235.17	144.47
	Objective	2311.86	1486.05

The most important observation from Table 9.5 is that neither `s0_r1_7` nor `s5*_r1_7` succeed in meeting demand on all shifts. As seen from the frequency_post measure, DNIC is still understaffed on 8 % and 5 % of all weekend shifts after actions have been taken for `s0_r1_7` and `s5*_r1_7`, respectively. As the penalty for not meeting demand is very high, the reason for the weekend shortage when violations of governmental rules are not allowed, is that there simply are not enough employees who are willing to work unplanned shifts during weekends.

The number of actions taken are reduced in both of the `r1` instances compared to the corresponding `r0` instance, which makes sense because many of the actions taken in `s0_r0_7` and `s5*_r1_7` now are illegal. However, the objective function value is increased by a considerable amount because the demand is not met on all shifts, which is considered the worst violation of preferences of them all.

`s5*` still shows considerably improved results compared to `s0`, with the reduction in frequency_post during weekends being the most important contribution.

9.2.5 Validity of Results

We conclude the analysis by leaving some concluding remarks regarding the validity of the results.

The rescheduling model relies heavily on assumptions about the availability of employees. It could be that we have reduced the solution space too much, and that in reality there are more available employees during weekdays and weekends than what we assumed in Section 8.3.4. This can again affect the results of the robustness evaluation and which actions are taken. Although this might be true, we do not consider it too much of a problem, as the most robust strategies ideally should be able to cope with low availability rates.

The values of the weighting parameters are based on simplified cost estimates and a subjective perception of how inconvenient different actions and outcomes are. If the weights were changed, the results of running the rescheduling model on the different instances could have been very different. However, we have to the best of our ability tried to set the values such that they match the reality faced at DNIC. We therefore believe that the actions taken should fit well with the environment at the department.

Chapter 10

Concluding Remarks

The purpose of this Master's thesis was to analyze the robustness of nurse schedules in a real-life case at the Department of Neonatal Intensive Care (DNIC) at St. Olavs Hospital. To achieve this purpose, we have developed a system consisting of three components; a MIP scheduling model, models simulating demand and absence, and a MIP rescheduling model. The simulation and rescheduling models are used to assess the robustness of the schedules made by the scheduling model. As all three components are based on real-life data, we believe that this thesis is among the studies on the robustness of nurse or personnel schedules closest to reality to date.

The scheduling model was extended with four proactive strategies intended to improve the schedule robustness. Three out of four of these strategies were able to significantly improve the robustness, as compared to the base case model where no proactive strategies were used. The strategies can with benefit be used as rules of thumb in the real-life scheduling process:

1. *Allow employees to trade extra weekend work for extra off days.* There was a huge improvement in robustness when employees were allowed to trade extra weekend work for extra off days in the initial schedules. The policy change led to a more stable schedule during weekends, without any significant effects on the stability during weekdays although less work shifts where scheduled. The flexibility improvements were seen in a reduction of the overall rescheduling objective by 19 % compared to the base case. Weekend shifts are the lowest staffed shifts in the schedule today,

and we believe that accepting this policy change would be very beneficial.

2. *Assign surplus work hours evenly over all work shifts.* The employees at DNIC have more contracted work hours than what is required to meet the minimum demand on all shifts. By allocating these surplus hours evenly over all shifts, the schedule robustness was improved in terms of better stability during weekdays and a reduction of the rescheduling objective by 11 % compared to the base case, where they were assigned more arbitrarily.

3. *Consider the employees' vulnerability to absence when making the schedules.* Analysis of historical data revealed that some employees are more vulnerable to absence than others. Considering this in the scheduling process by adding a small buffer to the shifts where these employees are scheduled proved to improve the schedule robustness. This was seen in terms of a lower frequency of understaffing during weekdays and an improvement by 9 % in the rescheduling objective compared to the base case. This way, the employees are better protected from work overload, and the department is more prepared if the employees are absent.

The most robust schedule was obtained when the three strategies were combined, with an overall reduction of the rescheduling objective function value of 37 % compared to the base case instance and a much better stability on both weekdays and weekdays.

We also assessed how the flexibility of the schedules varied when we gradually increased the duration of the replanning period from one to seven days. A seven-day replanning period proved to be the most robust choice, decreasing the overall rescheduling objective by 10 % compared to a one-day period and by 5 % compared to a four-day period, which is the daily practice at DNIC. The main reason for the improvement when considering several days is a better utilization of resources by assigning more favorable actions such as swaps. However, in the technical analysis we saw that the size of the rescheduling problem increases with the number of days in the replanning period, making it difficult to manage all the options that become available when considering a whole week in practice.

Chapter 11

Future Research

Through the work in this thesis, we have uncovered several interesting areas for future research. These are connected to both practical and theoretical applications.

Several assumptions and simplifications were made in the mathematical formulations and implementations of both the scheduling and rescheduling models. For the tools to be used as support to the manual planning processes today, both would have to be developed further. For the scheduling model, holidays, personal inclinations and relevant fairness measures are interesting to include. When it comes to the rescheduling model, interesting aspects to incorporate include fairness, patterns and preferences. Examples of each of them are to ensure that the number of schedule changes are fairly distributed over all employees, that undesirable patterns are avoided, and that removing employees from shifts they requested in the initial schedule is penalized. It could also be interesting to study the effects of having a more dynamic replanning period, e.g. with a different duration depending on whether the day in question was a weekday or weekend day, and to test the effects of having a replanning period longer than a week.

If the models are developed further, we believe they have potential as management decision tools in multiple phases of the hierarchical planning process. On the tactical level they could be used to test the effects of policy changes, such as altering the staff size or changing the rules regarding weekend work, on the robustness of schedules before the policies are approved. On the offline operational level the scheduling model can aid the managers by quickly producing initial schedules, which the managers thereby can adapt as they see fit. The rescheduling model together with the simulation model can then be employed to test the robustness of the final schedule. If a rescheduling model can easily be integrated with existing IT systems in which absence, extra shift requests, the initial schedule and other relevant information is recorded, we believe it could aid in quickly identifying which employees to ask to change their schedules. It is likely that the model could process more information than the managers can in a short time span, which could lead to better decisions.

It is interesting to conduct a more detailed analysis on how the robustness potentially is affected by the various skill categories and in turn propose proactive strategies dedicated specifically to skills. For example, it might be that the proactive strategy of assigning extra work weekends could have increased the performance even further if only employees with the most critical skills were chosen. Studying schedule robustness on a shift level is another interesting extension of our work.

The underlying assumptions of the simulation models in this thesis were that employee absence is independent of demand, and that the condition of the patients was independent of the workload per nurse. We have also assumed no seasonal variations. An interesting extension of our work is to research whether these assumptions are valid or not. If they are not, developing proactive strategies taking these insights into account could further improve the robustness of the real-life schedules. For example, planning for seasonal variations by adding greater buffers during periods when the demand or absence is expected to be at peak levels could be one such approach.

Bibliography

Aickelin, U. and Dowsland, K. A. (2004). An indirect genetic algorithm for a nurse-scheduling problem. *Computers & Operations Research*, 31(5):761–778.

Andersen, E. (2018). 05531: Levendefødte, etter måned 1966 - 2017. https://www.ssb.no/statbank/table/05531.

Arbeids- og sosialdepartementet (2017). Lov om arbeidsmiljø, arbeidstid og stillingsvern mv. (arbeidsmiljøloven). https://lovdata.no/dokument/NL/lov/2005-06-17-62.

Azaiez, M. and Al Sharif, S. (2005). A 0-1 programming model for nurse scheduling. *Computers & Operations Research*, 32(1):491–507.

Bagheri, M., Devin, A. G., and Izanloo, A. (2016). An application of stochastic programming method for nurse scheduling problem in real word hospital. *Computers & Industrial Engineering*, 96(Supplement C):192 – 200.

Bard, J. F. and Purnomo, H. W. (2005a). Hospital-wide reactive scheduling of nurses with preference considerations. *IIE Transactions*, 37(7):589–608.

Bard, J. F. and Purnomo, H. W. (2005b). Preference scheduling for nurses using column generation. *European Journal of Operational Research*, 164(2):510 – 534.

Bäumelt, Z., Dvořák, J., Šcha, P., and Hanzálek, Z. (2016). A novel approach for nurse rerostering based on a parallel algorithm. *European Journal of Operational Research*, 251(2):624–639.

Burke, E. K., De Causmaecker, P., Berghe, G. V., and Van Landeghem, H. (2004). The State of the Art of Nurse Rostering. *Journal of Scheduling*, 7(6):441–499.

Clark, A. and Walker, H. (2011). Nurse rescheduling with shift preferences and minimal disruption. *Journal of Applied Operational Research*, 3(3):148–162.

Efron, B. and Hastie, T. (2016). *Computer Age Statistical Inference*. Cambride University Press.

Ernst, A., Jiang, H., Krishnamoorthy, M., and Sier, D. (2004). Staff scheduling and rostering: A review of applications, methods and models. *European Journal of Operational Research*, 153(1):3–27.

Glass, C. A. and Knight, R. A. (2010). The nurse rostering problem: A critical appraisal of the problem structure. *European Journal of Operational Research*, 202(1):379–389.

Halsteinli, V. (2017). Rapport Nyfødt intensiv: Driftsgjennomgang - organisering, aktivitet og bemanning.

Hans, E. W., van Houdenhoven, M., and Hulshof, P. J. H. (2012). *A Framework for Healthcare Planning and Control*, pages 303–320. Springer US, Boston, MA.

Harper, P. R., Powell, N. H., and Williams, J. E. (2010). Modelling the size and skill-mix of hospital nursing teams. *Journal of the Operational Research Society*, 61(5):768–779.

Helse Midt-Norge (2018). Styrets beretning og årsregnskap 2017. `https://ekstranett.helse-midt.no/1001/Foretaksmter/Forms/St%20Olavs%20Hospital%20%20alle%20dokumenter.aspx`.

Ingels, J. and Maenhout, B. (2015). The impact of reserve duties on the robustness of a personnel shift roster: An empirical investigation. *Computers & Operations Research*, 61:153 – 169.

Ingels, J. and Maenhout, B. (2017). Employee substitutability as a tool to improve the robustness in personnel scheduling. *OR Spectrum*, 39(3):623–658.

Ingels, J. and Maenhout, B. (2018). The impact of overtime as a time-based proactive scheduling and reactive allocation strategy on the robustness of a personnel shift roster. *Journal of Scheduling*, 21(2):143–165.

Ionescu, L. and Kliewer, N. (2011). Increasing flexibility of airline crew schedules. *Procedia-Social and Behavioral Sciences*, 20:1019–1028.

Jones, M. M. T. (2005). *Estimating Markov transition matrices using proportions data: an application to credit risk*. Number 5-219. International Monetary Fund.

Kitada, M. and Morizawa, K. (2010). A Heuristic Method in Nurse Rerostering Following a Sudden Absence of Nurses. In *Proceedings of the 11th Asia Pacific Industrial Engineering & Management Systems Conference*, Melaka.

Kitada, M. and Morizawa, K. (2013). A Heuristic Method for Nurse Rerostering Problem with a Sudden Absence for Several Consecutive Days. *International Journal of Emerging Technology and Advanced Engineering*, 3(11):353–361.

Klyve, K. K. and Beckmann, F. R. (2016). Optimisation-Based Nurse Scheduling for Real-Life Instances. Master's thesis, Norwegian University of Science and Technology.

Kokangul, A., Akcan, S., and Narli, M. (2017). Optimizing nurse capacity in a teaching hospital neonatal intensive care unit. *Health Care Management Science*, 20(2):276–285.

Lachin, J. M. (2016). Fallacies of last observation carried forward analyses. *Clinical Trials*, 13(2):161–168.

Lim, G. and Mobasher, A. (2011). Robust Nurse Scheduling Problem. *IIE Annual Conference.Proceedings*, pages 1–8.

Løyning, I. N. and Melby, L. M. H. (2017). Evaluating the Robustness of Optimization Generated Nurse Schedules Using Simulation. Project report, Norwegian University of Science and Technology.

Maenhout, B. and Vanhoucke, M. (2011). An evolutionary approach for the nurse rerostering problem. *Computers & Operations Research*, 38(10):1400 – 1411.

Maenhout, B. and Vanhoucke, M. (2013a). An Artificial Immune System Based Approach for Solving the Nurse Re-rostering Problem. In Middendorf, M. and Blum, C., editors, *Evolutionary Computation in Combinatorial Optimization*, pages 97–108, Berlin, Heidelberg. Springer Berlin Heidelberg.

Maenhout, B. and Vanhoucke, M. (2013b). Reconstructing nurse schedules: Computational insights in the problem size parameters. *Omega*, 41(5):903 – 918.

MedlinePlus (2017). Positive airway pressure treatment. `https://medlineplus.gov/ency/article/001916.htm`.

Mikalsen, I., Davis, P., and Øymar, K. (2016). High flow nasal cannula in children: a literature review. *Scandinavian Journal of Trauma, Resuscitation and Emergency Medicine*, 24:93.

Moz, M. and Pato, M. V. (2003). An integer multicommodity flow model applied to the rerostering of nurse schedules. *Annals of Operations Research*, 119(1-4):285–301.

Moz, M. and Pato, M. V. (2004). Solving the problem of rerostering nurse schedules with hard constraints: new multicommodity flow models. *Annals of Operations Research*, 128(1-4):179–197.

Moz, M. and Pato, M. V. (2007). A genetic algorithm approach to a nurse rerostering problem. *Computers & Operations Research*, 34(3):667–691.

Mutingi, M. and Mbohwa, C. (2017). The nurse rerostering problem: An explorative study. *Proceedings of the International Conference on Industrial Engineering and Operations Management, Rabat, Morocco.*

National Heart Lung and Blood Institute (2011). What Is a Ventilator. `https://www.nhlbi.nih.gov/health/health-topics/topics/vent`.

Pato, M. V. and Moz, M. (2008). Solving a bi-objective nurse rerostering problem by using a utopic pareto genetic heuristic. *Journal of Heuristics*, 14(4):359–374.

Rauhala, A., Kivimäki, M., Fagerström, L., Elovainio, M., Virtanen, M., Vahtera, J., Rainio, A.-K., Ojaniemi, K., and Kinnunen, J. (2007). What degree of work overload is likely to cause increased sickness absenteeism among nurses? evidence from the rafaela patient classification system. *Journal of Advanced Nursing*, 57(3):286–295.

RSHU (2018). Medarbeidere. `https://stolav.no/fag-og-forskning/kompetansetjenester-og-sentre/rshu/medarbeidere`.

Rönnberg, E. and Larsson, T. (2010). Automating the self-scheduling process of nurses in Swedish health care: a pilot study. *Health Care Management Science*, 13(1):35–53.

Snedecor, G. W. C. and William, G. (1989). *Statistical Methods*. Iwoa State University Press.

Soyster, A. L. (1973). Technical Note—Convex Programming with Set-Inclusive Constraints and Applications to Inexact Linear Programming. *Operations Research*, 21(5):1154–1157.

SSB (2018). Om sesongjustering av sykefravær, egen- og legemeldt. `https://www.ssb.no/a/kortnavn/sykefratot/sesongjustering.html`.

St. Olavs Hospital (2018). Nøkkeltall for St. Olavs hospital. `https://stolav.no/om-oss/nokkeltall-for-st-olavs-hospital`.

Stein, H., Beck, J., and Dunn, M. (2016). Non-invasive ventilation with neurally adjusted ventilatory assist in newborns. *Seminars in Fetal and Neonatal Medicine*, 21:154–161.

Unio (2018). § 6 Overtid – forskjøvet arbeidstid. `https://www.unio.no/no/dokumenter/hovedtariffavtalen-ks-omraadet-2016-2018/6-overtid--forskjoevet-arbeidstid`.

Van den Bergh, J., Beliën, J., De Bruecker, P., and Demeulemeester, E. (2012). Personnel scheduling: A literature review. *European Journal of Operational Research*, 226(1):367–385.

Appendices

Appendix A

Compressed Models

This appendix contains compressed versions of the mathematical scheduling and rescheduling models and their corresponding model extensions, which were defined and explained in Chapters 5 and 7, respectively.

A.1 Scheduling Model

In this Section, we provide the entire scheduling model as well as the model extensions presented in Chapter 5, but without any of the explanations.

A.1.1 Definitions

Indices

n	employee
s	shift
t	day
k	week
c	skill
p	shift pattern

Sets

$\mathcal{N}$	set of employees
$\mathcal{C}$	set of skills, $\mathcal{C} = \{1, 2, 3, 4\}$, where 1 = Emergency skills, 2 = Intensive Care skills, 3 = Monitoring skills, 4 = Assistant nurse skills
$\mathcal{N}_c$	set of employees with skill c as their highest ranked skill, $\bigcup_{c\in\mathcal{C}} \mathcal{N}_c = \mathcal{N}$
$\mathcal{N}^{GEN}$	generic set of employees, explained further whenever used
$\mathcal{S}$	set of shifts
$\mathcal{S}^W$	set of work shifts, $\mathcal{S}^W = \{D, E, N\}$, $\mathcal{S}^W \subset \mathcal{S}$
$\mathcal{S}^O$	set of off shifts, $\mathcal{S}^O = \{F1, F\}$, $\mathcal{S}^O \subset \mathcal{S}$
$\mathcal{T}$	set of days in the current planning period
$\mathcal{T}^{SUN}$	set of Sundays
$\mathcal{T}^{\mathcal{P}}$	set of days in the current and previous planning period, where $t \leq 0$ indicates days in the previous period and $t = 1$ is the first day in the current period
$\mathcal{K}$	set of weeks in the planning period
$\mathcal{T}_k$	set of days in week k, $\bigcup_{k\in\mathcal{K}} \mathcal{T}_k = \mathcal{T}$
$\mathcal{P}^{DW}$	set of shift patterns occurring during a weekend which are desirable to assign
$\mathcal{P}^{GEN}$	set of shift patterns that are considered undesirable, and should only be assigned if the employee specifically requests it
$\mathcal{P}^{ILL}$	set of shift patterns illegal to assign

Parameters

Limit Parameters

$\underline{D}_{st}$	minimum number of employees required to cover total demand at shift s on day t
$\overline{D}_{st}$	maximum number of employees allowed to work at shift s on day t
$\underline{D}^C_{cst}$	minimum number of employees required to cover demand for skill c at shift s on day t
$\overline{M}^{CW}$	maximum number of consecutive work shifts for each employee

$\overline{M}^N$	maximum number of consecutive Night shifts for each employee
$\overline{H}_n^{7D}$	maximum number of hours employee n can work during a 7-day period

Weighing Parameters

W^R	reward for assigning a requested shift
W^P	reward for assigning a desirable shift pattern
W^S	penalty if demand for an employee of a particular skill is covered by an employee with a more advanced skill type as their main skill

Indicator Parameters

β_{nst}^{PA}	1 if employee n should have shift s preassigned on day t, 0 otherwise
β_{nst}^{NA}	1 if employee n should *never* have shift s assigned on day t, 0 otherwise
β_{nt}^{N}	1 if employee n can cover demand on day t, 0 otherwise
$\beta_{s_1 s_2}^{F1}$	1 if there is sufficient time between shifts s_1 and s_2 on days $t-2$ and t, respectively, for an employee to be assigned an 'F1'-day on day $t-1$, 0 otherwise
β_{nst}^{R}	1 if employee n has requested shift s on day t, 0 otherwise

General Parameters

H_s	duration of shift s in hours
H_n^{CW}	number of hours employee n should work during the planning period
H^{DEV}	allowed deviation in percent from H_n^{CW} for the number of hours assigned
M^{NW}	employees work every M^{NW} weekend
U	minimum amount in percent of shifts which, for each employee, must be an Evening or a Night shift
B	work load of a Night shift relative to an Evening shift
L_p	duration of shift pattern p in days
L_p^S	start day of shift pattern p relative to Sunday, where $L_p^S = 2$ is a Saturday, $L_p^S = 3$ is a Friday, and so on, where $p \in \mathcal{P}^{DW}$
S_{tp}	shift type on day t in shift pattern p

Variables

$$x_{nst} = \begin{cases} 1 & \text{if employee } n \text{ is assigned shift } s \text{ on day } t \\ 0 & \text{otherwise} \end{cases}$$

$y_{cst} =$ number of employees in a different skill group than group c who have to rank down to cover demand for that skill on shift s on day t

$$w_{ntp} = \begin{cases} 1 & \text{if employee } n \text{ works desirable shift pattern } p \text{ containing } t\text{, where } t \in \mathcal{T}^{SUN} \\ 0 & \text{otherwise} \end{cases}$$

A.1.2 Objective Function

$$\max Z = W^R \sum_{n\in\mathcal{N}} \sum_{s\in\mathcal{S}} \sum_{t\in\mathcal{T}} \beta^R_{nst} x_{nst} + W^P \sum_{n\in\mathcal{N}} \sum_{t\in\mathcal{T}^{SUN}} \sum_{p\in\mathcal{P}^{DW}} w_{ntp} - W^S \sum_{c\in\mathcal{C}} \sum_{s\in\mathcal{S}^W} \sum_{t\in\mathcal{T}} y_{cst} \tag{A.1}$$

A.1.3 Constraints

Covering Demand

$$\sum_{s\in\mathcal{S}} x_{nst} = 1 \qquad n \in \mathcal{N}, t \in \mathcal{T} \tag{A.2}$$

$$\underline{D}_{st} \leq \sum_{n\in\mathcal{N}} \beta^N_{nt} x_{nst} \leq \overline{D}_{st} \qquad s \in \mathcal{S}^W, t \in \mathcal{T} \tag{A.3}$$

$$\sum_{i=1}^{c} \sum_{n\in\mathcal{N}_i} \beta^N_{nt} x_{nst} \geq \sum_{i=1}^{c} \underline{D}^C_{ist} \qquad c \in \mathcal{C}, s \in \mathcal{S}^W, t \in \mathcal{T} \tag{A.4}$$

$$y_{cst} + \sum_{n\in\mathcal{N}_c} \beta^N_{nt} x_{nst} \geq \underline{D}^C_{cst} \qquad c \in \mathcal{C}, s \in \mathcal{S}^W, t \in \mathcal{T} \tag{A.5}$$

Weekends

$$\sum_{s \in \mathcal{S}^W} (x_{ns(t-1)} - x_{nst}) = 0 \qquad n \in \mathcal{N}, t \in \mathcal{T}^{SUN} \tag{A.6}$$

$$\sum_{s \in \mathcal{S}^W} \sum_{\tau=0}^{M^{NW}-1} x_{ns(t-7\tau)} \leq 1 \qquad n \in \mathcal{N}, t \in \mathcal{T}^{SUN} \tag{A.7}$$

$$2x_{nN(t-1)} - x_{nNt} - x_{nN(t+1)} = 0 \qquad n \in \mathcal{N}, t \in \mathcal{T}^{SUN} \tag{A.8}$$

Work Hours

$$H_n^{CW}(1 - H^{DEV}) \leq \sum_{s \in \mathcal{S}^W} \sum_{t \in \mathcal{T}} H_s x_{nst} \leq H_n^{CW}(1 + H^{DEV}) \qquad n \in \mathcal{N} \tag{A.9}$$

$$\sum_{s \in \mathcal{S}^W} \sum_{\tau=t-6}^{t} H_s x_{ns\tau} \leq \overline{H}_n^{7D} \qquad n \in \mathcal{N}, t \in \mathcal{T} \tag{A.10}$$

Required Rest

$$x_{ns_1(t-2)} + x_{n'F1'(t-1)} + \sum_{s_2 \in S | \beta^{F1}_{s_1 s_2} = 0} x_{ns_2 t} \leq 2 \qquad n \in \mathcal{N}, s_1 \in \mathcal{S}, t \in \mathcal{T} \tag{A.11}$$

$$\sum_{t \in \mathcal{T}_k} x_{n'F1't} = 1 \qquad n \in \mathcal{N}, k \in \mathcal{K} \tag{A.12}$$

Shift Patterns

$$\sum_{d=1}^{L_p} x_{nS_{dp}(t-L_p^S+d)} - L_p w_{ntp} \geq 0 \qquad n \in \mathcal{N}, t \in \mathcal{T}^{SUN}, p \in \mathcal{P}^{DW} \tag{A.13}$$

$$\sum_{d=1}^{L_p} x_{nS_{dp}(t-L_p+d)} - \prod_{d=1}^{L_p} \beta^R_{nS_{dp}(t-L_p+d)} \leq L_p - 1 \qquad n \in \mathcal{N}, t \in \mathcal{T}, p \in \mathcal{P}^{GEN} \tag{A.14}$$

$$\sum_{d=1}^{L_p} x_{nS_{dp}(t-L_p+d)} \leq L_p - 1 \qquad n \in \mathcal{N}, t \in \mathcal{T}, p \in \mathcal{P}^{ILL} \tag{A.15}$$

Other Scheduling Requirements

$$\sum_{s\in\mathcal{S}^W} \sum_{\tau=t-\overline{M}^{CW}}^{t} x_{ns\tau} \leq \overline{M}^{CW} \qquad n \in \mathcal{N}, t \in \mathcal{T} \tag{A.16}$$

$$\sum_{\tau=t-\overline{M}^{N}}^{t} x_{nN\tau} \leq \overline{M}^{N} \qquad n \in \mathcal{N}, t \in \mathcal{T} \tag{A.17}$$

$$\sum_{t\in\mathcal{T}} (x_{nEt} + Bx_{nNt}) \geq U \sum_{s\in\mathcal{S}^W} \sum_{t\in\mathcal{T}} x_{nst} \qquad n \in \mathcal{N}^{GEN} \tag{A.18}$$

Variable Declarations and Fixations

$$x_{nst} \in \{0,1\} \qquad n \in \mathcal{N}, s \in \mathcal{S}, t \in \mathcal{T} \tag{A.19}$$

$$y_{cst} \in \mathbb{N}_0 \qquad c \in \mathcal{C}, s \in \mathcal{S}^W, t \in \mathcal{T} \tag{A.20}$$

$$w_{ntp} \in \{0,1\} \qquad n \in \mathcal{N}, t \in \mathcal{T}^{SUN}, p \in \mathcal{P}^{DW} \tag{A.21}$$

$$x_{nst} = 1 \qquad n \in \mathcal{N}, s \in \mathcal{S}, t \in \mathcal{T}^{\mathcal{P}} \mid \beta_{nst}^{PA} = 1 \tag{A.22}$$

$$x_{nst} = 0 \qquad n \in \mathcal{N}, s \in \mathcal{S}, t \in \mathcal{T} \mid \beta_{nst}^{NA} = 1 \tag{A.23}$$

A.1.4 Proactive Model Extensions

Extension 1: Buffer

Definitions

r	index of variable in specially ordered set of type 1
F	maximum buffer size rewarded in objective function
$\mathcal{R}$	set of integers, $\mathcal{R} = \{1, \ldots, F\}$
W_r^B	reward of buffer of size r in excess of the minimum demand

W_r^{SB} reward of buffer of size r in excess of the minimum demand for a certain skill

$$\lambda_{rst}^{B} = \begin{cases} 1 & \text{if there are } r \text{ employees in excess of the minimum demand on shift } s \text{ on day } t \\ 0 & \text{otherwise} \end{cases}$$

$$\lambda_{rcst}^{SB} = \begin{cases} 1 & \text{if there are } r \text{ employees with skill } c \text{ in excess of the minimum demand for that skill on shift } s \text{ on day } t \\ 0 & \text{otherwise} \end{cases}$$

Objective and Constraints

$$z^D = \sum_{r\in\mathcal{R}} \sum_{s\in\mathcal{S}^W} \sum_{t\in\mathcal{T}} \left(W_r^B \lambda_{rst}^B + W_r^{SB} \sum_{c\in\mathcal{C}} \lambda_{rcst}^{SB} \right) \tag{A.24}$$

$$\sum_{n\in\mathcal{N}} \beta_{nt}^N x_{nst} - \sum_{r\in\mathcal{R}} r\lambda_{rst}^B \geq \underline{D}_{st} \qquad s\in\mathcal{S}^W, t\in\mathcal{T} \tag{A.25}$$

$$\sum_{n\in\mathcal{N}_c} \beta_{nt}^N x_{nst} + y_{cst} - \sum_{r\in\mathcal{R}} r\lambda_{rcst}^{SB} \geq \underline{D}_{cst}^C \qquad c\in\mathcal{C}, s\in\mathcal{S}^W, t\in\mathcal{T} \tag{A.26}$$

$$\sum_{r\in\mathcal{R}} \lambda_{rst}^B \leq 1 \qquad s\in\mathcal{S}^W, t\in\mathcal{T} \tag{A.27}$$

$$\sum_{r\in\mathcal{R}} \lambda_{rcst}^{SB} \leq 1 \qquad c\in\mathcal{C}, s\in\mathcal{S}^W, t\in\mathcal{T} \tag{A.28}$$

$$\lambda_{rst}^B \in \{0,1\}, SOS1 \qquad r\in\mathcal{R}, s\in\mathcal{S}^W, t\in\mathcal{T} \tag{A.29}$$

$$\lambda_{rcst}^{SB} \in \{0,1\}, SOS1 \qquad r\in\mathcal{R}, c\in\mathcal{C}, s\in\mathcal{S}^W, t\in\mathcal{T} \tag{A.30}$$

Extension 2: Ghost

Definitions

$\mathcal{S}^G$ set of ghost shifts, $\mathcal{S}^G = \{GN\}$, $\mathcal{S}^G \subset \mathcal{S}$

H Hours in full-time work week

W^G Reward for assigning a ghost shift

$$g_{nk}^H = \begin{cases} 1 & \text{if employee } n \text{ is assigned less than } H \text{ hours of work in week } k \\ 0 & \text{otherwise} \end{cases}$$

Objective and Constraints

$$z^G = W^G \sum_{n \in \mathcal{N}} \sum_{s \in \mathcal{S}^G} \sum_{t \in \mathcal{T}} x_{nst} \tag{A.31}$$

$$\sum_{s \in \mathcal{S}^W \cup \mathcal{S}^G} \sum_{\tau = t - \overline{M}^{CW}}^{t} x_{ns\tau} \leq \overline{M}^{CW} \qquad n \in \mathcal{N}, t \in \mathcal{T} \tag{A.32}$$

$$\sum_{\tau = t - \overline{M}^N}^{t} (x_{nN\tau} + x_{n'GN'\tau}) \leq \overline{M}^N \qquad n \in \mathcal{N}, t \in \mathcal{T} \tag{A.33}$$

$$\sum_{s \in \mathcal{S}^W \cup \mathcal{S}^G} \sum_{t \in \mathcal{T}_k} H_s x_{nst} \leq H + (\overline{H}_n^{7D} - H)(1 - g_{nk}^H) \qquad n \in \mathcal{N}, k \in \mathcal{K} \tag{A.34}$$

$$0 \leq \sum_{s \in \mathcal{S}^G} \sum_{t \in \mathcal{T}_k} H_s x_{nst} \leq H g_{nk}^H \qquad n \in \mathcal{N}, k \in \mathcal{K} \tag{A.35}$$

$$g_{nk}^H \in \{0, 1\} \qquad n \in \mathcal{N}, k \in \mathcal{K} \tag{A.36}$$

Extension 3: Absence

Definitions

$\mathcal{N}_t^{HR}$ set of nurses with high risk of being absent on day t

$\mathcal{T}^W$ set of Saturdays and Sundays

C_{nt} fraction of how much employee n contributes to covering demand on day t

W^{HR} penalty per additional employee with high risk of being absent that is assigned the same shift during weekends

S^{HR} number of employees in high-risk group who can be assigned the same shift without getting penalized in the objective function

s^{HR}_{st} = number of nurses with high risk of being absent exceeding S^{HR} that are assigned to shift s on day t

Objective and Constraints

$$z^{HR} = -W^{HR} \sum_{s \in \mathcal{S}^W} \sum_{t \in \mathcal{T}} s^{HR}_{st} \tag{A.37}$$

$$\underline{D}_{st} \leq \sum_{n \in \mathcal{N}} C_{nt} x_{nst} \leq \overline{D}_{st} \qquad s \in \mathcal{S}^W, t \in \mathcal{T} \tag{A.38}$$

$$\sum_{n \in \mathcal{N}^{HR}_t} x_{nst} - s^{HR}_{st} \leq S^{HR} \qquad s \in \mathcal{S}^W, t \in \mathcal{T}^W \tag{A.39}$$

$$s^{HR}_{st} \in \mathbb{N}_0 \qquad s \in \mathcal{S}^W, t \in \mathcal{T}^W \tag{A.40}$$

Extension 4: Extra Weekends

Definitions

$\overline{M}^{EW}$ maximum number of extra weekends an employee can work

H^{EO} number of extra off hours gained by working an extra weekend

β^{EW}_n 1 if employee n can work more weekends than normally contracted, 0 otherwise

$$e_{nt} = \begin{cases} 1 & \text{if employee } n \text{ is works an extra weekend containing Sunday } t \\ 0 & \text{otherwise} \end{cases}$$

Objective and Constraints

$$\sum_{t \in \mathcal{T}^{SUN}} e_{nt} \leq \overline{M}^{EW} \beta_n^{EW} \qquad n \in \mathcal{N} \tag{A.41}$$

$$\sum_{s \in \mathcal{S}^W} (x_{ns(t-1)} + x_{nst}) \geq 2\beta_n^{EW} e_{nt} \qquad n \in \mathcal{N}, t \in \mathcal{T}^{SUN} \tag{A.42}$$

$$\sum_{s \in \mathcal{S}^W} \sum_{\tau=0}^{M^{NW}-1} x_{ns(t-7\tau)} - \sum_{\tau=0}^{M^{NW}-1} e_{n(t-7\tau)} \leq 1 \qquad n \in \mathcal{N}, t \in \mathcal{T}^{SUN} \tag{A.43}$$

$$\sum_{\tau=0}^{M^{NW}-1} e_{n(t-7\tau)} \leq \beta_n^{EW} \qquad n \in \mathcal{N}, t \in \mathcal{T}^{SUN} \tag{A.44}$$

$$\sum_{s \in \mathcal{S}^W} \sum_{t \in \mathcal{T}} H_s x_{nst} + \sum_{t \in \mathcal{T}^{SUN}} H^{EO} e_{nt} \geq H_n^{CW}(1 - H^{DEV}) \qquad n \in \mathcal{N} \tag{A.45}$$

$$\sum_{s \in \mathcal{S}^W} \sum_{t \in \mathcal{T}} H_s x_{nst} + \sum_{t \in \mathcal{T}^{SUN}} H^{EO} e_{nt} \leq H_n^{CW}(1 + H^{DEV}) \qquad n \in \mathcal{N} \tag{A.46}$$

$$e_{nt} \in \{0, 1\} \qquad n \in \mathcal{N}, t \in \mathcal{T}^{SUN} \tag{A.47}$$

A.2 Rescheduling Model

A.2.1 Definitions

Indices

n	employee
c	skill
s	shift
t	day
k	week
q	double shift type

Sets

$\mathcal{T}^{PRE}$	set of days in pre-period

$\mathcal{T}^R$	set of days in the replanning period
$\mathcal{T}^{POST}$	set of days in post-period
$\mathcal{T}^{ALL}$	set of days, $\mathcal{T}^{ALL} = \{\mathcal{T}^{PRE} \cup \mathcal{T}^R \cup \mathcal{T}^{POST}\}$
$\mathcal{T}^{SN}$	set of days which trigger a short notice ahead of swapping a shift
$\mathcal{T}^{LN}$	set of days which trigger a long notice ahead of swapping a shift
$\mathcal{K}$	set of weeks containing the days $\mathcal{T}^R \cup \mathcal{T}^{POST}$
$\mathcal{T}_k$	set of days in week k
$\mathcal{S}$	set of shifts
$\mathcal{S}^W$	set of work shifts, $\mathcal{S}^W = \{D, E, N\}$, $\mathcal{S}^W \subset \mathcal{S}$
$\mathcal{S}^O$	set of off shifts, $\mathcal{S}^O = \{F\}$, $\mathcal{S}^O \subset \mathcal{S}$
$\mathcal{N}$	set of employees
$\mathcal{N}_t^W$	set of available employees who are assigned a work shift on day t, where $t \in \{\mathcal{T}^R \cup \mathcal{T}^{POST}\}$
$\mathcal{N}_t^O$	set of available employees who are assigned to an off shift on day t, where $t \in \{\mathcal{T}^R \cup \mathcal{T}^{POST}\}$
$\mathcal{N}_t^A$	set of available employees on day t, $\mathcal{N}_t^W \bigcup \mathcal{N}_t^O = \mathcal{N}_t^A$
$\mathcal{C}$	set of skills, $\mathcal{C} = \{1, 2, 3, 4\}$, where 1 = Emergency skills, 2 = Intensive Care skills, 3 = Monitoring skills, 4 = Assistant nurse skills
$\mathcal{N}_c$	set of employees with skill c as their highest ranked skill, $\bigcup_{c \in \mathcal{C}} \mathcal{N}_c = \mathcal{N}$
$\mathcal{Q}$	set of double shift types

Parameters

General Parameters

H	number of hours in a full-time work week
H_s	duration of shift s in hours
H_{nk}^{MAX}	number of work hours employee n can work in week k without incurring overtime pay
$\overline{M}^{CW}$	maximum number of consecutive work shifts without penalty per employee
$\overline{M}^{N}$	maximum number of consecutive Night shifts without penalty per employee

$\underline{M}^B$	minimum size of buffer assigned each shift in order for swaps and exchanges from the shift to be allowed
$\underline{D}_{st}^{RE}$	minimum online operational demand for employees for shift s on day t
$\underline{D}_{cst}^{RE}$	minimum online operational demand for employees with c as their highest skill for shift s on day t
D_{st}^{SIM}	real demand for employees for shift s on day t
D_{cst}^{SIM}	real demand for employees with c as their highest skill for shift s on day t

Indicator Parameters

X_{nst}	1 if employee n was initially scheduled to work shift s on day t, 0 otherwise
α_{nst}^{EX}	1 if employee n has requested to work an extra shift s on day t, 0 otherwise
α_{nt}^{OFF}	1 if employee n has requested to exchange a work day on day t, 0 otherwise
α_{nst}^{PRE}	1 if employee n should be preassigned shift s on day t, 0 otherwise
α_{nt}^{CW}	1 if employee n previously has been assigned a number of consecutive work shifts ending on day t that exceeds $\overline{M}^{CW}$, 0 otherwise
α_{nt}^{N}	1 if employee n previously has been assigned a number of consecutive Night shifts ending on day t that exceeds $\overline{M}^{N}$, 0 otherwise

Weighing Parameters

W^{REX}	penalty per assigned extra shift requested by an employee
W^{NEX}	penalty per assigned extra shift not requested by an employee
W^{EXC}	penalty per exchanged shift
W^{SN}	penalty per shift swapped on short notice
W^{LN}	penalty per shift swapped on long notice
$W^{DB,(q)}$	penalty per double shift of type q worked
W^{CW}	penalty per consecutive shift worked that exceeds the governmental maximum limit
W^{N}	penalty per consecutive Night shift worked that exceeds the governmental maximum limit
W_t^{D}	penalty per employee in shortage of covering overall demand on day t
W_t^{C}	penalty per employee in shortage of covering skill specific demand on day t

W^{CO} penalty per hour of overtime worked

Variables

Decision variables

$$x'_{nst} = \begin{cases} 1 & \text{if employee } n \text{ is assigned shift } s \text{ on day } t \\ 0 & \text{otherwise} \end{cases}$$

$$u_{nt_1t_2} = \begin{cases} 1 & \text{if employee } n \text{ was initially scheduled to a work shift on day } t_2\text{, but swapped or exchanged shifts to get a work shift on day } t_1 \text{ and an off day on day } t_2 \\ 0 & \text{otherwise} \end{cases}$$

$$v_{nst} = \begin{cases} 1 & \text{if employee } n \text{ takes on an extra shift of type } s \text{ on day } t \\ 0 & \text{otherwise} \end{cases}$$

$$z^{(1)}_{nt} = \begin{cases} 1 & \text{if employee } n \text{ works two consecutive shifts without a break in between, where the second shift occurs on day } t \\ 0 & \text{otherwise} \end{cases}$$

$$z^{(2)}_{nt} = \begin{cases} 1 & \text{if employee } n \text{ works two consecutive shifts with a break in between, but without getting sufficient rest according to rules and regulations, where the second shift occurs on day } t \\ 0 & \text{otherwise} \end{cases}$$

Slack Variables

s^{C}_{cst} = unsatisfied demand for skill c for shift s on day t

s^{D}_{st} = unsatisfied demand for employees for shift s on day t

$$s^{CW}_{nt} = \begin{cases} 1 & \text{if employee } n \text{ is assigned a pattern of consecutive work shifts exceeding } \overline{M}^{CW} \text{ that ends on day } t\text{, incurred by changes in the work schedule which were approved on the current day} \\ 0 & \text{otherwise} \end{cases}$$

$$s_{nt}^N = \begin{cases} 1 & \text{if employee } n \text{ is assigned a pattern of consecutive Night shifts exceeding } \overline{M}^N \text{ that ends on day } t\text{, incurred by changes in the work schedule which were approved on the current day} \\ 0 & \text{otherwise} \end{cases}$$

$s_{nk}^H =$ overtime worked by employee n in week k incurred by changes in the work schedule which were approved on the current day

Indicator Variables

$$d_{st} = \begin{cases} 1 & \text{if there is sufficient staff assigned shift } s \text{ on day } t \text{ to allow a swap or exchange from the shift} \\ 0 & \text{otherwise} \end{cases}$$

A.2.2 Objective Function

$$w_t^{DEM} = \sum_{s \in \mathcal{S}^W} \left(W^D s_{st}^D + \sum_{c \in \mathcal{C}} W^C s_{cst}^C \right) \qquad t \in \mathcal{T}^R \qquad \text{(A.48)}$$

$$w_t^{GOV} = \sum_{n \in \mathcal{N}} \left(W^{CW} s_{nt}^{CW} + W^N s_{nt}^N \right) \qquad t \in \mathcal{T}^R \qquad \text{(A.49)}$$

$$w_t^{EX} = \sum_{n \in \mathcal{N}} \sum_{s \in \mathcal{S}^W} \left(W^{REX} \alpha_{nst}^{EX} v_{nst} + W^{NEX} (1 - \alpha_{nst}^{EX}) v_{nst} \right) \qquad t \in \mathcal{T}^R \qquad \text{(A.50)}$$

$$w_t^{DB} = \sum_{q \in \mathcal{Q}} W^{DB,(q)} \sum_{n \in \mathcal{N}} z_{nt}^{(q)} \qquad t \in \mathcal{T}^R \qquad \text{(A.51)}$$

$$w_t^{EXC} = W^{EXC} \sum_{n \in \mathcal{N}} \sum_{t_2 \in \mathcal{T}^R \setminus \{t\}} \alpha_{nt_2}^{OFF} u_{ntt_2} \qquad t \in \mathcal{T}^R \qquad \text{(A.52)}$$

$$\begin{aligned} w^{SWAP} &= W^{SN} \sum_{n \in \mathcal{N}} \sum_{t_1 \in \mathcal{T}^{SN}} \sum_{t_2 \in \mathcal{T}^R \setminus \{t_1\}} (u_{nt_1t_2} + u_{nt_2t_1})(1 - \alpha_{t_2}^{OFF}) \\ &+ W^{LN} \sum_{n \in \mathcal{N}} \sum_{t_1 \in \mathcal{T}^{LN}} \sum_{t_2 \in \mathcal{T}^{LN} \setminus \{t_1\}} u_{nt_1t_2}(1 - \alpha_{t_2}^{OFF}) \\ &+ W^{SN} \sum_{n \in \mathcal{N}} \sum_{t_1 \in \mathcal{T}^{SN}} u_{nt_1t_1} + W^{LN} \sum_{n \in \mathcal{N}} \sum_{t_1 \in \mathcal{T}^{LN}} u_{nt_1t_1} \end{aligned} \qquad \text{(A.53)}$$

$$w_k^{OVER} = W^{CO} \sum_{n \in \mathcal{N}} s_{nk}^H \qquad k \in \mathcal{K} \qquad \text{(A.54)}$$

$$\min Z = \sum_{t \in \mathcal{T}^R} \left(w_t^{DEM} + w_t^{GOV} + w_t^{EX} + w_t^{DB} + w_t^{EXC} \right) + w^{SWAP} + \sum_{k \in \mathcal{K}} w_k^{OVER} \quad \text{(A.55)}$$

A.2.3 Constraints

Covering Real Demand

$$\sum_{n \in \mathcal{N}_t^A} x'_{nst} + s_{st}^D - \underline{M}^B d_{st} \geq \max\{\underline{D}_{st}^{RE}, D_{st}^{SIM}\} \qquad s \in \mathcal{S}^W, t \in \mathcal{T}^R \quad \text{(A.56)}$$

$$\sum_{i=1}^{c} \sum_{n \in \mathcal{N}_i \cap \mathcal{N}_t^A} x'_{nst} + s_{cst}^C \geq \max\{\underline{D}_{cst}^{RE}, D_{cst}^{SIM}\} \qquad c \in \mathcal{C}, s \in \mathcal{S}^W, t \in \mathcal{T}^R \quad \text{(A.57)}$$

Technical Constraints for Actions

$$x'_{nst_1} - v_{nst_1} - \sum_{t_2 \in \mathcal{T}^R} u_{nt_1t_2} - \sum_{q \in \mathcal{Q}} z_{nt_1}^{(q)} \leq X_{nst_1} + \alpha_{nst_1}^{PRE} \qquad t_1 \in \mathcal{T}^R, n \in \mathcal{N}_{t_1}^A, s \in \mathcal{S}^W \quad \text{(A.58)}$$

$$\sum_{s \in \mathcal{S}} x'_{nst} - \sum_{q \in \mathcal{Q}} z_{nt}^{(q)} \leq 1 \qquad t \in \mathcal{T}^R, n \in \mathcal{N}_t^A \quad \text{(A.59)}$$

$$\sum_{t_1 \in \mathcal{T}^R} u_{nt_10} + \sum_{q \in \mathcal{Q}} z_{n0}^{(q)} \leq 1 \qquad n \in \mathcal{N}_0^A \quad \text{(A.60)}$$

Double Shift

$$x'_{nN(t-1)} + x'_{nDt} - z_{nt}^{(1)} \leq 1 \qquad t \in \{0 \ldots |\mathcal{T}^R|\}, n \in \mathcal{N}_t^A \quad \text{(A.61)}$$

$$x'_{nDt} + x'_{nEt} - z_{nt}^{(1)} \leq 1 \qquad t \in \mathcal{T}^R, n \in \mathcal{N}_t^A \quad \text{(A.62)}$$

$$x'_{nEt} + x'_{nNt} - z_{nt}^{(1)} \leq 1 \qquad t \in \mathcal{T}^R, n \in \mathcal{N}_t^A \quad \text{(A.63)}$$

$$x'_{nDt} + x'_{nNt} - z_{nt}^{(2)} \leq 1 \qquad t \in \mathcal{T}^R, n \in \mathcal{N}_t^A \quad \text{(A.64)}$$

$$x'_{nN(t-1)} + x'_{nEt} - z_{nt}^{(2)} \leq 1 \qquad t \in \{0 \ldots |\mathcal{T}^R|\}, n \in \mathcal{N}_t^A \quad \text{(A.65)}$$

$$\sum_{q \in \mathcal{Q}} z_{nt}^{(q)} \leq 1 \qquad t \in \mathcal{T}^R, n \in \mathcal{N}_t^A \quad \text{(A.66)}$$

Swap and Exchange

$$\sum_{t_1 \in \mathcal{T}^R} u_{nt_1t_2} + \sum_{s \in \mathcal{S}^W} \left(X_{nst_2} x'_{nst_2} + \alpha^{PRE}_{nst_2} \right) = 1 \qquad t_2 \in \mathcal{T}^R, n \in \mathcal{N}^W_{t_2} \tag{A.67}$$

$$\sum_{t_2 \in \mathcal{T}^R} u_{nt_1t_2} \leq 1 \qquad t_1 \in \mathcal{T}^R, n \in \mathcal{N}^O_{t_1} \tag{A.68}$$

$$\sum_{t_1 \in \mathcal{T}^R} u_{nt_1t_2} \leq 1 \qquad t_2 \in \mathcal{T}^R, n \in \mathcal{N}^W_{t_2} \tag{A.69}$$

$$\sum_{t_1 \in \mathcal{T}^R} u_{nt_1t_2} - \sum_{s \in \mathcal{S}^W} d_{st_2} \left(X_{nst_2} + \alpha^{PRE}_{nst_2} \right) \leq 0 \qquad t_2 \in \mathcal{T}^{LN}, n \in \mathcal{N}^W_{t_2} \tag{A.70}$$

Consecutive Work

$$\sum_{\tau = t - \overline{M}^{CW}}^{t} \left(\sum_{s \in \mathcal{S}^W \setminus \{N\}} x'_{ns\tau} + x'_{nN(\tau-1)} \right) - s^{CW}_{nt} \leq \overline{M}^{CW} + \alpha^{CW}_{nt} \quad n \in \mathcal{N}, t \in \mathcal{T}^R \cup \mathcal{T}^{POST} \tag{A.71}$$

$$\sum_{\tau = t - \overline{M}^{N}}^{t} x'_{nN\tau} - s^{N}_{nt} \leq \overline{M}^{N} + \alpha^{N}_{nt} \qquad n \in \mathcal{N}, t \in \mathcal{T}^R \cup \mathcal{T}^{POST} \tag{A.72}$$

$$\sum_{s \in \mathcal{S}^W} \sum_{t \in \mathcal{T}_k} H_s x'_{nst} - s^{H}_{nk} \leq H^{MAX}_{nk} \qquad n \in \mathcal{N}, k \in \mathcal{K} \tag{A.73}$$

Variable Declarations and Fixations

$$x'_{nst} \in \{0,1\} \qquad t \in \mathcal{T}^{ALL}, n \in \mathcal{N}^A_t, s \in \mathcal{S} \tag{A.74}$$

$$u_{nt_1t_2} \in \{0,1\} \qquad t_1, t_2 \in \mathcal{T}^R, n \in \mathcal{N}^A_{t_1} \cap \mathcal{N}^W_{t_2} \tag{A.75}$$

$$v_{nst} \in \{0,1\} \qquad t \in \mathcal{T}^R, n \in \mathcal{N}^O_t, s \in \mathcal{S}^W \tag{A.76}$$

$$z^{(q)}_{n0} \in \{0,1\} \qquad n \in \mathcal{N}^W_0, q \in \mathcal{Q} \tag{A.77}$$

$$z^{(q)}_{nt} = 0 \qquad t \in \{1 \dots |\mathcal{T}^R|\}, n \in \mathcal{N}^A_t, q \in \mathcal{Q} \tag{A.78}$$

$$x'_{nst} = 1 \qquad n \in \mathcal{N}, s \in \mathcal{S}, t \in \mathcal{T}^{ALL} \mid \alpha^{PRE}_{nst} = 1 \tag{A.79}$$

$$x'_{nst} = X_{nst} \qquad t \in \mathcal{T}^R, n \in \mathcal{N} \backslash \mathcal{N}_t^A, s \in \mathcal{S} \mid \alpha_{nst}^{PRE} = 0 \tag{A.80}$$

$$s_{cst}^C \in \mathbb{N}_0 \qquad c \in \mathcal{C}, s \in \mathcal{S}^W, t \in \mathcal{T}^R \tag{A.81}$$

$$s_{st}^D \in \mathbb{N}_0 \qquad s \in \mathcal{S}^W, t \in \mathcal{T}^R \tag{A.82}$$

$$s_{nt}^{CW} \in \{0,1\} \qquad n \in \mathcal{N}, t \in \mathcal{T}^R \cup \mathcal{T}^{POST} \tag{A.83}$$

$$s_{nt}^N \in \{0,1\} \qquad n \in \mathcal{N}, t \in \mathcal{T}^R \cup \mathcal{T}^{POST} \tag{A.84}$$

$$s_{nk}^H \geq 0 \qquad n \in \mathcal{N}, k \in \mathcal{K} \tag{A.85}$$

$$d_{st} \in \{0,1\} \qquad s \in \mathcal{S}^W, t \in \mathcal{T}^{LN} \tag{A.86}$$

A.2.4 Extension 1: Stricter Rescheduling

Objective and Constraints

$$x'_{nst_1} - v_{nst_1} - \sum_{t_2 \in \mathcal{T}^R} u_{nt_1t_2} \leq X_{nst_1} + \alpha_{nst_1}^{PRE} \quad t_1 \in \mathcal{T}^R, n \in \mathcal{N}_{t_1}^A, s \in \mathcal{S}^W \tag{A.87}$$

$$\sum_{s \in \mathcal{S}} x'_{nst} = 1 \qquad t \in \mathcal{T}^{\mathcal{R}}, n \in \mathcal{N}_t^A \tag{A.88}$$

$$\sum_{t_1 \in \mathcal{T}^R} u_{nt_10} \leq 1 \qquad n \in \mathcal{N}_0^A \tag{A.89}$$

$$x'_{nN(t-1)} + x'_{nDt} \leq 1 \qquad n \in \mathcal{N}_t^A, t \in \{0 \dots |\mathcal{T}^R|\} \tag{A.90}$$

$$x'_{nN(t-1)} + x'_{nEt} \leq 1 \qquad n \in \mathcal{N}_t^A, t \in \{0 \dots |\mathcal{T}^R|\} \tag{A.91}$$

$$\sum_{\tau = t - \overline{M}^{CW}}^{t} \left(\sum_{s \in \mathcal{S}^W \backslash \{N\}} x'_{ns\tau} + x'_{nN(\tau-1)} \right) \leq \overline{M}^{CW} \qquad n \in \mathcal{N}, t \in \mathcal{T}^R \cup \mathcal{T}^{POST} \tag{A.92}$$

$$\sum_{\tau = t - \overline{M}^N}^{t} x'_{nN\tau} \leq \overline{M}^N \qquad n \in \mathcal{N}, t \in \mathcal{T}^R \cup \mathcal{T}^{POST} \tag{A.93}$$

Appendix B

Examples

B.1 Scheduling Example

Table B.1: An example schedule for 9 nurses over 9 days, where D, E, N, F and F1 denote Day, Evening, Night, Off and the mandatory weekly off day, respectively

Day Index	Sun 0	Mon 1	Tue 2	Wed 3	Thu 4	Fri 5	Sat 6	Sun 7	Mon 8
Nurse 1	F1	D	D	E	F	F1	E	D	F
Nurse 2	F1	E	E	F	F1	E	D	E	D
Nurse 3	F1	E	D	E	D	F	F	F1	D
Nurse 4	F1	F	E	D	E	E	F	F1	F
Nurse 5	F1	F	F	F	F1	F	N	N	N
Nurse 6	F1	F	N	N	F	D	F	F1	E
Nurse 7	E	D	F	F	N	N	F	F1	F
Nurse 8	D	F	F	D	E	F	F	F1	F
Nurse 9	N	N	F	F	D	D	F	F1	E
#D	1	2	2	2	2	2	1	1	2
#E	1	2	2	2	2	2	1	1	2
#N	1	1	1	1	1	1	1	1	1
#F	6	4	4	4	4	4	6	6	4
$\underline{D}_{st}$	1	1	1	1	1	1	1	1	1

B.2 Uncertainty Example

Table B.2: An example schedule, where the demand on Wednesday and Thursday is higher than expected, Nurse 4 is absent on Wednesday, and Nurse 8 hands in a sickness note of 7 days

Day	Sun	Mon	Tue	Wed	Thu	Fri	Sat	Sun	Mon
Index	-3	-2	-1	0	1	2	3	4	5
Nurse 1	F1	D	D	E	F	F1	E	D	F
Nurse 2	F1	E	E	F	F1	E	D	E	D
Nurse 3	F1	E	D	E	D	F	F	F1	D
Nurse 4	F1	F	E	~~D~~	E	E	F	F1	F
Nurse 5	F1	F	F	F	F1	F	N	N	N
Nurse 6	F1	F	N	N	F	D	F	F1	E
Nurse 7	E	D	F	F	N	N	F	F1	F
Nurse 8	D	F	F	~~D~~	~~E~~	~~F~~	~~F~~	~~F1~~	~~F~~
Nurse 9	N	N	F	F	D	D	F	F1	E
#D	1	1	2	1	2	2	1	1	2
#E	1	2	2	2	1	2	1	1	2
#N	1	1	1	1	1	1	1	1	1
Real demand	1	1	1	2	2	1	1	1	1

B.3 Rescheduling Example

Table B.3: An example schedule, where the current day is denoted by index 0 and the replanning period consists of three days. Nurse 4 is absent today, Nurse 8 is long-term absent, and the real demand today and tomorrow is higher than expected

	Pre-period			Replanning period			Post-Period		
Day	Sun	Mon	Tue	Wed	Thu	Fri	Sat	Sun	Mon
Index	-3	-2	-1	0	1	2	3	4	5
Nurse 1	F	D	D	E	F	F	E	D	F
Nurse 2	F	E	E	F	F	E	D	E	D
Nurse 3	F	E	D	E	D	F	F	F	D
Nurse 4	F	F	E	~~D~~	E	E	F	F	N
Nurse 5	F	F	F	F	F	N	N	N	F
Nurse 6	F	N	N	F	F	D	F	F	E
Nurse 7	E	D	F	N	N	F	F	F	F
Nurse 8	D	F	F	~~D~~	~~E~~	~~F~~	~~F~~	~~F~~	~~F~~
Nurse 9	N	F	F	F	D	D	F	F	E
#D	1	2	2	0	2	2	1	1	2
#E	1	2	2	2	1	2	1	1	2
#N	1	1	1	1	1	1	1	1	1
#F	6	4	4	6	4	4	6	6	4
Real demand	1	1	1	2	2	1	1	1	1

Table B.4: Example of one possible assignment of actions when the real demand on day 0 and 1 is two, Nurse 4 is absent today and Nurse 8 has a long-term absence. A change to an employee's initial shift is illustrated by a slash cancellation and highlighted by blue

Day Index	Tue -1	Wed 0	Thu 1	Fri 2	Action	Variable
Nurse 1	D	~~E~~ E + N	F	F	Double shift of type 1	$z_{10}^{(1)} = 1$
Nurse 2	E	F	F	E	No actions are made	
Nurse 3	D	E	D	F	No actions are made	
Nurse 4	E	~~D~~	E	E	No actions are possible	
Nurse 5	F	F	~~F~~ N	N	Extra shift Thursday	$v_{5N1} = 1$
Nurse 6	N	F	~~F~~ E	~~D~~ F	Day shift Friday swapped	$u_{612} = 1$
Nurse 7	F	~~N~~ D + N	N	F	Double shift of type 2	$z_{70}^{(2)} = 1$
Nurse 8	F	~~D~~	~~E~~	~~F~~	No actions are possible	
Nurse 9	F	~~F~~ D	D	D	Extra shift today	$v_{9D0} = 1$
#D #E #N #F	2 2 1 4	2 2 2 3	2 2 2 3	1 2 1 5		
D_{st}^{SIM}	1	2	2	1		

B.4 Calculating Real Demand

The real demand is calculated using the values for the estimated number of nurses required to treat one patient at each level displayed in Table 8.18. Additionally, there is always demand for one coordinator and one employee with Emergency skills on all shifts.

Table B.5: Estimated need for nurses per patient per level (Halsteinli, 2017), where skill 2 is Intensive Care and 3 is Monitoring

Level	1	2	3	4	5
Nurses per patient	0.3	0.4	0.7	1.0	1.5
Minimum skill	3	3	3	2	2

To illustrate how the number of patients is related to the demand of nurses per skill and the overall demand, assume the scenario in Table B.6, where 20 patients are admitted to DNIC. Based on this patient mix, there is a demand for 12.8 nurses, where 6.3 should have skills minimally corresponding to the Monitoring Unit and 6.5 should have Intensive Care skills or higher. Additionally, there is demand for one coordinator and one employee with Emergency skills on all shifts, giving an overall demand for 14.8 employees, with an average need for care per patient of 0.74.

Table B.6: Historical patient data from February 1st 2016

Level	1	2	3	4	5	Total
Patients per level	2	9	3	5	1	20
Nurses with Monitoring skills required	0.6	3.6	2.1			6.3
Nurses with Intensive Care skills required				5.0	1.5	6.5
Nurses with pre-defined tasks						2.0
Total demand for nurses						14.8

Appendix C

Data Analysis and Simulation Results

C.1 Absence Probability Distributions

Table C.1: Probabilities for transitioning between the states non-absent (a_N), short-term absent (a_S), and long-term absent (a_L)

(a) Low-risk group on weekdays

		To		
		a_N	a_S	a_L
From	a_N	0.97	0.03	0.00
	a_S	0.43	0.56	0.01
	a_L	0.05	0.00	0.95

(b) High-risk group on weekdays

		To		
		a_N	a_S	a_L
From	a_N	0.89	0.10	0.01
	a_S	0.35	0.64	0.01
	a_L	0.03	0.01	0.96

(c) Low-risk group on weekends

		To		
		a_N	a_S	a_L
From	a_N	0.99	0.01	0.00
	a_S	0.40	0.60	0.00
	a_L	0.02	0.02	0.96

(d) High-risk group on weekends

		To		
		a_N	a_S	a_L
From	a_N	0.94	0.06	0.01
	a_S	0.18	0.80	0.01
	a_L	0.01	0.01	0.98

C.2 Results from Hypothesis Testing

Table C.2: Results from two-tailed, two-samples t-tests with a significance level equal to 5 % for short-term absences on various days of the week. H_0 is that the means of variable 1 and 2 are equal.

Variable 1	**Variable 2**	p**-value**	**Reject** H_0**?**
Monday	Sunday	0.046	Yes
Monday	Saturday	0.020	Yes
Monday	Friday	0.922	No
Monday	Thursday	0.509	No
Monday	Wednesday	0.711	No
Monday	Tuesday	0.274	No
Tuesday	Sunday	0.002	Yes
Tuesday	Saturday	0.001	Yes
Tuesday	Friday	0.234	No
Tuesday	Thursday	0.721	No
Tuesday	Wednesday	0.475	No
Wednesday	Sunday	0.019	Yes
Wednesday	Saturday	0.008	Yes
Wednesday	Friday	0.640	No
Wednesday	Thursday	0.755	No
Thursday	Sunday	0.013	Yes
Thursday	Saturday	0.005	Yes
Thursday	Friday	0.452	No
Friday	Sunday	0.057	No
Friday	Saturday	0.026	Yes
Saturday	Sunday	0.755	No

Table C.3: Results from two-tailed, two-samples t-tests with a significance level equal to 5 % for short-term absences on various shift types. H_0 is that the means of variable 1 and 2 are equal.

Variable 1	**Variable 2**	p**-value**	**Reject** H_0**?**
Weekdays_D	Weekdays_E	0.069	No
Weekdays_D	Weekdays_N	0.124	No
Weekdays_E	Weekdays_N	0.827	No
Weekends_D	Weekends_E	0.067	No
Weekends_D	Weekends_N	0.124	No
Weekends_E	Weekends_N	0.827	No

Table C.4: Results from two-tailed, two-samples t-tests with a significance level equal to 5 % for long-term absences and various days of the week. H_0 is that the means of variable 1 and 2 are equal.

Variable 1	**Variable 2**	p**-value**	**Reject** H_0**?**
Monday	Sunday	0.343	No
Monday	Saturday	0.357	No
Monday	Friday	0.704	No
Monday	Thursday	0.649	No
Monday	Wednesday	0.950	No
Monday	Tuesday	0.699	No
Tuesday	Sunday	0.579	No
Tuesday	Saturday	0.599	No
Tuesday	Friday	1.000	No
Tuesday	Thursday	0.940	No
Tuesday	Wednesday	0.739	No
Wednesday	Sunday	0.365	No
Wednesday	Saturday	0.380	No
Wednesday	Friday	0.743	No
Wednesday	Thursday	0.687	No
Thursday	Sunday	0.640	No
Thursday	Saturday	0.661	No
Thursday	Friday	0.941	No
Friday	Sunday	0.585	No
Friday	Saturday	0.605	No
Saturday	Sunday	0.975	No

Table C.5: Results from two-tailed, two-samples t-tests with a significance level equal to 5 % for long-term absences on various shift types. H_0 is that the means of variable 1 and 2 are equal.

Variable 1	**Variable 2**	p**-value**	**Reject** H_0**?**
Weekdays_E	Weekdays_N	0.897	No
Weekdays_D	Weekdays_E	0.835	No
Weekdays_D	Weekdays_N	0.936	No
Weekends_E	Weekends_N	0.897	No
Weekends_D	Weekends_E	0.834	No
Weekends_D	Weekends_N	0.936	No

Appendix D

Technical Analysis of Instances `s5a-c`

Table D.1: Values of the weighing parameters in instances `s5a-c`

Parameter	`s5a`	`s5b`	`s5c`
W^R	1	1	1
W^P	0.8	0.8	0.8
W^S	2	0.3	2
W_1^B	1.5	–	1.5
W_1^{SB}	1.5	–	1.5
W^G	–	0.5	–
W^{HR}	–	1.2	1.2

The upper limits for overall maximum demand during weekdays for instances `s5a-c` were the same as in the base case instance, namely $\overline{D}_{nDt} = 30$, $\overline{D}_{nEt} = 20$ and $\overline{D}_{nNt} = 20$. The minimum demand for weekend shifts for instances `s5a` and `s5c` was $\underline{D}_{nst} = 14$.

Table D.2: Key variable values, run time and optimality gap for instances `s5a-c`

Instance	`s5a`	`s5b`	`s5c`
Requests granted [%]	84.76	87.79	85.24
Work shifts	5545	5530	5549
Desirable patterns [%]	58.67	63.74	60.46
Rank-downs in skills	29	29	10
First sol under 10 % gap [s]	1145	10680	3370
Run time [s]	86400	86400	88000
Optimality gap [%]	2.39	1.13	1.95