Strategic Airline Crew Sizing

A Two-Stage Stochastic Optimisation Approach

Sam Hofman

TU Delft

Strategic Airline Crew Sizing

A Two-Stage Stochastic Optimisation Approach

by

Sam Hofman

to obtain the degree of Master of Science
at the Delft University of Technology,
to be defended publicly on December 7th, 2021 at 09:00.

Student number:	4353862	
Project duration:	March 2020 – December 2021	
Thesis committee:	Prof.dr.ir. Mirjam Snellen	TU Delft, Chair
	Dr.ir. Bruno F. Santos,	TU Delft, Supervisor
	Ir. Alessandro Bombelli	TU Delft, Examiner
	M. de Haas,	TUI fly, External Supervisor

Cover image by Thijs van Zwol, 2019.

An electronic version of this thesis is available at `http://repository.tudelft.nl/`.

TU Delft   TUI

Acknowledgements

Dear reader,

This report marks the end of my studies at the Delft University of Technology and is the result of more than one and a half year of work. Delivering this final piece of work during a global pandemic has been the most challenging task during my studies, and therefore I am very proud I can present it to you.

I would like to take this chance and thank some people who have motivated and supported me during this time. First of all, I would like to thank my thesis supervisor Dr.ir. Bruno Santos, who has always given me very helpful feedback and gently steered me in the right direction, while still giving me all the freedom to make decisions. Your critical but constructive guidance has been a real help and I am very happy you agreed to being my supervisor. I would also like to thank Michael de Haas at TUI fly. Our bi-weekly meetings made sure I stayed motivated and connected to the company amidst all work-at-home advices. Michael also provided me with all the inside information I needed and was enthusiastic about the research and its opportunities from the beginning.

In addition, I would like to thank all my friends for pushing me to keep going and for giving me the necessary distraction from time to time, both in Belgium and The Netherlands. Wouter, thank you for this amazing time in Delft. You have been with me from the start and were always there to hang out or grab a beer. Ruth and Kato, thank you for making me laugh at home and understanding that I could not be there as much as I would have wanted. Also thank you for not graduating before your older brother. Guul, thank you for always being there for me, for motivating me and for listening to my complaints. You are my rock.

Finally, I cannot say thank you enough to my parents, who gave me the opportunity to study abroad when I was only 17 and have always supported me, even though they do not always understand what I am doing. Without your support, trust and encouragement, this would have been impossible.

Thank you!

Sam Hofman
Delft, December 2021

Contents

List of Figures

List of Tables

Nomenclature

AHP	Analytic hierarchy process
ARIMA	Autoregressive integrated moving average
CP	Captain
DMP	Deterministic manpower problem
DSS	Decision support system
DV	Decision variable
FO	First officer
GA	Genetic algorithm
GM	Grey model
IATA	International Air Transport Association
LA	Learning automata
LH	Long-haul
LHS	Latin hypercube sampling
LP	Linear programming
MAE	Mean absolute error
MAPE	Mean absolute percentage error
MILP	Mixed-integer linear programming
MOLP	Multi-objective linear programming
MSE	Mean squared error
MU	Monetary units
NN	Neural network
OF	Objective function
PDF	Probability density function
PMF	Probability mass function
RSS	Residual sum of squares
SARIMA	Seasonal autoregressive integrated moving average
SH	Short-haul
SMP	Stochastic manpower problem

Introduction

The research presented in this report is situated in the strategic crew planning field. Crew costs are the second biggest expense for airlines, hence a small improvement in workforce planning can lead to significant savings. Within the long and complex process of crew planning, the strategic phase focuses on forecasting future crew demand and supply, and on defining strategies to close the gap between them, approximately one and a half year in advance.

Crew demand is mainly determined by the flight schedule, but is naturally increased to account for holidays, compulsory rest time, training days, crew absence, etc. Crew supply on the other hand is dictated by the current workforce, but also fluctuates due to e.g. retirements, illness and the economic situation. Once crew demand and supply are determined, the aim is to match them by means of transitions, hiring new crew, or in the worst case dismissing employees. In the strategic planning phase, the flight schedule is often still unknown and thus in practice airlines tend to not look more than one year ahead when it comes to crew planning.

Because of this uncertainty, most research seems to take crew demand as a given, *deterministic*, input and then investigates the best strategy of closing the gap between supply and demand, or focuses on the more urgent short-term operational planning phase. The research presented in this report focuses on developing a flexible, cost-optimal strategic crew plan for airlines, and assumes the crew demand is of a stochastic nature.

The research has been carried out at the Air Transport & Operations department at the Delft University of Technology, and in cooperation with TUI fly Western Region, under the supervision of Dr.ir. Bruno F. Santos (TU Delft) and Michael de Haas (TUI fly).

This thesis report is split into three parts. Part I contains the central part of the thesis project, the research article. Part II contains the literature study written before the start of the research. It includes relevant literature, identifies a research gap and introduces the the research question. Finally, Part III contains all the appendices. Here, supporting work related to the research article of Part I can be found.

I

Scientific Paper

A Two-Stage Stochastic Optimisation Approach for Strategic Airline Cockpit Crew Sizing

S. Hofman*

Delft University of Technology, Delft, The Netherlands

Abstract

This paper introduces a new two-stage stochastic optimisation approach for the strategic airline crew planning problem with uncertain demand. The developed model provides the airline with a cockpit crew composition plan per crew position before the flight schedule and crew demand are known. This is done by estimating future crew demand and by treating it as a stochastic variable. Historical crew demand is analysed and is assumed to follow a Beta probability distribution. Since demand at different crew positions is correlated, demand scenarios are generated from these distributions by using Latin hypercube sampling (LHS) for correlated variables. The generated scenarios are multiplied with a trend value to account for a predicted demand increase or decrease. In the first stage, the model determines the number of permanent employees, while in the second stage we consider hiring temporary crew members, the transition of crew members between compatible positions and the option to fire permanent crew members. This second stage comes with a recourse cost. Three case studies are performed for a holiday airline with both scheduled and charter flights to validate the model and test its possibilities. Results show that the model provides a cost reduction of 2.1% with respect to the airline's current practice and a further 4.5% when affluent crew members can be fired. Results also show that the model is adjustable to assist the airline with its strategic crew plan in the post-Covid-19 recovery phase. Furthermore, it is concluded that the model is flexible and that it can handle demand fluctuations.

Keywords: *Two-stage Stochastic Optimisation · Strategic Airline Planning · Scenario Generation · Latin Hypercube Sampling · Uncertain Crew Demand*

1 Introduction

Crew costs are the second biggest expense for airlines, hence a small improvement in workforce planning can lead to significant savings (Belobaba et al., 2009). Airline workforce planning can start as soon as five years before the day of operation and lasts until the very day itself. Within this long and complex process, strategic workforce planning focuses on forecasting future crew demand and supply, and on defining strategies to close the gap between them, approximately one and a half year in advance.

Crew demand is mainly determined by the flight schedule, but is naturally increased to account for holidays, compulsory rest time, training days, crew absence, etc. Crew supply on the other hand is dictated by the current workforce, but also fluctuates due to e.g. retirements, illness and the economic situation. Once crew demand and supply are determined, the aim is to match them by means of transitions, hiring new crew, or in the worst case dismissing employees. However, forecasting crew demand already in the strategic planning phase has been called the hardest part of the crew planning problem because of the high uncertainty (Holm, 2008). In this phase, the flight schedule is often still unknown and thus in practice airlines tend to not look more than one year ahead when it comes to crew planning (Hofman, 2020).

Not only the lack of a flight schedule makes forecasting crew demand so difficult, also external factors can profoundly affect air travel and thus crew demand. In recent years, the airline industry has seen among others economic crises, the grounding of the Boeing 737 MAX and a global pandemic. All these events heavily weigh on airlines and are often impossible to foresee.

*MSc Student, Air Transport and Operations, Faculty of Aerospace Engineering, Delft University of Technology

Because of this high uncertainty, most research seems to take crew demand as a given, *deterministic*, input and then investigates the best strategy of closing the gap between supply and demand, or focuses on the more urgent short-term operational planning phase. Overall, little research has been performed that investigates the optimisation of long-term, strategic crew planning for airlines.

Next to this, there has been a growing interest in planning robustness from the airlines over the years. A robust plan can deal with the realisation of uncertainties without having a significant negative impact on the solution (Laumanns, 2011). In other words, taking into account possible disruptions already in the (strategic) planning phase can result in higher planned costs, but lower operational costs (Ehrgott and Ryan, 2002). In spite of that, research focuses on improving robustness in the operational rather than in the strategic planning phase.

Following from these identified gaps and needs, the main research objective of this paper is *"to contribute to solving the long-term, strategic workforce planning problem for the cockpit crew of an airline with uncertain flight schedules while considering planning robustness and flexibility by creating an optimisation model to analyse different demand realisations."* The model should provide the airline with a cost-optimal crew plan already in the strategic planning phase, so before the flight schedule is known. Furthermore, the tool should give the airline an overview of possible future scenarios in terms of crew supply and demand and guide the planning department in anticipating these possible outcomes.

In this paper, we propose such a strategic crew planning tool by creating a two-stage stochastic linear programming model with recourse. In two-stage stochastic optimisation, uncertainty is represented by random variables of which the distribution is known or can be estimated. The general idea is that decisions that need to be taken now (first stage), should be based on data that is available now, while taking into account unknown future observations (second stage) (Ahuja et al., 2009). Decisions are thus optimised while considering possible scenarios and their probability, but without knowing which scenario will eventually occur. In this case, crew demand will be treated as a stochastic input and this is made possible by analysing the airline's historical demand data and finding the best fitting probability distribution. From these distributions, correlated demand scenarios are constructed by using a Latin hypercube sampling approach based on Magini et al. (2019). Lastly, in order to allow the model to react to fluctuating demand, temporary employees can be hired, transitions between some positions are allowed, and firing permanent crew members can be made possible.

Paper Contribution

This paper is the first to use a stochastic programming approach for the strategic airline crew planning problem. Where previous research considered crew demand as a fixed, deterministic input, this paper acknowledges the high level of uncertainty that goes hand in hand with long-term planning by treating crew demand as a stochastic input. Hence, the model makes crew planning possible before the flight schedule is known. On top of this, it is the first time the Latin hypercube sampling (LHS) approach, as used by Magini et al. (2019), is used in an airline planning problem. Since demand at different crew positions is found to be correlated, the LHS approach allows to take this correlation into account when generating different crew demand scenarios.

While the model is developed and tested for a charter and holiday airline, the formulation is deliberately kept as generic as possible so that it can also serve as a useful strategic planning tool for more traditional airlines. All costs can easily be adapted and most constraints can be switched on or off, or can be modified, to adjust the model to the user's needs.

Paper Structure

This paper is structured as follows. Section 2 gives an overview of the relevant literature. Section 3 presents the used methodology and model formulation in a detailed way. Section 4 elaborates on the case studies, used to validate the model and show its possibilities, and discusses the results per case. Finally, Section 5 draws conclusions from the obtained results and provides recommendations for further research.

2 Relevant Literature

The overall objective of the workforce planning process is to have the right number of people with the right skills at the right time. This is done by making forecasts of supply and demand, taking decisions that will close the gap between these two and by making optimal work schedules (Altenstedt et al., 2017). Sohoni et al. (2004) call the workforce planning problem one of the most important ones in the airline industry. However, most articles focus on the actual construction and optimisation of crew schedules, and not on the steps that precede this phase: how many employees will be needed, how many will retire, how many need to be hired, at what position are they needed, etc. Below, the relevant literature with respect to strategic workforce planning is given, both for airlines as for other industries. Next, an overview is given of the literature concerning forecasting methods and two-stage stochastic programming, the solution approach used in this paper.

2.1 Strategic Workforce Planning

As stated before, very little research has been done in strategic workforce planning for airlines. Safarishahrbijari (2018) analyses 275 papers on workforce forecasting models and concludes that since 1980, only two percent of research in this field is done in the marine and airlines industry.

In articles that do investigate the strategic workforce planning for airlines, the objective always is to minimise costs and for this mixed-integer linear programming (MILP) models are mostly used (Verbeek, 1991; Holm, 2008; Morén, 2012). Differently, Hooijen (2019) uses a heuristic planning model. All of these articles also use the transitioning of crew from one position to another to close the gap between supply and demand. In addition, Morén (2012) allows for pilots to fly below rank and to fly on multiple aircraft types. However, all of the aforementioned articles consider crew demand as a given, deterministic input and therefore the implication of different demand realisations are not considered in the solution. On the other hand, the approach presented in this paper treats crew demand as a stochastic input by generating multiple demand scenarios based on historical data. This also means the optimal crew size is determined using stochastic optimisation, rather than deterministic optimisation.

Also in other industries, such as the healthcare sector and the army, strategic workforce planning is very important. In most articles, multiple objectives are identified and thus other methods are used next to MILP models to determine the optimal staff size: goal programming (Trivedi, 1981; Horn et al., 2016), using a multiple objective linear programming (MOLP) model (Li et al., 2007), using chance constraints (Ganguly et al., 2014) or using a combination of techniques (Weigel and Wilcox, 1993). Here too, staff demand is always used as an input and either based on historical data, or just a parameter without more details given.

2.2 Forecasting Methods and Two-Stage Stochastic Programming

Airlines usually do not like to disclose their workforce forecasting methods and therefore not a lot of literature has been written on this topic. Ciriani et al. (2013) identify two possibilities to determine crew demand: the airline's long-term fleet plan when looking more than a year ahead, and the flight schedules as soon as they are known (usually a year before the day of operation). Holm (2008) also identifies the possibility of estimating demand based on crew utilisation. In this case, an aircraft type's expected number of block hours, based on either the (preliminary) flight schedule or on historical data, is divided by the expected crew utilisation to come to a rough crew demand estimate (Ciriani et al., 2013; Yu et al., 2003).

Dealing with the high demand uncertainty that comes naturally with strategic planning is often done by using two-stage stochastic programming. Ng et al. (2008) compared six different approaches for human resource planning and concluded that two-stage stochastic programming provided the best results to find optimal staffing levels before the employees' attendance rates are known. Zhu and Sherali (2009) present a two-stage stochastic mixed-integer programming model to tackle fluctuating and uncertain workforce demand for service centres in a 12 month horizon. They conclude that compared to the deterministic model, the stochastic model proposes significantly fewer alterations to the prescribed workforce plan. Stochastic programming is also used to increase solution robustness. Yen and Birge (2006) use this approach to realise a more robust airline crew schedule in the operational planning phase that can better withstand disruptions. Bard et al. (2007) on the other hand use it to deal with wide workforce demand fluctuations in a distribution centre.

3 Methodology

The goal of the model is to provide a cost-optimised crew composition plan for the cockpit crew of a multi-fleet airline. This plan needs to be provided already in the strategic planning phase. At this point in time, the flight schedule is unknown and thus the crew demand is still uncertain. However, the airline already wants to know how many permanent crew they will need to have on each crew position, so that the long process of recruiting, hiring and training permanent crew members can already be started. At a later point in time, when the destinations and flight schedules are determined, the crew demand becomes known. If at this point it is observed that crew supply and demand do not match, it can be decided to hire temporary crew, or to transition crew from one position to another. In the ultimate case, the airline can consider firing permanent crew members to match supply and demand.

In order to solve this problem, a two-stage stochastic modelling approach is used. This approach allows to take into account the uncertain *stochastic* nature of the crew demand, and to make decisions at two moments in time (*stages*). In two-stage stochastic models, the uncertainty is represented by random variables of which the distribution is known or can be estimated. The general idea is then that decisions that need to be taken now (first stage), should be based on data that is available now, while taking into account unknown future observations (second stage) (Ahuja et al., 2009). Decisions are thus optimised while considering possible scenarios and their probability, but without knowing which scenario will eventually occur. Decision variables are split up in first-stage variables and second-stage variables.

In this case, in the first-stage it is decided how many permanent crew there should be per crew position and per time step. The model also has the option to hire new permanent crew members in the first stage. In the second stage the model has the possibility to transition crew to other positions, or hire temporary crew; both these possibilities come with a recourse cost. The model's second stage is also extended to allow the possibility of firing crew members. This option also comes with a cost. The main assumptions are listed below:

- The output is provided in full-time equivalents (FTE). This gives the airline the possibility to work with part-time crew.
- The time step of the model is one month. Hence, the model's output provides an overview of the optimal number of crew per month and per position, and gives insight into what decisions to take each month in different demand scenarios.
- The fleet size is constant throughout the simulation period. A changing fleet size would influence the crew demand.
- There is a new hire capacity. New hires need to be trained and this requires resources (money, staff, simulators, classrooms, etc.). Since these resources are limited, the number of new hires is limited accordingly.
- There is a natural outflow rate of staff. This rate accounts for crew retiring or leaving the airline for other reasons.

In Figure 1 the high-level flowchart for the solution process is given. The different steps are explained in detail in the subsections below. Subsection 3.1 explains the first two blocks of Figure 1: reading and preparing all input parameters and data. Next, Subsection 3.2 presents the analysis of the different factors contributing to crew demand. Subsection 3.3 explains how the demand scenarios are generated that are used in the model. Then, in Subsection 3.4 the mathematical model formulation can be found, used to construct the stochastic model. Finally, Subsection 3.5 explains how the solution is obtained and how the results need to be interpreted.

3.1 Reading and Preparing Input Data

The first step in the model is to load and prepare all the needed input parameters and data. The input parameters, such as the number of aircraft types and crew positions, salary costs, training costs, initial crew pool size, etc. are defined by the user and can easily be changed. These parameters will be discussed in detail in Subsection 3.4. The historical crew data and block hour data on the other hand are

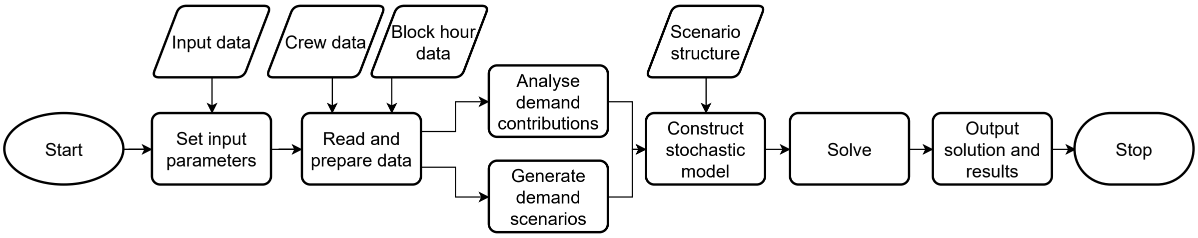

Figure 1: Flowchart of the proposed model.

provided by the airline. The former gives a historical overview of each crew member's daily scheduled activity and their type rating, i.e. which aircraft type(s) they are qualified to fly on. This is needed to later analyse the contributions of each of these activities to the total crew demand (Subsection 3.2). Lastly, the block hour data contains the number of block hours flown for each day and each crew member, together with the aircraft type on which the block hours were flown. This data is needed in order to generate demand scenarios (Subsection 3.3).

Preparing the data consists of removing empty or incomplete entries, removing irrelevant data (e.g. crew age and name), converting all dates to the same notation and filtering the data to the desired time period. This way the data is ready to be used further on in the model.

3.2 Analysing Demand Contributions

The demand for crew comes primarily from the need to operate flights. Besides this, there are many other factors that increase the number of needed crew. In Figure 2 an overview is given of these factors for the airline in question; they are not to scale. Here it can be seen that crew can be scheduled to many other duties besides flights, such as standby, positioning (i.e. *deadheading*), ground duties or training courses. Besides this, crew can be absent or be scheduled to have a holiday or legally required weekend days. These contributions to total crew demand are approximated by analysing historical data from the airline. For each day, it is calculated what percentage of the crew is assigned to each category mentioned in the bottom bar of Figure 2. This results in values indicating relatively how much crew is assigned to each category each day. The values can then be averaged per month. It is found that the values are rather constant over the years. For this reason it is assumed that the values observed in the past for a respective month are also representative for the future. However, one category in Figure 2 is not considered here: the crew demand for flights is treated as a stochastic variable and is generated using the method explained in the next subsection.

Figure 2: Overview of the different factors contributing to crew demand. Not to scale.

3.3 Generating Demand Scenarios

As can be seen in Figure 1, generating crew demand scenarios is done independently from the demand contribution analysis. This is because the generated scenarios contain the crew demand *for flights*. However, the demand contribution analysis is done for all other categories mentioned in Figure 2. Multiple crew demand scenarios are generated to serve as an input to the two-stage stochastic model.

This is done by following five steps:

1. Analyse the daily historical block hours, per month and per crew position and find the best fitting probability distribution.
2. Convolute the daily block hour distributions to come to a monthly distribution per crew position.
3. Analyse the correlation between block hour demand at the different crew positions.
4. Apply Latin hypercube sampling to generate demand scenarios with correlated variables.
5. Multiply the demand scenarios with a trend forecast to account for future demand increase or decrease.

Step 1: Block Hour Analysis Block hours are the time between an aircraft leaving the departure gate and arriving at the destination gate. The number of block hours a pilot may fly every year is bounded by regulations (EU, 2014). Knowing this, it is assumed that the demand for pilots can be computed by estimating the total block hours needed in the future.

The airline's historical daily block hours are analysed per crew position (6) and per month (12), resulting in 72 histograms. An example can be seen in Figure 3a. Next, the best fitting mathematical probability distribution is selected out of nine possible options: the Normal, Gamma, Beta, Chi, Chi-squared, Cauchy, Exponential and Weibull distributions. These distributions are chosen since they are well-researched and hence well-understood. Furthermore, together they can take a wide range of shapes and this increases the chance of finding a fitting distribution. The best fitting distribution is assumed to be the one with with the smallest residual sum of squares (RSS). This is calculated using Equation 1, where y_i is the variable to be predicted, and $f(x_i)$ the predicted value of y_i (Draper and Smith, 1998).

$$RSS = \sum_{i=1}^{n} (y_i - f(x_i))^2 \tag{1}$$

It is found that the Beta probability distribution is the most frequent best fitting distribution. The Beta distribution can take many shapes, is supported on a fixed range and its location and range can be altered by means of two extra parameters (Johnson et al., 1995). For this reason, the daily block hours are assumed to be Beta distributed (Figure 3b) and its characteristics are estimated for each of the 72 distributions. Furthermore, by assuming the Beta distribution for all cases, the implementation of the next step is made easier, since the convolution needs to be programmed for only one type of probability distribution.

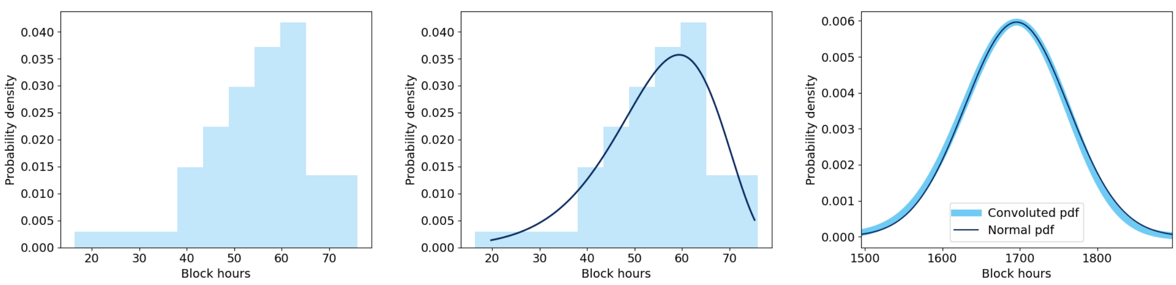

(a) Original histogram of the data. (b) Beta probability distribution fitted to the data. (c) Convoluted Beta distributions and theoretical normal distribution.

Figure 3: Example process of going from an empirical block hour distribution to a convoluted one, approximating a Gaussian distribution. Data comes from the airline between October 2016 and February 2020. This example shows data for the CP SH crew position in an arbitrary month.

Step 2: Convolution of Daily Block Hour Distributions Since the model works with a time step of one month, these daily block hour distributions need to be transformed into a monthly distribution. This is done by taking the convolution of the obtained Beta distributions for each day of the respective month, because the sum of random variables is the convolution of their individual distributions (Holmes, 1998). As can be seen in the example of Figure 3c, the convoluted Beta probability density function approximates a normal distribution. This can be expected from the central limit theorem (Montgomery and Runger, 2010). For this reason, it is safe to assume the monthly block hour distributions are normally distributed. For every of the 72 distributions, the mean and standard deviation are determined. The mean μ is equal to the block hour value (x) where the probability density (y) is the highest (y_{max}). The standard deviation σ can then be derived from the probability density function (pdf) for a normal distribution (Johnson et al., 1995):

$$f(x) = \frac{1}{\sigma\sqrt{2\pi}} e^{-\frac{1}{2}\left(\frac{x-\mu}{\sigma}\right)^2} \tag{2a}$$

At mean μ, we have: $x = \mu$, and $f(x) = y_{max}$, so that:

$$y_{max} = \frac{1}{\sigma\sqrt{2\pi}} \Leftrightarrow \sigma = \frac{1}{y_{max}\sqrt{2\pi}}$$

Since y_{max} is known, both the mean and standard deviation can easily be calculated. These are used in Step 4.

Step 3: Correlation Analysis It is important to analyse the correlation between the block hour distributions. If the block hour distributions are correlated, it is necessary to generate correlated demand scenarios too. For this, the Pearson correlation coefficient is calculated between every crew position for every month. This coefficient measures the linear correlation between two data sets and takes values between -1 and 1 to indicate a negative or positive correlation respectively. The pearson correlation is calculated by using Equation 3, with r_{xy} the Pearson correlation coefficient, n the sample size, x_i and y_i the individual sample points of the two sets under investigation, and $\bar{x}$ and $\bar{y}$ the sample means (Ross, 2014).

$$r_{xy} = \frac{\sum_{i=1}^{n}(x_i - \bar{x})(y_i - \bar{y})}{\sqrt{\sum_{i=1}^{n}(x_i - \bar{x})^2}\sqrt{\sum_{i=1}^{n}(y_i - \bar{y})^2}} \tag{3}$$

By calculating r_{xy} for every combination of crew positions, a correlation matrix is obtained for each month. An example can be found in Figure 4 for the month of April.

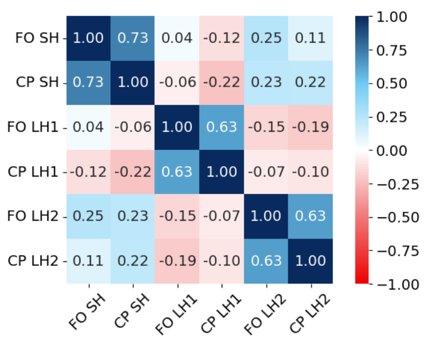

Figure 4: Example of a Pearson correlation matrix for block hour values for different crew positions. Data comes from the airline between October 2016 and February 2020. This example shows data for April.

It is clear that the block hour distributions are indeed correlated, both positively and negatively, between the different crew positions. As expected, the highest correlations can be found between FO and CP on the same aircraft type.

Step 4: Applying Latin Hypercube Sampling Using the monthly block hour distributions obtained in Step 2 and the correlation matrices obtained in Step 3, we can now generate the demand scenarios needed for the model. This is done by using Latin hypercube sampling (LHS), an approach first described by McKay et al. (1979). LHS generates samples from multi-dimensional distributions by dividing them into k intervals with the same probability. k sample points are then put into a Latin hypercube, i.e. one sample for each dimension. When dealing with correlated variables however, some extra steps are needed. Here, we use the approach explained by Golub and Van Loan (1996) and Ross (2014), and used in Magini et al. (2019) to come to correlated samples. This approach uses lower triangular Cholesky decomposition to impose the desired correlation on the samples. The inverse transformation method is applied to ensure that the generated samples have the same distribution as the original distributions (Ross, 2014). The detailed steps are explained below (Magini et al., 2019):

1. Create a random $(k \times n)$ matrix $\mathbf{Z}^*$, with k the number of scenarios and n the number of crew positions. This matrix contains k Latin hypercube samples of size n from a standardised normal distribution. The correlation matrix $\mathbf{I}^*$ of the samples and the identity matrix $\mathbf{I}$ do not coincide, hence they are not independent. To induce the desired correlation, the the $(k \times n)$ matrix $\mathbf{Z}$ is created by using lower triangular Cholesky decomposition (Golub and Van Loan, 1996):

$$\mathbf{I} = \mathbf{C} \cdot \mathbf{C}^{\mathrm{T}} \tag{4}$$

$$\mathbf{I}^* = \mathbf{E} \cdot \mathbf{E}^{\mathrm{T}} \tag{5}$$

$$\mathbf{Z} = \mathbf{Z}^* \cdot \mathbf{C} \cdot \mathbf{E}^{-1} \tag{6}$$

   $\mathbf{Z}$ contains k independent samples of size n from a standardised normal distribution.

2. Create a random $(k \times n)$ matrix $\mathbf{G}$. This matrix contains k samples from a standardised normal distribution with the previously obtained correlation matrix $\mathbf{B}$ containing the correlation between crew positions. In other words, $\mathbf{B}$ contains the *desired* correlation. By applying lower triangular Cholesky decomposition, this desired correlation is induced in $\mathbf{Z}$:

$$\mathbf{B} = \mathbf{P} \cdot \mathbf{P}^{\mathrm{T}} \cdot \mathbf{G} = \mathbf{Z} \cdot \mathbf{P} \cdot \mathbf{C}^{-1} \tag{7}$$

3. Apply the inverse transformation method to $\mathbf{G}$ to create $(k \times n)$ matrix $\mathbf{D}$. The matrix $\mathbf{D}$ complies with the desired marginal distributions at each crew position. The inverse transformation method states that applying the inverse cumulative distribution function of any distribution F to a random variable with $U(0, 1)$ distribution results in a random variable whose distribution is exactly F (Ross, 2014). In this case, F is always a normal distribution, since we know that the monthly block hours are normally distributed. The correlation matrix of the generated samples is now equal to the desired correlation matrix $\mathbf{B}$.

In Figure 5 an example of the result of this process is shown for the block hour distributions of the FO SH and CP SH crew positions in an arbitrary month. For each month, the correlation values of the generated samples are verified to be the same as the correlation values from Step 3.

We apply this LHS process per month in n dimensions (with n the number of crew positions) and this results in a set of k correlated samples (i.e. scenarios) per month, each with the same probability p_k. Each scenario specifies the expected block hour demand per crew position. The block hour values are then divided by the legal maximum of block hours per crew member per month, to come to crew demand values in FTE.

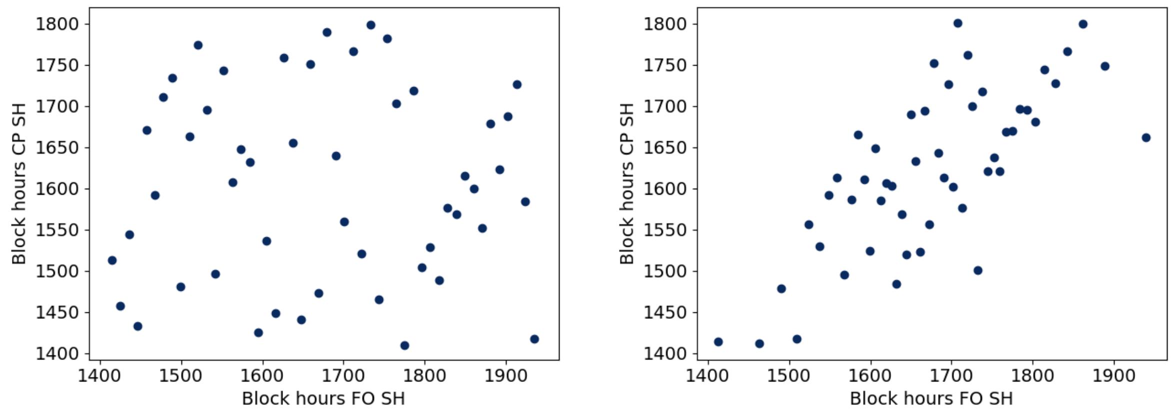

(a) LHS without correlation between the random variables.

(b) LHS with correlation between the random variables.

Figure 5: Example of two-dimensional LHS with 50 samples, illustrating the difference between LHS without and with (positive) correlation respectively.

Step 5: Applying Trend Forecast The last step is to apply a trend forecast to the generated crew demand scenarios. Since the scenarios are generated using historical data, applying a trend forecast accounts for the expected future demand increase or decrease. This is done by multiplying all the generated demand values with a trend value. This trend value can come from the airline itself (e.g. by decomposing the block hour data into trend and seasonality), or from professional forecasts available on the market. In this paper, the demand trend forecasts of IATA and EUROCONTROL are used, since no clear trend is visible when decomposing the airline's block hour data. The forecasts of both IATA and EUROCONTROL are widely used in the airline industry, and are also used by the airline in question to forecast future crew demand.

3.4 Model Formulation

A two-stage stochastic linear programming model with recourse is used to determine the cost-optimal crew composition for each time step. This is done by considering various crew demand scenarios and factors influencing total demand, discussed in the previous subsections. The objective and constraints of the model are discussed here. First, an overview is given of all the sets, variables and parameters and their notation, followed by the mathematical model formulation. The explanation of the model can be found at the end of this section.

Sets

T	set of time steps, indexed by t, where $T = \{t_{start}, ..., t_{end}\}$
H	set of time steps in which temporary FTE can be hired, where $H \subseteq T$
K	set of scenarios, indexed by k
P	set of crew positions, indexed by p
Q_p	set of crew positions which can transition to crew position p, where $Q_p \subseteq P$
R	set of crew positions that allow temporary FTE, where $R \subseteq P$

First-stage decision variables

h^l_{tp}	number of permanent FTE newly hired at time step t at position p, $t \in T$, $p \in P$
x_{tp}	number of permanent FTE employed at time step t at position p, $t \in T$, $p \in P$

Second-stage decision variables

h^s_{tp}	number of temporary FTE newly hired at time step t at position p, $t \in H$, $p \in R$
r_{tp}	number of permanent FTE fired at time step t at position p, $t \in T$, $p \in P$
y_{tp}	number of temporary FTE employed at time step t at position p, $t \in T$, $p \in P$
z_{tpq}	number of permanent FTE transitioned at time step t from position p to position q, $t \in T$, $p \in P$, $q \in Q_p$, $p \neq q$

Parameters

a	fixed new hire capacity in FTE for one time step t
c^{fire}_p	firing cost of one permanent FTE at position p, $p \in P$
$c^{salary,l}_p$	salary cost of one permanent FTE at position p for one time step t, $p \in P$
$c^{salary,s}_p$	salary cost of one temporary FTE at position p for one time step t, $p \in P$
$c^{training,i}_p$	average initial crew training cost at position p for one time step t, $p \in P$
$c^{training,r}_p$	average recurrent crew training cost at position p for one time step t, $p \in P$
$c^{transition}_{pq}$	average cost of one permanent FTE transitioning from position p to position q, $p \in P$, $q \in Q_p$
d_{tpk}	gross crew demand in FTE for position p at time t in scenario k, $k \in K$, $t \in T$, $p \in P$
$f^{absence}$	average daily historical fraction of absent crew
f^{unused}	average daily historical fraction of crew with an unused production day
$f^{non-regular}_t$	average daily historical fraction of crew performing non-regular activities
$f^{training}_t$	average daily historical fraction of crew performing training, in time step t, $t \in T$
$f^{vacation}_t$	average daily historical fraction of crew on holiday, in time step t, $t \in T$
$f^{weekend}_t$	average daily historical fraction of crew with a rostered weekend day, in time step t, $t \in T$
$f^{position}_{tp}$	average daily historical fraction of crew at position p with a rostered hotel or positioning day, in time step t, $t \in T$, $p \in P$
$f^{standby}_{tp}$	average daily historical fraction of crew at position p with a rostered standby day, in time step t, $t \in T$, $p \in P$
l^l	length of a permanent contract expressed in time steps t
l^r_{pq}	length of a transition course in days from position p to q, $p \in P$, $q \in Q_p$
l^s	length of a temporary contract expressed in time steps t
n_t	number of days in time step t, $\forall t \in T$
p_k	realisation probability of scenario k, $k \in K$
u^l	natural outflow rate of permanent crew members per time step t
x^{start}_p	number of permanent FTE at position p at $t = t_{start}$, $p \in P$

Two-stage stochastic model

$$\text{Minimise} \sum_{t\in T}\sum_{p\in P}\left(c^{salary,l}_p + \frac{c^{training,i}_p}{l^l} + c^{training,r}_p\right) \times x_{tp} + \sum_{k\in K} p_k \times f_k(x) \tag{8a}$$

where, for each scenario $k \in K$, we have

$$f_k(x) = \min \sum_{t\in T}\sum_{p\in P}\left(\left(c^{salary,s}_p + \frac{c^{training,i}_p}{l^s}\right) \times y_{tp} + c^{fire}_p \times r_{tp}\right) + \sum_{t\in T}\sum_{p\in P}\sum_{q\in Q_p} c^{transition}_{pq} \times z_{tpq} \tag{8b}$$

subject to

$$\textit{Crew balance:} \qquad x_{tp} = \begin{cases} \left(1-u^l\right) \times x_p^{start} + h_{tp}^l - r_{tp} + \sum\limits_{q \in Q_p} z_{tqp} - \sum\limits_{q \in Q_p} z_{tpq} & \text{if } t = t_{start} \quad \forall p \in P \\ \left(1-u^l\right) \times x_{t-1,p} + h_{tp}^l - r_{tp} + \sum\limits_{q \in Q_p} z_{tqp} - \sum\limits_{q \in Q_p} z_{tpq} & \text{if } t \in T : t \neq t_{start}, \forall p \in P \end{cases} \tag{8c}$$

$$\textit{Demand:} \qquad x_{tp} + y_{tp} - d_{tpk} - \frac{l_{qp}^r}{n_t} \times \sum_{q \in Q_p} (z_{tqp}) - \Bigg(\frac{1}{n_t} \Big(f_{tp}^{standby} + f_{tp}^{position} + f_t^{non-regular} + f_t^{training}$$

$$+ f_t^{vacation} + f_t^{weekend} + f^{absence} + f^{unused} \Big) \Bigg) \times (x_{tp} + y_{tp}) \geq 0$$

$$\forall k \in K, \forall t \in T, \forall p \in P, \forall q \in Q_p \tag{8d}$$

$$\textit{Temporary:} \qquad y_{tp} = \begin{cases} \sum\limits_{i=t_{start}}^{t} h_{ip}^s & \text{if } t \in [t_{start}, l^s - 1] \quad \forall p \in P \\ \sum\limits_{i=t-l^s+1}^{t} h_{ip}^s & \text{if } t \in [l^s, t_{end}] \quad \forall p \in P \end{cases} \tag{8e}$$

$$\textit{Capacity:} \qquad \sum_{t \in T} \sum_{p \in P} \left(h_{tp}^l + h_{tp}^s\right) - \sum_{t \in T} a \leq 0 \tag{8f}$$

$$\textit{Variables:} \qquad h_{tp}^l, h_{tp}^s, r_{tp}, x_{tp}, y_{tp} \in \mathbb{R}_+^2 \quad \forall t \in T, \forall p \in P$$

$$z_{tpq} \in \mathbb{R}_+^3 \quad \forall t \in T, \forall p \in P, \forall q \in P \tag{8g}$$

The objective function (8a) minimises the total crew costs over all the considered crew positions and time steps. These costs consist of salary costs and initial and recurrent training costs, and they are dependent on the crew position. The initial training costs are divided by the contract length (l^l and l^s) to come to a cost per time step, while the recurrent training costs are calculated by taking into account the duration, staffing costs, simulator costs, etc. The recurrent training costs disappear in the cost calculation for the second stage (8b), since temporary employees do not need recurrent training. On the other hand, a firing cost is added in the second stage and multiplied with the number of fired FTE. The last term in the second stage accounts for the total costs of transitioned crew members. This transition cost only occurs if a crew member is transitioned to a lower-paid position. In this case, the transition cost is the difference in salary. In case a person is moved to a higher-paid position, the model accounts for this in the calculation of the salary costs. Finally, the second-stage costs are multiplied with the scenario probability.

The first constraint (8c) ensures that there is a crew balance for permanent employees. It does so by respectively calculating the natural outflow from the previous time step, increasing it by the number of FTE hired, decreasing it by the number of FTE fired, increasing it with the number of FTE transitioned *to* that crew position, and decreasing it with the number of FTE transitioned *from* that crew position. Constraint (8d) ensures that the demand is satisfied. The first two terms are the number of permanent and temporary FTE, the third term is the gross crew demand in scenario k (so the number of crew needed *only* for flights). The fourth term provides a demand penalty for a transition to a crew position, since transitioned crew has to follow a training course and is not deployable for a period of time. The length of this course depends on which transition is made. The entire fifth term increases the size of the crew pool to account for all the different factors discussed in Subsection 3.2. Constraint (8e) defines the total number of temporary FTE in terms of hired FTE and takes into account the temporary contract length (l^s). In fact, this constraint assures a crew balance for temporary employees similar to constraint (8c), only now there are no natural outflow and transitions but only new hires. Constraint (8f) limits the number of new hires for the entire planning period. This limit mainly depends on the availability of training staff and is discussed with the airline. Finally, constraint (8g) states that all decision variables are non-negative real numbers. They are not set to be integers, since they are expressed in FTE and this way the airline has the liberty to work with part-time employees. Depending on the airline using this tool, most of these constraints and parameters can easily be adjusted.

3.5 Model Solution and Output

Once the model is fully defined and constructed, its solution is obtained by means of the primal simplex method. This is a very common algorithm for solving linear optimisation problems, and used by many commercial solvers (Saunders, 2019; Hillier and Lieberman, 2015). This results in an optimal objective value, consisting of both first-stage and second-stage costs. This is calculated using Equation 9, where OF_{tot}, OF_1 and OF_2 are the total, first-stage and second-stage objective value respectively, and p_k is the scenario probability.

$$OF_{tot} = OF_1 + \sum_{k \in K} OF_2 \times p_k \tag{9}$$

The output also provides the optimal values of the decision variables per time step for each scenario. Since it is a two-stage optimisation problem, it naturally follows that the first-stage decision variables values are the same for all scenarios, whereas the second-stage decision variables values depend on the scenario. This output gives the user an overview of the optimal permanent crew composition plan for each time step, and helps them with taking decisions once the crew demand becomes known.

4 Case Studies

Three case studies are used to validate the model and show its options and capabilities. First, a historical real-world case is solved and compared to the airline's solution if the same generated demand is considered. By doing so, the developed model is validated and conclusions can be drawn with respect to the model's robustness, solution flexibility and cost. Next, the second case allows the model to consider firing crew members to see if this can further improve solutions. This case also investigates if the airline is overstaffed at certain crew positions. Finally, the model is used to solve a summer and winter season combined in the recovery phase of the Covid-19 pandemic. This case is used to both test its ability to deal with the transition from one season to another, and to explore if the model is useful in the post-Covid-19 period.

For all three cases, real-world data is used from a European airline with both scheduled and chartered flights. The airline uses three aircraft types: one type for short-haul (SH) and two for long-haul flights (LH1 and LH2). There are two cockpit positions: First Officer (FO) and Captain (CP). The year is divided in a summer season (from April to October) and a winter season (from November to March). The airline allows to hire temporary employees only on the FO SH position in April, with a contract length (l^s) for the entire summer season. Transitions are only allowed between FO SH and CP SH, and between FO SH and FO LH1. The time step for all three cases is one month. This means that all input (salary costs, training costs, demand scenarios, etc.) is given per month and that the resulting values of the decision variables are also provided per month. Hence, the model gives a monthly insight into the optimal decisions to be taken by the airline regarding hiring new hires, transitions or dismissing crew.

Table 1 gives an overview of the permanent and temporary contract lengths, the natural outflow rate, the monthly new hire capacity and the transition course length. These parameters are constant for all positions and case studies and are taken from the airline's records and interviews with the airline's crew planners. Table 2 gives an overview of all the costs: salary, training and transition costs. These costs are averages and calculated using the airline's labour agreement and interviews with the airline's crew training schedulers. These parameters are also constant for all three case studies, but depend on the crew position. Parameters that are case-specific are discussed separately below.

Table 1: Input parameters constant for all crew positions and case studies.

l^l [months]	l^s [months]	u^l [%/month]	a [FTE/month]	l^r [days]
420	7	0.83	2.7	3

For each case study the solution is obtained after 20 repetitions and each repetition 10 demand scenarios are generated, hence $k = 10$ and $p_k = 0.1$ for each k. These demand scenarios contain the crew demand for each crew position and each month in the simulation period. In Table 3 this process is illustrated.

Table 2: Input parameters depending on the crew position, but constant for all case studies. MU are monetary units.

Position (p)		$c_p^{salary,l}$ [MU/month]	$c_p^{salary,s}$ [MU/month]	$c_p^{training,i}$ [MU]	$c_p^{training,r}$ [MU/month]	$c_{qp}^{transition}$ ($\forall q \in Q$) [MU]
FO	SH	42.4	20.0	125.0	0.9	13.1
	LH1	42.4	/	125.0	0.8	13.1
	LH2	42.4	/	150.0	1.6	/
CP	SH	55.5	/	125.0	0.9	0.0
	LH1	55.5	/	125.0	0.8	/
	LH2	55.5	/	150.0	1.6	/

There are 20 solution repetitions, and in each repetition 10 crew demand scenarios are generated. This results in first-stage decision variables (DVs) that are the same for the 10 scenarios, and in second-stage DVs that vary per scenario. The first-stage objective function OF_1 is constant for all scenarios within a repetition, while the second-stage objective function OF_2 depends on the scenario. The total objective function per repetition OF_{tot} is then calculated with Equation 9. After 20 repetitions, the overall total objective function is calculated by taking the average of the 20 obtained values. The same can be done for the decision variables.

Table 3: Illustration of how the results are obtained

Repetition (m)	Scenario (k)	1st-stage DVs	2nd-stage DVs	OF_1	OF_2	$OF_{tot,m}$
1	1		$h^s_{tp,1,1}; r_{tp,1,1}; y_{tp,1,1}; z_{tpq,1,1}$		$OF_{2,(1,1)}$	
	$\vdots$	$h^l_{tp,1}; x_{tp,1}$	$\vdots$	$OF_{1,(1)}$	$\vdots$	$OF_{1,(1)} + \sum_{k=1}^{10} OF_{2,(1,k)} \times 0.1$
	10		$h^s_{tp,1,10}; r_{tp,1,10}; y_{tp,1,10}; z_{tpq,1,10}$		$OF_{2,(1,10)}$	
$\vdots$	...	...	...	...	...	...
20	1		$h^s_{tp,20,1}; r_{tp,20,1}; y_{tp,20,1}; z_{tpq,20,1}$		$OF_{2,(20,1)}$	
	$\vdots$	$h^l_{tp,20}; x_{tp,20}$	$\vdots$	$OF_{1,(20)}$	$\vdots$	$OF_{1,(20)} + \sum_{k=1}^{10} OF_{2,(20,k)} \times 0.1$
	10		$h^s_{tp,20,10}; r_{tp,20,10}; y_{tp,20,10}; z_{tpq,20,10}$		$OF_{2,(20,10)}$	
Total OF over all repetitions: $\frac{1}{m} \sum_{m=1}^{20} OF_{tot,m}$						

It is chosen not to generate more than 10 scenarios since this would result in scenarios with extreme demand values, sampled from the outer edges of the normal distributions. These extreme values act as shock events and severely influence the solution. It is also found that running each case 20 times provides a stable average solution. This can be seen in Figure 6, where the progression of the average OF value is shown for all cases. It can be seen that after 20 solution repetitions, the average OF value is very stable. For all three case studies, the model's average solution time for one repetition is around 55 seconds.

All simulations are run on an Intel Core i7-4710MQ CPU with 8 GB RAM using Python 3.8 and Pyomo's PySP package, an open-source software package used for (two-stage) optimisation problems (Watson et al., 2012). In PySP, the scenario structure is defined and the whole model is constructed. The solution is obtained by means of the GLPK solver package. This is a linear programming kit that uses the primal simplex method by default to solve linear programming problems [1].

The rest of this section is structured as follows. Subsection 4.1 first shows the results of the crew demand contribution analysis, as explained earlier in Subsection 3.2. Next, Sections 4.2 and 4.3 show the results for the summer 2019 cases, without and with the option to fire permanent crew members respectively. Finally, Subsection 4.4 presents the results of Case 3, in which two consecutive seasons are solved at once.

[1] `https://www.gnu.org/software/glpk/`

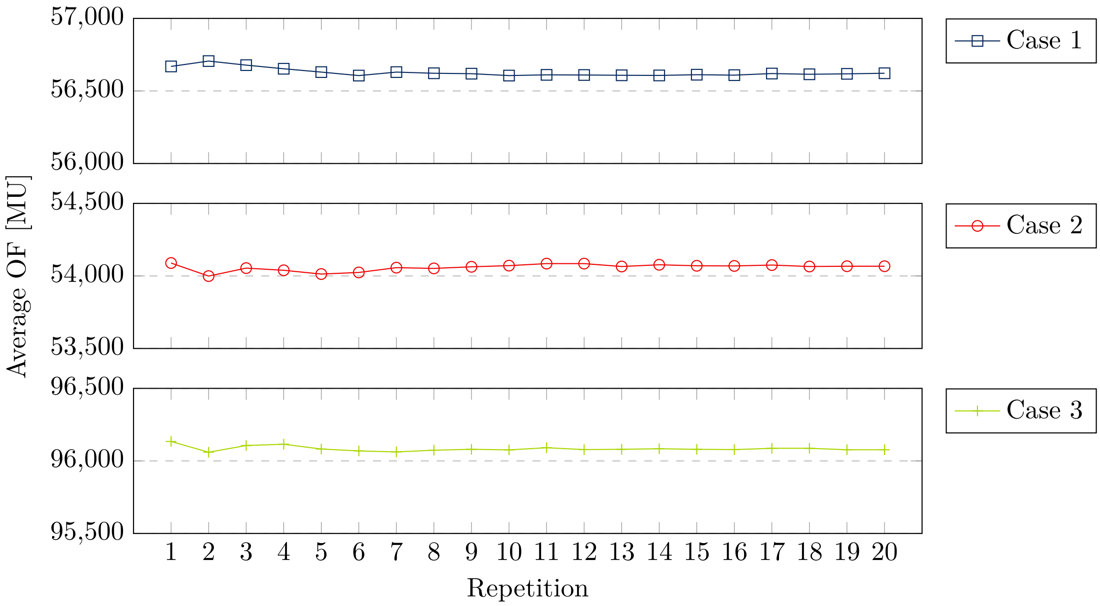

Figure 6: Progression of average OF value for all cases with increasing solution repetitions.

4.1 Crew Demand Contributions

The analysis of the crew demand contributions, as explained in Subsection 3.2, can be found in Figure 7. For this, four years of historical data is used. The graphs show the monthly average of the daily fraction of total crew assigned to the categories shown in the bottom bar of Figure 2. The left graph shows the results for two production related categories: *standby* and *hotel/positioning*. These results are split up per fleet type. The right graph shows the results for the other regular activities and non-regular activities. These results are analysed for all fleet types together, since it is found that these values do not differ much per type. It can be seen that legal weekend days, i.e. compulsory days off, are the biggest contributor to increased crew demand. This is because crew at the airline is entitled to 11 days off per 28 days. The results obtained here are used in the model to increase the total number of crew needed per position and per month. They are indicated by $f^{absence}$, f^{unused}, $f_t^{non-regular}$, $f_t^{training}$, $f_t^{vacation}$, $f_t^{weekend}$, $f_{tp}^{position}$ and $f_{tp}^{standby}$. For the sake of simplicity, it is assumed that the *absence* and *unused production day* contributions are constant throughout the year.

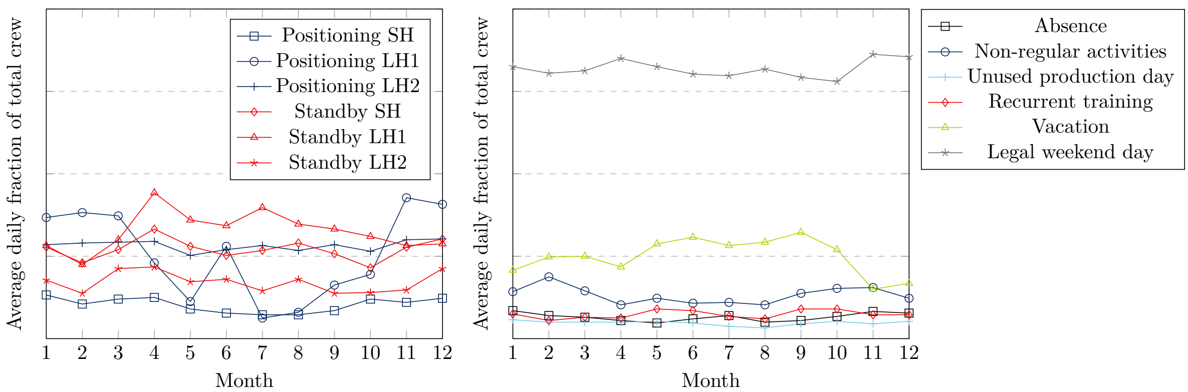

Figure 7: Average daily fraction of total crew assigned to different categories, per month. Both graphs have the same scale. Data comes from the airline between October 2016 and February 2020.

4.2 Case 1: Summer 2019 and Validation

To validate the model, the summer season of 2019 is solved and compared to the airline's solution in case the same generated demand scenarios are considered. To find the airline's solution, the airline's planned permanent FTE at the start of the summer 2019 season are fixed as first-stage variables. Next, the second stage is solved using the same demand scenarios as used in the model's summer 2019 solution. This way the model's solution and the airline's solution are compared in a fair way.

The summer 2019 season is used since this was the last full season before Covid-19 impacted airline operations. For the block hour analysis and scenario generation, block hour data from summer season 2015 up to and including summer season 2017 is used. This is because the model is meant as a strategic planning tool, hence the airline will use it three to four seasons in advance. For this case we do not allow the model to fire FTE yet, i.e. $r_{tp} = 0, \forall t \in T, \forall p \in P$. This is because firing permanent crew members is a very exceptional occasion at the airline.

4.2.1 Model Parameters

Since most parameters are already defined for all three cases, two case-specific parameters are left to define: the starting conditions for each crew position (x_p^{start}) and the trend value used in the demand scenario generation. In Table 4, the former parameter can be found. These values are the airline's *planned* permanent FTE values at the end of March 2019, so at the start of our simulation. To account for the future trend, the generated demand scenarios are multiplied with $(1.025)^2$, since 2.5% was the predicted European yearly increase in flight demand in 2016 (IATA, 2016). This value is squared since it is assumed that the analysis is done in 2017, two years in advance.

Table 4: Starting conditions and case-specific parameters for Case 1.

Position		x_p^{start} [FTE]	Position		x_p^{start} [FTE]
FO	SH	35.0	CP	SH	32.0
	LH1	18.9		LH1	10.5
	LH2	45.1		LH2	26.1

4.2.2 Results

In Table 5, the model's minimum, maximum and average objective value, and average stage solutions of 20 solution repetitions can be found. These results suggest that the model's average solution is 2.1% cheaper than the airline's solution when considering the same generated demand scenarios. This cost difference comes both from the first-stage and the second-stage costs: in both cases the model provides cheaper solutions than the airline. The second-stage cost is very low compared to the first-stage cost and this indicates that the model does not use the second-stage options frequently. In fact, for the model's solution, *all* second-stage costs come from transitions, and none from hiring temporary employees. This is shown in Table 6. In the airline solution on the other hand, 0.6 temporary FTE are hired, but here too the main contribution to the second-stage cost are transitions. Transition costs are lower than hiring temporary FTE, and that is why the second-stage costs are low compared to the first-stage costs. This also indicates that there is enough crew in total to meet demand, but they are not stationed at the optimal positions.

Table 5: Objective value results for Case 1 for both the model and the airline.

Case	Objective value [MU]				First-stage cost [MU]	Second-stage cost [MU]
	Min	Max	Average	Δ(min,max) [%]	Average	Average
Model	56483	56804	56622	0.57	56532	90
Airline	57753	58087	57862	0.58	57710	152

Table 6: Average number of temporary FTE hired in April on the FO SH position for the model's and airline's solution.

Position		Temporary FTE hired in April (avg)	
		Model solution	Airline solution
FO	SH	0.0	0.6

Figures 8 and 9 illustrate why the model provides a cheaper solution than the airline. The blue plots show the model's generated demand scenarios for each month, the red plots show its supply solution and the green plots show the airline's solution in case of the same generated demand. It can be seen that for the CP positions (Figure 8), the model's supply solution is lower than the airline's for every month. The CP positions are the biggest cost for the airline and this difference between the model's and airline's solution is the main cause of the airline's higher objective value in Table 5. For the CP SH and FO SH positions (8a and 9a), the model manages the demand fluctuations very well. This is not the case for the CP LH1 and CP LH2 positions (8b and 8c). At these positions, no transitions are allowed and thus the model is limited to the natural outflow to lower the supply. Here, the difference between demand and supply suggests that the starting conditions are too high and overestimate the demand. Figures 9a and 9b show that the model indeed uses the allowed transitions between FO SH and FO LH1. A clear decline in supply is visible in July for the FO LH1 position, while the supply increases in the same month for the FO SH position. Figure 9c shows that the model also reacts to a demand increase from one month to another in case no transitions or temporary employees are allowed, as is the case for the FO LH2 position. In this case, new permanent employees are hired, as can be seen in the graph for the months of July and August. It can also be seen that in all graphs, the range of demand scenarios is far larger than the range of supply solutions. This is because the supply fluctuation is restricted by natural outflow, a new hire capacity and the fact that temporary crew members can only be hired on the FO SH position in April.

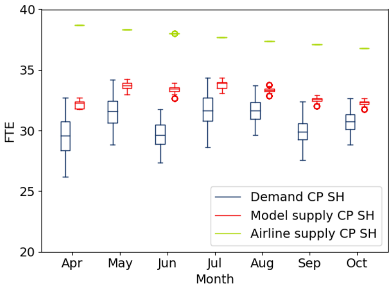

(a) CP SH.

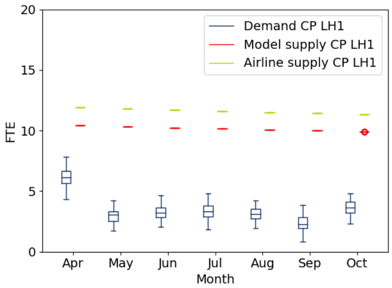

(b) CP LH1.

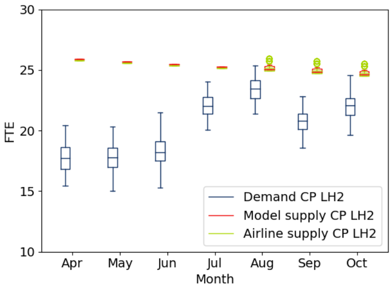

(c) CP LH2.

Figure 8: Boxplots of demand generated by the model, its supply solution, and the airline's (second-stage) supply solution for the same demand scenarios, for CP positions for Case 1.

The airline's higher supply also results in a more rigid solution. From Figures 8 and 9 it can be seen that the airline often provides more crew than needed to respond to the generated crew demand. This results in an expensive, rigid solution with almost no room for fluctuations.

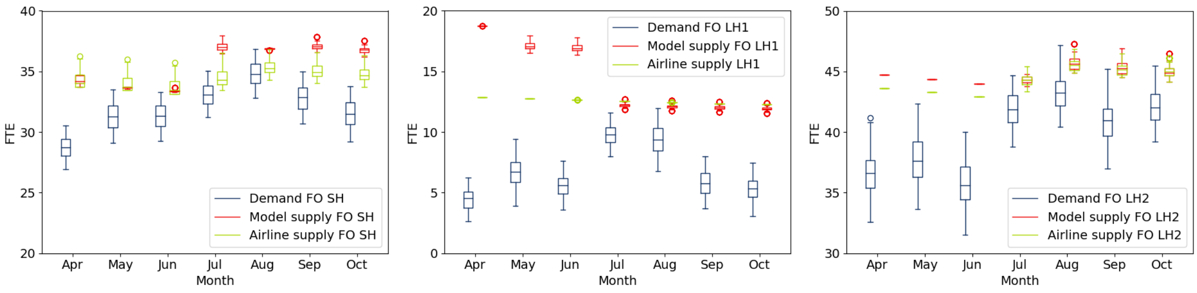

(a) FO SH. (b) FO LH1. (c) FO LH2.

Figure 9: Boxplots of demand generated by the model, its supply solution, and the airline's (second-stage) supply solution for the same demand scenarios, for FO positions for Case 1.

4.3 Case 2: Lay-off Possibility

Since there are indications that the airline is overstaffed, we now activate the model's possibility to fire crew members with a permanent contract to see if we can further optimise the crew composition and reduce costs. Firing permanent crew members is almost never done at the airline under consideration and is considered an expensive option. Similar to Case 1, the model's solution is again compared to the airline's solution in case the same generated demand is considered.

4.3.1 Model Parameters

For this case, we will use the same settings, starting conditions and block hour data as in Case 1, with the only difference that now the second-stage decision variable r_{tp} does no longer have to be zero for all time steps and crew positions. This also means the cost of firing one permanent FTE has to be defined in the model. This firing cost c_p^{fire} is equal to 102.0 MU for all FO positions, and 130.8 MU for all CP positions. These costs are averages and are calculated using the airline's labour agreement and the average contract length.

4.3.2 Results

In Table 7 the objective value results are found for Case 2. We see that by allowing the model to dismiss crew, the cost has reduced a further 4.5% compared to the model's solution in Case 1 (from 56622 to 54067). On the other hand, the average second-stage cost has increased from 90 to 2926. This is no surprise, since the model now has more flexibility to deal with demand fluctuations in the second stage. In Table 8 it can be seen that the model indeed uses this increased flexibility by firing LH1 and LH2 FTE in the first two months of the simulation period. This confirms the idea that the airline is overstaffed compared to the expected, generated demand, as can be seen for instance in Figure 8b.

Table 7: Objective value results for Case 2.

Case	Objective value [MU]				First-stage cost [MU]	Second-stage cost [MU]
	Min	Max	Average	Δ(min,max) [%]	Average	Average
Model Case 1	56483	56804	56622	0.58	56532	90
Model Case 2	53828	54253	54067	0.78	51131	2926
Airline Case 2	54243	54633	54360	0.72	51253	3107

The result graphs in Figures 10 and 11 show that the model's supply solution can now catch the generated demand way better than in Case 1 for all crew positions (Figures 8 and 9). Especially for both the LH1 and LH2 positions the gap between generated demand and proposed supply has significantly decreased compared to Case 1. Once again, the model's limitation to only hire temporary

Table 8: Average number of FTE fired per month in the model's second-stage solution for Case 2.

Position (p)		Average F_{tp} per month (t) [FTE]						
		Apr	May	Jun	Jul	Aug	Sep	Oct
FO	SH							
	LH1	8.7						
	LH2	3.7						
CP	SH							
	LH1	3.1	2.8					
	LH2	5.2						

FO SH FTE in April is visible in Figure 11a, where it can be seen that temporary FTE are hired at the start of the season, even though they are only needed later on.

It can also be seen that when dismissing crew members is an option, the airline's solution follows demand fluctuations way better (the green plots in Figures 10 and 11). Even though we know that the airline plans to have more crew members than needed in the first stage, the firing option in the second stage now enables a less rigid solution than in Case 1, resulting in more flexibility.

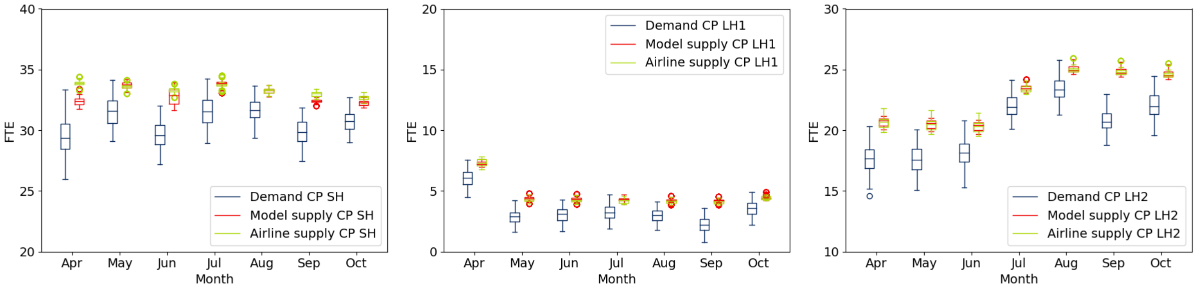

(a) CP SH. (b) CP LH1. (c) CP LH2.

Figure 10: Boxplots of demand generated by the model, its supply solution, and the airline's (second-stage) supply solution for the same demand scenarios, for CP positions for Case 2.

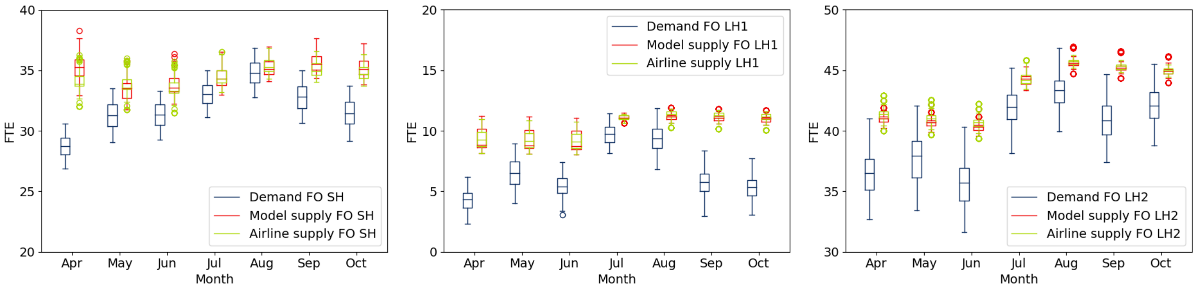

(a) FO SH. (b) FO LH1. (c) FO LH2.

Figure 11: Boxplots of demand generated by the model, its supply solution, and the airline's (second-stage) supply solution for the same demand scenarios, for FO positions for Case 2.

4.4 Case 3: Two Consecutive Seasons

To see how the model reacts to a transition from one season to another, we now analyse a summer and winter season combined. On top of this, the case investigates the model's usefulness for strategic crew planning in the recovery phase after the Covid-19 pandemic. For this case study, the summer season and consecutive winter season of 2022 are analysed (i.e. April 2022 to March 2023).

4.4.1 Model Parameters

For this case the option to fire permanent crew members is again disabled, since it is not a common practice at the airline. This means that again $r_{tp} = 0, \forall t \in T, \forall p \in P$.

New starting conditions and trend values need to be defined. In Table 9 the input starting conditions can be found. These values are provided by the airline and are the *planned* FTE at the end of the winter 2021 season, so at the start of our simulation.

The trend values for this case are provided by EUROCONTROL's Covid-19 recovery forecast and give a demand variation compared to 2019 (EUROCONTROL, 2021). This forecast assumes a flight demand of 84% in 2022 compared to 2019 for the Visiting Friends and Relatives (VFR) market, a part of the tourism market. Hence, block hour data from 2019 is used for the scenario generation and multiplied with 0.84.

Table 9: Starting conditions for Case 3 in terms of permanent crew on the different positions.

Position		x_p^{start} [FTE]	Position		x_p^{start} [FTE]
FO	SH	32.0	CP	SH	32.0
	LH1	12.0		LH1	12.0
	LH2	50.0		LH2	26.0

4.4.2 Results

In Table 10 the model's minimum, maximum and average objective value, and average stage solutions for Case 3 can be found. Once again, the average values are calculated over 20 solution repetitions. Since the number of time steps has gone up from 7 to 12 months compared to Case 1 and 2, the average objective value has increased significantly. In this case, the high second-stage cost stands out. This means that the model makes good use of the second-stage decision variables (transitions and temporary FO SH contracts).

In fact, the model proposes to hire 13.2 temporary FTE on average for the FO SH position. These temporary FTE have to be hired in April, but are mainly used in the summer holiday months when the expected demand is highest, as can be seen in Figure 12a. Hence, the big difference between supply and demand for April to June does not necessarily mean that the starting conditions are too high, but mainly results from the fact that new hires are already hired in April. For the CP SH position (Figure 12b) the model is very well able to follow the expected demand fluctuations. This can be done because at this position transitions are allowed. The change in demand from the summer season to the winter season is also clearly visible in November. This change in demand is managed by transitioning CPs back to the FO position, where in November the temporary FTE hired in April are no longer employed.

However, this flexibility does not exist for all positions. Figure 12c shows the model's solution for the FO LH2 position. Here, no transitions are allowed and the model can only lower the supply by means of natural outflow. It can be seen that the starting conditions (x_p^{start}) are too high for the expected, generated demand for both the summer and the winter season.

Table 10: Objective value results for Case 3.

Case	Objective value [MU]				First-stage cost [MU]	Second-stage cost [MU]
	Min	Max	Average	Δ(min,max) [%]	Average	Average
Model	95905	96248	96077	0.36	92416	3661

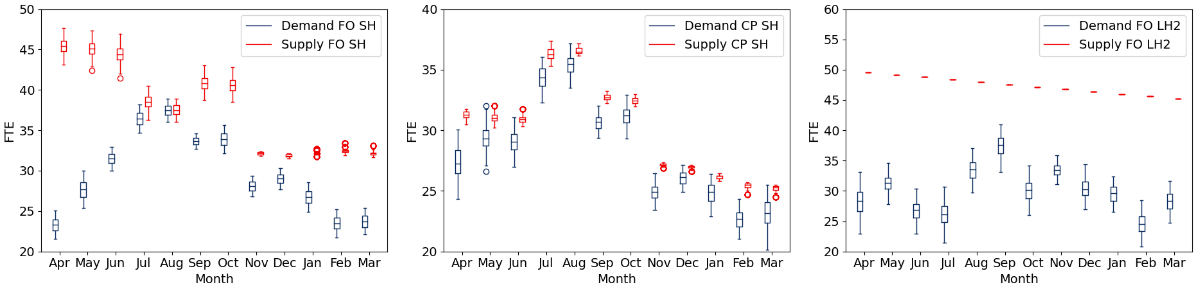

(a) FO SH. (b) CP SH. (c) FO LH2.

Figure 12: Boxplots of generated demand and the supply solutions for some crew positions for Case 3.

5 Conclusions

This paper proposes a new approach for the strategic airline crew planning problem. This problem is faced by airlines before the flight schedule and crew demand are known and thus involves a lot of uncertainty. The methodology uses a two-stage stochastic linear programming model to consider this uncertainty. Crew demand was assumed to be directly related to block hour demand, since the annual number of block hours per crew member is limited. For this reason, the airline's historical block hours were analysed to find the daily distributions, which were assumed to follow a Beta probability distribution. These distributions were then convoluted into Gaussian monthly crew demand distributions. In order to generate the crew demand scenarios that are revealed in the second stage, Latin hypercube sampling was used. This approach allows to take into account the demand correlations between different crew positions. This resulted in correlated crew demand scenarios per crew position and per month of the simulation period. Next to this, it was investigated which other factors besides flights contribute to the total crew demand and by how much. It was found that, among others, absence, holidays and standby shifts cause the airline to need more crew than just for the flights.

Three case studies of real-world situations demonstrated that the adapted approach can support the airline with their strategic crew plan, in both normal years and in the post-Covid-19 recovery phase. The resulting crew composition plan is more flexible than the airline's current practice and is well able to follow crew demand fluctuations. This flexibility is caused by efficiently transitioning crew members between different crew positions and by hiring permanent or temporary crew members in case no transitions are available. It was also shown that the airline's crew supply is too high compared to the estimated future demand. For this reason, the model was allowed to consider firing permanent crew members to further optimise the crew plan. The proposed solution results in cost reductions ranging from 2.1 to 4.5%. These cost reductions mainly come from a lower proposed crew supply. In the final case, it was demonstrated that the developed model is able to solve two consecutive seasons at once and provide a flexible solution. The model dealt well with change in demand from the summer to the winter season. This case also proved that the model can help the airline in the post-Covid-19 recovery phase.

These case studies have shown that the model can already serve as a helpful decision support system for composing an airline's strategic crew plan. However, this paper also enables further research on various topics. For instance, it can be investigated how future crew demand can be modelled more precisely. This can be done by finding the best fitting distribution for each crew position and month instead of assuming the Beta distribution for all cases. As a consequence, the convoluted distribution will not necessarily follow a normal distribution and thus the Latin hypercube sampling will become more complex. It can also be investigated how the trend forecast can consider the specific airline situation. Now the quality of the forecasted demand scenarios depends partially on the quality of the commercial trend forecast. Including airline-specific insights will improve the trend forecast. The model in futures studies can also be extended to account for a better cost calculation. This can be done by looking into the airline's current crew composition with more detail. When crew is considered at an individual level, salary costs, firing costs and transition possibilities will be more exact. Finally, it would be interesting to see how easily the model can be converted to provide a strategic crew plan for cabin crew. Although there are many different rules for cabin crew, an integrated model for both cockpit and cabin crew will help the crew planner even further.

References

Ahuja, R. K., Möhring, R. H., and Zaroliagis, C. D. (2009). *Robust and online large-scale optimization: models and techniques for transportation systems*, volume 5868. Springer.

Altenstedt, F., Thalén, B., and Sjögren, P. (2017). Solving the Airline Manpower Planning Problem. In *Proceedings of the 13th Workshop on Models and Algorithms for Planning and Scheduling Problems.*

Bard, J. F., Morton, D. P., and Wang, Y. M. (2007). Workforce planning at USPS mail processing and distribution centers using stochastic optimization. *Annals of Operations Research*, 155(1):51–78.

Belobaba, P., Odoni, A., and Barnhart, C. (2009). *The Global Airline Industry.* Aerospace Series. Wiley.

Ciriani, T. A., Fasano, G., Gliozzi, S., and Tadei, R., editors (2013). *Operations research in space and air*, volume 79, pages 413–414. Springer Science & Business Media.

Draper, N. R. and Smith, H. (1998). *Applied regression analysis*, volume 326. John Wiley & Sons.

Ehrgott, M. and Ryan, D. (2002). Constructing robust crew schedules with bicriteria optimization. *Journal of Multi-Criteria Decision Analysis*, 11:139–150.

EU (2014). Commission Regulation (EU) No 83/2014, OJ L 28, 31.01.2014, p.17. Official Journal of the European Union .

EUROCONTROL (2021). New EUROCONTROL Four-Year Forecast finds air traffic not expected to reach 2019 levels until 2024 at earliest. `https://www.eurocontrol.int/press-release/new-eurocontrol-four-year-forecast-finds-air-traffic-not-expected-reach-2019-levels`.

Ganguly, S., Lawrence, S., and Prather, M. (2014). Emergency department staff planning to improve patient care and reduce costs. *Decision Sciences*, 45(1):115–145.

Golub, G. H. and Van Loan, C. F. (1996). Matrix Computations, 3rd ed.

Hillier, F. S. and Lieberman, G. J. (2015). *Introduction to operations research, international edition.* McGraw-Hill Education, tenth edition.

Hofman, S. (2020). Strategic Airline Crew Sizing, Literature Study. Master's thesis, Delft University of Technology.

Holm, Å. (2008). Manpower Planning in Airlines: Modeling and Optimization. Master's thesis, Linköpings Universitet.

Holmes, S. (1998). Sums of Random Variables: Statistics 116. Stanford. https://statweb.stanford.edu/ susan/courses/s116/node114.html.

Hooijen, A. A. J. (2019). Cockpit crew transition planning optimisation. Master's thesis, Delft University of Technology.

Horn, M. E. T., Elgindy, T., and Gomez-Iglesias, A. (2016). Strategic workforce planning for the Australian Defence Force. *Journal of the Operational Research Society*, 67(4):664–675.

IATA (2016). IATA Forecasts Passenger Demand to Double Over 20 Years. `https://www.iata.org/en/pressroom/pr/2016-10-18-02/`.

Johnson, N. L., Kotz, S., and Balakrishnan, N. (1995). *Continuous univariate distributions, volume 2*, volume 289. John Wiley & sons.

Laumanns, M. (2011). Robust Planning and Optimization. Course 351-0860-00 Lecture Notes, ETH Zurich.

Li, Y., Chen, J., and Cai, X. (2007). An integrated staff-sizing approach considering feasibility of scheduling decision. *Annals of Operations Research*, 155(1):361–390.

Magini, R., Boniforti, M. A., and Guercio, R. (2019). Generating Scenarios of Cross-Correlated Demands for Modelling Water Distribution Networks. *Water*, 11(3):493.

McKay, M., Beckman, R., and Conover, W. (1979). Comparison of three methods for selecting values of input variables in the analysis of output from a computer code. *Technometrics*, 21(2):239–245.

Montgomery, D. C. and Runger, G. C. (2010). *Applied statistics and probability for engineers*. John Wiley & Sons.

Morén, B. (2012). Utilizing problem specific structures in branch and bound methods for manpower planning. Master's thesis, Linköpings Universitet.

Ng, T. S., Huang, H. C., and Ng, J. Y. (2008). Human resource planning with worker attendance uncertainty. In *2008 IEEE International Conference on Industrial Engineering and Engineering Management*, pages 364–368. IEEE.

Ross, S. M. (2014). *Introduction to probability and statistics for engineers and scientists*. Academic Press, fifth edition.

Safarishahrbijari, A. (2018). Workforce forecasting models: A systematic review. *Journal of Forecasting*, 37(7):739–753.

Saunders, M. (2019). The Primal Simplex Method. Stanford University CME338 Lecture Notes.

Sohoni, M. G., Johnson, E. L., and Bailey, T. G. (2004). Long-Range Reserve Crew Manpower Planning. *Management Science*, 50(6):724–739.

Trivedi, V. M. (1981). A mixed-integer goal programming model for nursing service budgeting. *Operations Research*, 29(5):1019–1034.

Verbeek, P. J. (1991). Decision support systems—an application in strategic manpower planning of airline pilots. *European Journal of Operational Research*, 55(3):368–381.

Watson, J.-P., Woodruff, D. L., and Hart, W. E. (2012). Pysp: modeling and solving stochastic programs in python. *Mathematical Programming Computation*, 4(2):109–149.

Weigel, H. S. and Wilcox, S. P. (1993). The Army's personnel decision support system. *Decision Support Systems*, 9(3):281–306.

Yen, J. W. and Birge, J. R. (2006). A Stochastic Programming Approach to the Airline Crew Scheduling Problem. *Transportation Science*, 40:3–14.

Yu, G., Pachon, J., and Thengvall, B. (2003). Optimization-based integrated manpower management for airlines. In *Operations Research in Space and Air*, pages 407–434. Springer.

Zhu, X. and Sherali, H. D. (2009). Two-stage workforce planning under demand fluctuations and uncertainty. *Journal of the Operational Research Society*, 60(1):94–103.

II

Literature Study

previously graded under AE4020

1 Introduction

Crew costs are the second biggest expense for airlines, hence a small improvement in manpower planning can lead to significant savings (Belobaba et al., 2009). Airline manpower planning can start as soon as five years before the day of operation and lasts until the very day itself. Within this long and complex process, strategic manpower planning focusses on forecasting future crew demand and supply, and on defining strategies to close the gap between them.

Crew demand is mainly determined by the flight schedule, but has to be increased to account for holidays, compulsory rest time, training days, crew absence, etc. Crew supply on the other hand is dictated by the current workforce, but also fluctuates due to e.g. retirements or illness. Once crew demand and supply are determined, the aim is to close the gap between them by means of transitions, hiring new crew, or in the worst case dismissing employees.

Since forecasting crew demand has been called the hardest part of the manpower planning problem (Holm, 2008), most research focusses on crew scheduling in the more urgent short-term, operational planning phase. Research that does focus on the strategic phase, mostly looks at the supply side of cockpit crew and tries to answer the question of which pilots to transition where in order to meet demand.

Nonetheless, forecasting crew demand, preferably with a high goodness of fit, is a crucial step in manpower planning. The closer a crew demand forecast is to the real demand on the day of operation, the less extra costs an airline incurs due to crew shortage or crew surplus. However, there are many factors that complicate crew demand forecasting in the strategic planning phase: the uncertainty is high, the flight schedule might not be available yet, the demand for reserve crew is unknown, etc. Therefore, most research tends to take crew demand as a given input and then investigates the best strategy of closing the gap.

In this report, an overview of the literature regarding manpower planning is given, both within as beyond the airline industry. The focus will mainly be on strategic manpower planning, since this was found to be a little researched area and thus one with still many opportunities. The goal of the report is first to analyse past research: which models and solution methods exist and what is the current state within (strategic) manpower planning? Based on these findings, the second goal is to formulate a research gap and formulate a proposed research question.

This literature study is structured as follows. Part II presents an elaborate overview of the manpower planning problem, both in the airline industry as well as in other industries. Models and approaches will be discussed that can be applicable to the strategic manpower planning problem at an airline. Section 2.3 will then present different forecasting methods which can be used to predict for instance crew demand or crew absence. Next, in Section 3.3 solution methods to the manpower planning problem are given. Again, articles dealing with manpower planning in other industries than the airline industry will be discussed as well. Finally, based on all these previous chapters, a research gap will be defined in Section 4.5. This research gap ultimately leads to a proposed research question and solution approach.

2
The Manpower Planning Problem

The objective of the manpower planning process is to have the right number of people with the right skills at the right time. This is done by making forecasts of supply and demand, taking decision that will close the gap between these two and by making optimal work schedules (Altenstedt et al., 2017).

In this chapter an overview is given of the literature on the manpower planning problem, with a focus on strategic (long-term) planning. Two main sections can be distinguished: Section 2.1 will discuss the manpower planning problem in the airline industry with regard to cockpit and cabin crew, and Section 2.2 will discuss the same problem in other fields, such as hospitals, the army and postal services.

2.1. The Airline Manpower Planning Process

Belobaba et al. (2009) state that labour is the second biggest expense for airlines, right behind fuel expenses. This means that even a small improvement in crew manpower planning can result in large amounts of money being saved, and this makes the manpower planning problem one of the most important ones in the airline industry (Sohoni et al., 2004).

This also means that a lot of research has been conducted on this problem, however, most articles focus on the actual construction and optimisation of crew schedules, and not on the steps that precede this phase: how many employees will be needed, how many employees will retire, how many employees need to be hired, at what position are they needed, etc. This whole process of workforce forecasting, planning and scheduling will be called the *airline manpower planning problem* throughout this report. As with the overall airline planning process, this problem can be divided into three chronological phases: a strategic, a tactical and an operational phase.

These three phases will be explained in detail in Section 2.1.3, 2.1.4 and 2.1.5 respectively. First some context about the overall airline planning process is given in Section 2.1.1 after which the goal and obstacles of the airline manpower planning problem are described in more detail in Section 2.1.2. Solution methods will be discussed separately in Chapter 3.3, and therefore the focus in this section will be on model assumptions and formulations.

2.1.1. The Overall Airline Planning Process

The overall objective of an airline's planning process is to maximise its profitability by making the right decisions at the right time. This objective can be reached by making strategic decisions (e.g. which aircraft types to buy), tactical decisions (e.g. prices of the tickets) and operational decisions (e.g. how many crew members to schedule on stand-by) (Santos, 2018a).

This full planning process is long and complex: it spans a period of approximately ten years, involves a lot of uncertainty at the beginning and is governed by an ample amount of factors. The economic situation, labour agreements, slots, curfews, available aircraft types, maintenance, legal requirements, etc. make the airline planning problem impossible to solve, or even formulate, in one go (Barnhart et al., 2003). Therefore, the whole problem is divided into multiple sub-problems that are then solved subsequentially. The output of one problem often is the input for the next problem (Clarke and Smith, 2004).

These different sub-problems are defined differently in literature, although not drastically. Belobaba et al. (2009) divide the overall planning process in three different categories: fleet planning, route planning and schedule development. Barnhart et al. (2003) further divide the schedule development phase into four categories: schedule design, fleet assignment, aircraft maintenance routing and crew scheduling. Finally, Clarke and Smith (2004) add airport resource management and revenue management to these four steps.

Even though there is no clear consensus about the definition and order of the different steps, it is easy to see a certain pattern in the numerous sub-problems. First a network and fleet plan are established. Then, the scheduling phase starts: a flight timetable is constructed, aircraft are assigned, major maintenance is planned, and finally a crew schedule is made.

Manpower planning does not come into play only at the crew scheduling phase, but is a part of the entire airline planning process, as will be explained in the next sections.

2.1.2. Goal and Obstacles of the Airline Manpower Planning Problem

The basic goal of the manpower planning process given at the beginning of this chapter, can be adapted to apply for an airline: to have the right number of people with the right skills at the right time *at the lowest cost*. This makes the problem consist of four elements: people, jobs, time and money (Grinold and Marshall, 1977). As simple as this may sound, there are some obstacles that complicate this objective:

- Uncertainty: uncertainty is the main difficulty in manpower planning since it generates inaccuracies in the forecasts. Crew's ambitions, future destinations, the economic situation, future supply and demand, crew absence, operational disruptions, etc. are all uncertain at some point or another and this needs to be taken into account in the manpower planning process. Not doing so may result in a shortage (or excess) on the day of operation, both resulting in unnecessary costs. The closer to the day of operation, the lower the uncertainty (Santos, 2018a).
- Complexity: airline crew is very heterogeneous when it comes to qualifications and employability. Most airlines have several aircraft types, and not all crew members have the same qualifications (*type ratings*). For instance, a crew member can be allowed to fly as a captain on the Boeing 737, and as a First Officer on the Boeing 787. In addition, large airlines usually have multiple hub airports. This makes that manpower planning is mostly done per aircraft type, per position and per hub.
- Regulations and labour agreements: an airline, just like anyone else, has to obey the law. For instance crew is entitled to rest and days off between their flights. On top of this, airlines mostly have different agreements with cockpit and cabin crew, and this further complicates the problem. Crew also requires training, which can be initial, transition or recurrent. Not taking this into account during the manpower planning may result in uncertified, and thus not employable, crew members.

As stated before, manpower planning is a long and complex process and therefore it is split up in different stages. One can easily understand that it would be impossible and absurd to make crew work schedules five years in advance. On the other hand, recruiting new crew needs to be done well on time. The different stages and respective research are explained in the next sections.

2.1.3. Strategic Manpower Planning

Strategic planning is the first phase in the manpower planning problem and its main goal is to forecast future crew supply and demand and to determine a strategy to close the gap between them. The process takes place from three or even five years until approximately one year before the day of operation (Holm, 2008; Hooijen, 2018). In practice however, most airlines do not look much more than one to one and a half year ahead when it comes to manpower planning.

Very little research has been done in strategic manpower planning for airlines. Safarishahrbijari (2018) analysed 275 papers on workforce forecasting models and concluded that since 1980, only two percent of research in this field is done in the marine and airlines industry, while over 50 percent of the research has been performed in the health care industry. Figure 2.1 shows the proportion of articles on this subject in other fields.

Already in 1991, Verbeek developed a strategic decision support system (DSS) for cockpit crew at KLM (Verbeek, 1991). The model had a planning horizon of ten years and was used to evaluate different

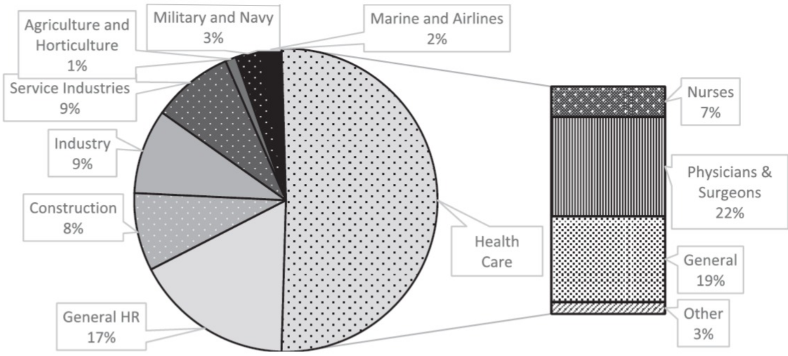

Figure 2.1: Percentage of articles on manpower forecasting models since 1980 related to different areas (Safarishahrbijari, 2018). The total number of articles is 275.

manpower scenarios, hence to help the planner in making faster and more efficient decisions. Among other things, a long-term fleet plan and forecasted demand were used as input. The article did not specify how this demand was determined. Flexible demand was used to account for vacations. In order to deal with the uncertainty in the long-term, Verbeek simulated pilot's behaviour, but again no details were given about this simulation. Long-term supply was simply determined by taking into account the workforce and expected retirements.

The model then determined the cockpit crew supply and demand, and the gap between them, for the entire planning horizon, per month and per position. The model was also capable of presenting a plan to solve possible imbalances between supply and demand at certain positions. This was done by proposing crew transitions, a method used by most airlines to close the gap between crew supply and demand, often called transition planning. Transition planning aims to solve the question of when to *transition* crew (mostly cockpit crew), and to which function. To illustrate this concept, a typical, hierarchical career path is shown in Figure 2.2, where the arrows indicate transitions.

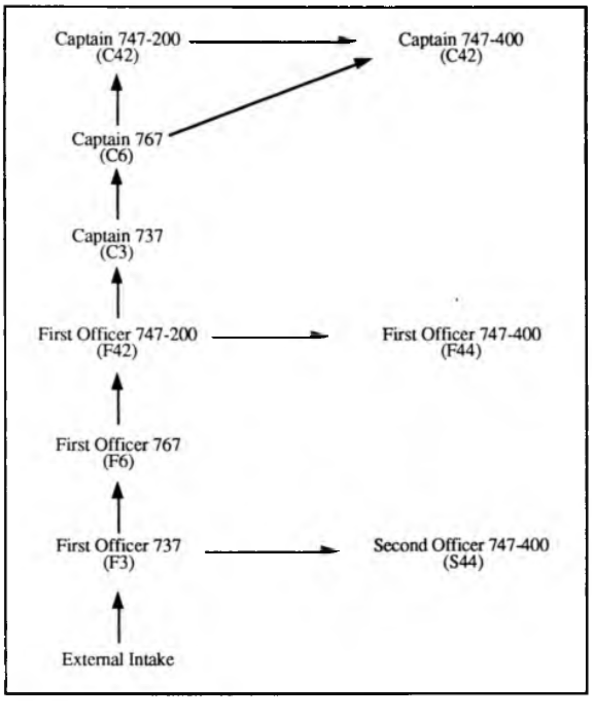

Figure 2.2: A typical, hierarchical career path for an airline pilot, with the arrows indicating a transition (Ives, 1992).

In order to come up with this transition plan, Verbeek's model made use of mathematical formulations in the form of a mixed-integer (MIP) model with the objective of minimising costs. Since the problem at hand was large, with 19 cockpit positions and over 900 pilots, it could not be solved with commercial software. Instead, heuristics were used to reach a near-optimal solution, but no further details were given.

Although Verbeek's DSS forms a good starting point to understand the working principles and factors behind strategic manpower planning, the article leaves out important information about the simulation, forecasting and heuristic solution methods. No other articles were found that investigate the strategic manpower planning in airlines with such a long planning horizon as in the article of Verbeek (1991).

In 2008, Holm tackled the same problem and developed mixed-integer linear programming (MILP) formulations for both the transition problem as for allocating training and vacation. Her model had a planning horizon of 70 weeks. Holm (2008) modelled the pilot transitioning problem assuming that a pilot's career path is not fixed and investigated how pilot's choices could be influenced by changing the salary distribution over different positions. This way the optimal, cheapest career ladders could be found. For instance, by placing career ladders in parallel, transition and training costs can be reduced drastically. This principle is illustrated in Figure 2.3.

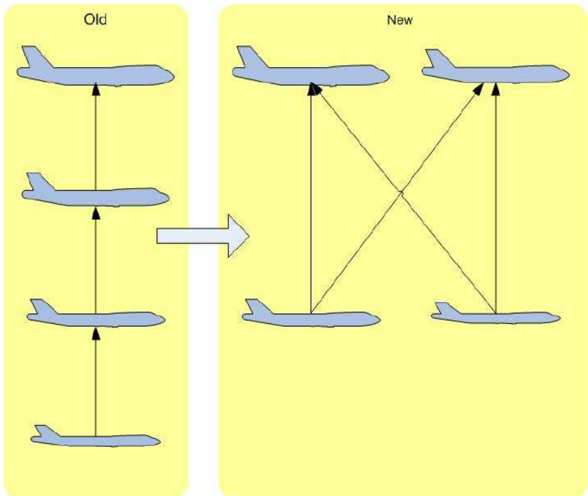

Figure 2.3: An example of how changing a career ladder structure influences the number of transitions needed for a certain aircraft type, and thus changes the costs associated with it (Holm, 2008).

The objective for both models was to minimise costs. For the transition model this was the sum of employee costs and transition costs. Constraints were set up that made sure demand was satisfied and the transition rules were followed. The model's output was the pilot supply for each position during a certain period, and the number of pilots trained from one position to another during that period.

For the training and vacation model, the costs were the sum of transition costs, course costs, shortage costs and pay-protection costs (a cost paid to a senior pilot if he did not get a transition but a less senior pilot did). Constraints put a limit on the number of course participants, made sure all labour agreements were complied with and tracked the shortage of block hours per position. The output here was the number of pilots assigned to a course, the number of pilots with vacation, the number of pilots under pay-protection and the block hour shortage per position, all calculated per period (e.g. one month).

Morén (2012) also developed a MILP formulation for the staffing and transitioning of pilots. The goal was to minimise costs while satisfying demand. The formulation and model output resembled these in Holm (2008), but Morén (2012) also allowed for pilots to fly below rank in case of a supply shortage. In this case, a captain can be used as a first officer or on a smaller aircraft type during a planning period, although this is seen as an expensive form of manpower. The model could also take into account pilots certified to fly on multiple aircraft types. The other models mentioned here before did not take this into account.

Both Holm (2008) and Morén (2012) used branch and bound as a solution method for their MILP formulation. This method, together with other solution methods for the manpower planning problem, will be discussed in Chapter 3.3.

More recently, Hooijen (2019) developed a DSS to help KLM with the strategic cockpit crew transition planning problem up to two years in advance, although the author states that cockpit crew transition planning is usually done three to five years in advance.

Hooijen (2019) developed both a heuristic planning model and an optimisation model. The goal of the heuristic planning model was to determine an optimal crew plan by planning such transitions. It uses supply and demand data, transition rules and some other parameters such as training capacity and transition characteristics as input.

The model consists of a local search algorithm that evaluates possible solutions in the neighbourhood of the current solution and iterates until a stopping criterion is met. The size of the problem was decreased by setting up a rule-based system to omit options that do not comply with certain rules and a tabu search method to avoid getting stuck in a local optimum. The best transition option in the neighbourhood was then picked by a selection algorithm using a tree search, combined with either a naive selection algorithm, a greedy algorithm or Dijkstra's shortest path algorithm to decrease search space and computation time. On the other hand, the optimisation model's goal was to minimise the cost of the gap between pilot supply and demand, be it negative or positive, for each position. The cost of this gap varied per position (different salaries) and increased if a crew shortage persisted for a longer period, since this has a negative effect on the operations.

As stated before, the uncertainty about the future during the strategic planning phase is very high and that makes designing an accurate strategic manpower plan quite hard. Airlines rarely have a flight schedule ready more than one and a half year in advance and thus it is hard to predict crew demand in the strategic phase. Morén (2012) and Verbeek (1991) used crew demand as a parameter in their models, but did not mention how that number was determined. Hooijen (2019) determined crew demand based on the flight schedule and then increased it to account for holidays, absence, training, etc. Finally, Holm (2008) determined crew demand based on expected crew utilisation, although she also stated that having a preliminary timetable would result in better estimates. No research was found that investigates the effect of different demand calculations or forecasting methods on the strategic manpower plan.

Next to determining crew demand, also determining crew supply is hard in this phase, since the crew's ambition is mostly unknown. In some airlines, pilots and cabin crew will have the opportunity once or twice a year to bid for a promotion (transition). For instance, a First Officer on the Boeing 737 may indicate that he wants to become a Captain on that aircraft type. If his bid is accepted, he will be trained and tested for that position. Since these bids only have to be made known approximately one year in advance, it is mostly not possible to include them with certainty in the strategic manpower planning phase (Hooijen, 2018). On top of this, crew absence (e.g. because of illness, holidays, etc.) is unknown at this point. In conclusion, when designing a long-term manpower plan, it is important to take the uncertainties in supply and demand into consideration (Holm, 2008).

One of the only aspects that can be predicted fairly accurately in the strategic planning phase, is the retirement rate in the oncoming years. Retirements are a determining factor in crew supply and are rather easy to predict, since pilots have a mandatory retirement age of 65 (EASA, 2019). Of course, retirement age can vary between companies, but in general an airline has a good overview of retirements in the coming years, assuming that elder pilots will stay at the company (Holm, 2008). Hooijen (2019) used the composition of the current workforce and forecasted retirements to calculate crew supply in the future and used that as an initial input to his model. Airlines usually also make a forecast of future crew absence based on historical data. The forecast is then used to estimate future crew supply.

The high level of uncertainty in crew supply and demand however, does not render the strategic manpower planning problem unimportant nor impossible. For an airline, it is still desirable to have an idea of manpower levels in the coming years. Not only will this determine the strategies to close the gap between supply and demand, this way the airline can also make sure to have sufficient training capacity, and it knows what is needed, both financially and logistically, to continue operations as planned.

2.1.4. Tactical Manpower Planning

Tactical planning comes after the strategic planning phase and takes place from approximately one year until six months before the day of operation. The main goal during this phase is to close the gap between forecasted supply and demand (Santos, 2018a). During this phase the uncertainty about the future is lower, and thus more concrete plans can be made about crew's vacation and training days.

Tactical manpower planning has been researched more than strategic manpower planning, especially transition planning in the tactical planning phase. As stated already in Section 2.1.3, in some airlines crew has the opportunity to bid for a promotion once or twice a year. This means that during the tactical phase, crew's ambitions are then known and this makes predicting crew supply easier.

Most research, however, considers the transition of a pilot as a given, and thus does not question if or whereto the pilot or crew member should optimally transition (Hooijen, 2018). Thalén (2010) integrated both the staffing problem (i.e. whether to hire pilots and which ones to promote) and the transition problem by designing a tabu search based algorithm that was 30 times faster than solving it with mixed-integer programming (MIP) software. The objective was to supply the right quantity of pilots with the right qualifications at the right time and at minimum cost. Thalén's model sometimes omitted the airline's seniority rules in order to find a more optimal transition solution. This was done by either using pay-protection (similar to Holm (2008)) or by prematurely moving pilots with a higher seniority to different positions.

Over several years, a tactical manpower planning DSS was developed for Continental Airlines by Gang Yu (Yu et al. (1998), Yu et al. (2003) and Yu et al. (2004)). In 1998 they first developed a mathematical framework to solve the pilot training assignment problem. Here, they only looked to optimise the training schedule, i.e. assigning pilots to training in time so that they can cover their assigned block hours in the future. Different possible objective functions were formulated: (1) maximise compliance with the business plan, (2) minimise pilot unavailability by minimising training cycles, and (3) minimise the total cost. They then designed a heuristic method to solve the problem. It is initialised by a manual solution and then looks at the entire planning horizon (discretised into time periods of one week) to make possible schedules for pilots that need training during that planning horizon. The main working principle of the heuristic is that, for each week, it assigns pilots to training if certain conditions are met, and then updates the list of unassigned pilots to prepare the heuristic for the next iteration (Yu et al., 1998). The model was run for a planning horizon of 52 weeks and resulted in a list of course dates, pilot course assignments and capability in terms of block hours. The model also had the option to determine how many new hires would be needed to satisfy the business plan.

In Yu et al. (2003) and Yu et al. (2004) a whole new model was presented that aimed to integrate various aspects of the manpower planning problem, such as demand forecasting, transition planning and absence management, for both cockpit and cabin crew, although they only present a mathematical formulation for the pilot transitioning and training problem. This time, the solution method used was optimisation-based and albeit stating that the goal of the paper is presenting an integrated manpower management DSS, the transitioning and training problem were solved separately.

2.1.5. Operational Manpower Planning

Operational manpower planning is the last step in the process and takes place from approximately six months before the day of operation until the very day itself. The objective during this phase is mostly to minimise costs (Santos, 2018a). Operational manpower planning is often called crew scheduling and is the most researched phase of the manpower planning problem. In most literature, the crew scheduling problem is split into two stages: the crew pairing problem and the crew rostering problem (Belobaba et al., 2009; Medard and Sawhney, 2007; Barnhart et al., 2003). In case of disruptions, also crew recovery needs to take place. Although this can be seen as a step that comes after operational planning, it will also be discussed in this section.

The crew pairing problem generates work schedules of one to five days. The goal is to construct crew itineraries (or *pairings*) so that all flights are covered at minimum cost. It starts by constructing a sequence of flights, which is called a duty period, and is generally made for one day. Then, by putting multiple duty periods after each other, a pairing is formed. On the other hand, the crew rostering problem combines multiple pairings, training days, standby duties and holidays into a crew roster of approximately 30 days that is then assigned to individual crew members (Santos, 2018b).

Of course, during these two steps, regulations and labour agreements have to be considered: there has to be enough rest time between flights, there is a maximum of work hours per day, the pairing has to start and end at the crew base, etc. Planners also have to make sure that there is legally sufficient cockpit and cabin crew assigned to each flight (Barnhart et al., 2003; Medard and Sawhney, 2007).

Over the last decades however, research on this topic has shifted from focusing solely on cost minimisation to focusing on both costs and schedule robustness. In a more robust schedule, disruptions (such as delays and cancellations) in the operations will propagate less into the future. This way a more robust schedule might result in higher *planned* costs but lower *actual* costs than the cost-minimised schedule (Ehrgott and Ryan, 2002).

For this reason, Ehrgott and Ryan (2002) developed a bi-criteria optimisation model to develop a crew schedule for a domestic airline. Two objectives were needed since maximising robustness conflicts with minimising the costs. An objective function was developed that penalised duty periods that were not robust, i.e. if the expected delay of the incoming flight was larger than the crew's scheduled ground time minus possible ground duty time. The expected delay was taken from historical data of over 46,000 flights and showed a clear linear increase throughout the day. The schedule robustness was then measured by calculating the cumulative expected delay. Ehrgott and Ryan (2002) concluded that a significant increase in schedule robustness is possible with only a small increase in planned operating costs.

Yen and Birge (2006) on the other hand developed a two-stage stochastic integer formulation and used a random variable to introduce random disruptions in the schedule to account for short-term changes in the long-term crew planning, but did not consider labour agreements, standby crews and flight cancellations. This was because the article aimed to identify the delays that were caused because of the crew schedule, and to use this info to make the crew schedule more robust. They concluded that serious cost savings could be made when considering disruption effects on the operations already during the planning phase.

In case of schedule disruptions, due to delays, absent crew, severe weather, etc. airlines want to resume normal operations as soon, but also as cheap as possible. This can be done by swapping crew or aircraft to new flights and by the use of reserve crew. Over the last years, research on the use of reserve crew has increased (Sohoni et al., 2006; Homaie-Shandizi et al., 2016; Bayliss et al., 2012, 2019). Determining how many reserve crew are optimally needed requires a model that can deal with future uncertainty and in one way or another predicts the probability of crew absence.

2.2. Manpower Planning in Non-Aviation Industries

In this section manpower planning models from other industries will be discussed. Section 2.2.1 will focus on hospital manpower planning, while Section 2.2.2 focuses on manpower planning in armies. In Section 2.2.3 some articles are discussed that could not be classified into one of these industries.

In the grey text boxes throughout this section, it will be discussed how the model(s) can be adapted to address the the airline manpower planning problem. Solution methods will be discussed separately in Chapter 3.3, and therefore the focus in this section will be on model assumptions and formulations.

2.2.1. Hospital Manpower Planning

As stated in Section 2.1.3 and Figure 2.1, Safarishahrbijari (2018) found that since 1980, 52 percent of articles on strategic manpower planning are related to the health care sector. Fortunately, the problem is in many ways similar to the airline manpower planning problem. Nurses and physicians have different skill levels, just like cabin and cockpit crew. Also uncertainty in staff demand plays a role in both industries. Moreover, hospital manpower planners also need to adhere to the law and agreements with unions when it comes to holidays, overtime and training (Trivedi, 1981).

Already in 1981, Trivedi made a mixed-integer goal programming model to determine the number of nurses needed in a hospital on each position for each shift. Goal programming uses target values for each objective and deviations from these values are minimised. In the model, staff is divided into different skill levels, there are constraints on time off and holidays, the amount of substitution among different skill levels is limited, etc. The model also takes the law and labour agreements into account.

Furthermore, Trivedi (1981) allows for part-time staff and implements a constraint that makes sure enough replacement staff is available to cover for absent staff.

Demand was derived from historical data to form a projected number of expected patients divided into four categories, for weekdays and weekend days, that was then converted into nursing hours. In the numerical example, this projected demand was assumed to be the same throughout the entire year and thus the model could be run for an entire year. The output then was the number of nurses needed per skill level, per shift (day, evening, night) per week day or per weekend day, for the entire year. Since the demand was the same throughout the year, the number of nurses needed each week for each position and shift was the same.

Five goals were defined: minimise the budget deficit (if the expenses are larger than the budget), the understaffing and the number of part-time employees, and maximise the budget surplus (if the expenses are smaller than the budget) and the number of full-time nurses.

Application to the airline manpower planning problem:
The constraints used by Trivedi (1981) to formulate holiday policies and limits on working hours could be used in the airline manpower planning problem. Also the implementation of part-time employees is a good addition, since not all airline crew members work full-time. A shortcoming however is that demand is assumed to be the same throughout the entire year and thus the output provides the same numbers throughout the year. It is uncertain how the model would deal with fluctuating staff demand per season, per week and even per day as is the case for airlines.

Li et al. (2007) developed an integrated staff-sizing approach that shows a lot of similarities with the airline manpower planning problem. The model they developed is meant for any service organisation although they use nurse manpower planning for their numerical example, and shows many similarities with the model of Trivedi (1981). Their planning horizon consisted of six planning periods of four weeks (so half a year). They described three stages in an integrated staffing decision model: forecasting, planning and scheduling, which they call stage I, stage II and stage III respectively. In Figure 2.4 a schematic overview of their integrated staffing decision model can be found.

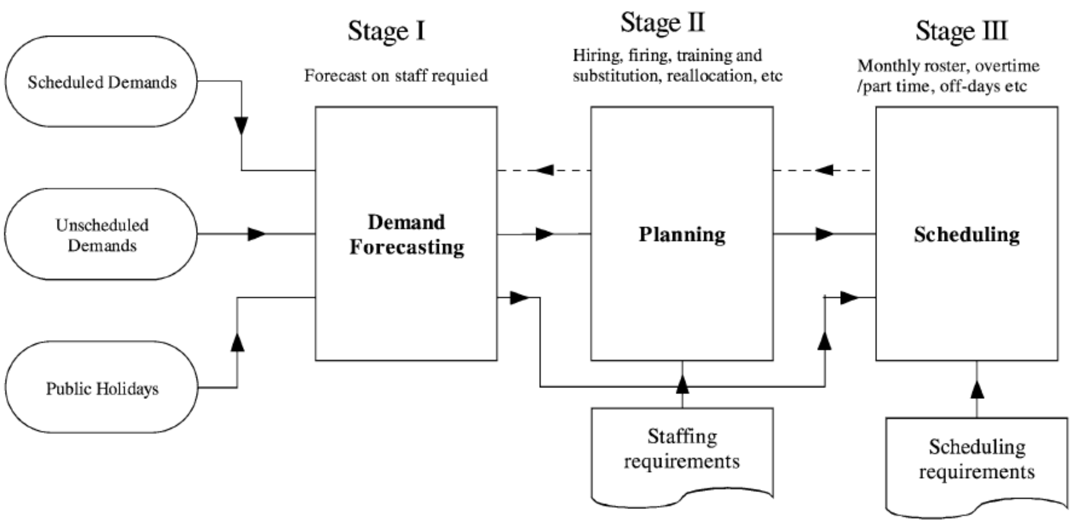

Figure 2.4: Overview of the integrated staffing decision model presented by Li et al. (2007).

In stage I the demand is determined. They distinguish between known demand (patients that have made an appointment) and unknown demand (patients that come in without an appointment). They do not specify further how this demand is determined, and simply use the sum of known and unknown demand as one demand parameter later in stage II and III. The output of stage I then is a monthly demand prediction and is the input for stage II.

In stage II these forecasts are used together with the staffing requirements to determine the desirable workforce size for each position. These requirements specify the maximum labour hours, overtime policies, required training, etc. The model also determines how many people to hire and dismiss, how many overtime hours will be performed and how many people will have to work on a lower-level position. This was done by developing a multi-objective staff planning model with linear objectives and six linear constraints to limit the overtime, training, etc.

Five separate objectives were identified to guarantee a certain service quality without neglecting cost issues, all minimisation objectives.An analytic hierarchy process (AHP) was used to convert the multi-objective model into a single objective model. This method is normally used for multi-criteria analysis but is here used in a multi-objective context. It calculates weights for the different objectives by determining their relative importance with respect to each other measured on an integer scale from 1 to 9. These values are then placed in a square matrix which is then normalised, resulting in the objective weights (Li et al., 2007; Winston and Goldberg, 2004).

Finally, in stage III the outputs from stage II are used together with scheduling requirements to generate a schedule. They stress that the success of the integrated system depends on an effective coordination between all stages, however they admit that the demand input is fixed and thus there is no interaction between stage I and the other two stages. Here again, a multi-objective model was developed.

Application to the airline manpower planning problem:

- Li et al. (2007) distinguish between known and unknown demand. Known demand could be derived from the flight schedule and unknown demand could come from staff absence, disruptions, etc.
- The staff is divided into different skill classes. This is also the case in an airline: captain and first officer in the cockpit; purser, assistant purser and cabin attendant in the cabin.
- The staffing and scheduling requirements used in stage II and III respectively, that specify the maximum amount of labour hours, training policies, holiday policies, etc. are with some adjustments applicable to an airline.

In conclusion, the models of Li et al. (2007) show many similarities with the strategic airline manpower planning problem. Especially the staff planning model seems like it could be applied to an airline with some adjustments. Their planning horizon would however need to be adapted in order to make it a strategic planning model.

More recently, Ganguly et al. (2014) developed a staff planning model for a medical emergency department. They used a mixed-integer linear programming (MILP) formulation to investigate the optimal number of staff. The objective was to minimise total staffing cost while still reaching a service level target. In contrast to Li et al. (2007), demand was fully unknown, due to the nature of a hospital's emergency department.

To cope with this, they aggregated historical patient demand into discrete time buckets of one hour and used this to model the stochastic distribution of demand within these buckets. Patients were divided into different acuity levels and care providers were divided into different skill levels (cf. Li et al. (2007); Trivedi (1981)). This way, they were able to formulate the distribution and density of work content, i.e. the relative amount of time a patient with a certain acuity level requires attention from a qualified provider.

Another interesting aspect is that Ganguly et al. (2014) made use of chance constraints as developed by Charnes and Cooper (1959) to guarantee that a certain service level is met. A chance constraint makes sure that the probability that a constraint will hold is higher than or equal to a set target p, in this case the target service level. The basic principle of a chance constraint can be formulated as follows:

If the inequality constraint is $h(x, \xi) \geq 0$, then the chance constraint is $P(h(x, \xi) \geq 0) \geq p$

with $p \in [0, 1]$, x the decision vector and ξ the vector of uncertainty (Li, 2015). A chance constraint can be rewritten as its deterministic equivalent if the probability distribution of the random variable (ξ) in the chance constraint is known (Hillier and Lieberman, 2015). Since the work content distribution was already developed (as explained above), it was possible to rewrite the chance constraints as their deterministic equivalent (Ganguly et al., 2014).

Finally, they also allowed the model to form teams of care providers. In this case, lower skilled providers are assisted by higher skilled staff to treat patients that could normally not be treated by the lower skilled provider. This further increased the flexibility and utilisation and decreased the total staffing cost.

In the numerical example, the planning horizon consisted of 24 planning periods of one hour (so one day). However, they found that a planning horizon consisting of 164 planning periods of one hour (so one week) resulted in two percent lower staffing costs.

Application to the airline manpower planning problem:
Ganguly et al. (2014) had to deal with a manpower planning problem in which all demand was unknown. Mostly, this is not the case for the airline manpower planning problem, since crew demand can be partially determined by either the flight schedule or future fleet size. Unknown crew demand in airlines mostly comes from absence and disruptions, and therefore this method could perhaps be used to model crew absence stochastically.

The use of chance constraints could be effective when planning for reserve crew. For instance: the probability that reserve crew can cover demand should be equal to or higher than p.

The planning horizon of one day, or even one week is not enough to make a long-term, strategic manpower plan. However, by for instance changing the time buckets from one hour to one week or one month, a longer planning horizon could be achieved with the same model.

2.2.2. Army Manpower Planning

Similar to hospital manpower planning, army manpower planning deals with getting the right people with the right skill at the right time under partially uncertain demand. Since one of the main characteristics of an army is that it has a hierarchical structure, manpower planning models often make use of promotions to transition staff. This is also the case in the airline manpower planning problem, as was explained in Section 2.1.3.

Weigel and Wilcox (1993) designed a personnel decision support system (DSS) for the United States Army. The goal was to model workforce supply by combining techniques of goal programming, network models, linear programming and Markov-type inventory projection. In total, three models were made and a hierarchy was used so that the bigger problem was represented as ”a sequence of linked models” (Weigel and Wilcox, 1993), where each model provides constraints for the next model. The reason to use such a hierarchical approach was that it was impossible to solve a problem of such size at once considering computer technology at that time. The hierarchy also allowed to focus on a different set of dimensions for each model: the top model performs the global optimisation, the lower models perform local optimisation.

Horn et al. (2016) designed a mixed-integer linear programming (MILP) model to support the Australian army in planning their workforce on a strategic level. A goal programming approach was used to formulate different objectives.

They started by designing a basic model that was extended later to reflect real-world conditions. In the basic model, army staff is divided into cohorts, trades and ranks, all characterised by a hierarchical structure. The objective of the basic model was to minimise total weighted deviations from the target personnel levels. These target levels were defined per trade, per rank and per year. The main output was the number of personnel in each cohort, trade and rank, the number of entries and departures, and promotion characteristics.

In the extended model more constraints were added to reflect the real-world situation. Real costs of hiring, attrition and promotions were added to the original objective function. The model's planning horizon consisted of 17 planning periods of one year (so 17 years). This means that proposed changes occurred at yearly intervals.

Škraba et al. (2016) used the principle of system dynamics to model a hierarchical human resources structure in organisations such as the army. A combination of genetic algorithms and stochastic local search was then used as an optimisation tool. They too made use of recruitments and promotions to attain the desired staffing levels for each rank, while also taking into account wastage.

They developed an iterative stochastic search algorithm that chose randomly between three actions in each of the possible defined situations. This was done in order to alleviate the complexity of dependencies between the different nodes. The algorithm looped over all ranks (which they called *classes*) and checked for each class if the staffing level was higher or lower than desired. If it was lower, the algorithm chose randomly between (1) increasing the recruitment, (2) decreasing the outflow or (3) decreasing the promotions. If it was higher, the algorithm chose randomly between the opposite of these three options.

Škraba et al. (2016) stated that the algorithm had to be launched repeatedly until two resulting iterations were equal. They further stated that the outcome of the algorithm depended on the initial values. Nonetheless, feasible solutions could be found. A genetic algorithm was further used as an optimisation technique to generate optimal values that were then used as initial values for the stochastic search algorithm.

Application to the airline manpower planning problem:
The discussed models all use recruitment and promotions to reach target levels in the hierarchical army structure. This is similar to what Holm (2008) and Hooijen (2019) did for the strategic airline manpower planning problem.

Weigel and Wilcox (1993) have shown that it is possible to combine different models in a hierarchical structure to solve a larger problem (which the airline manpower problem definitely is). Horn et al. (2016) demonstrated that using different penalty weights for different positions is a way to manually stress the importance of each position. Hooijen (2019) made use of this technique to account for problematic cockpit crew positions. Finally, it is interesting to see that Škraba et al. (2016) used an algorithm in which different options were chosen randomly at each step, and by combining it with a genetic algorithm optimal solutions could still be found.

2.2.3. Manpower Planning in Other Industries

This section will discuss articles and methods that could not be categorised into one of the industries mentioned above.

Cai et al. (2004) dealt with uncertain manpower demand by developing a multi-period stochastic decision model for organisations dealing with fluctuating demand. Although they did not specify their planning horizon, the developed model was meant for the long-term manpower problem, and it is up to the user to define the length of the planning periods. Similar to other articles mentioned before, they defined different skill levels so that high-skilled employees could be assigned to perform lower-skilled jobs.

They first formulated a stochastic manpower problem (SMP), where demand in the current period is zero and demand in the future is described by a random component. This means that only for the current period, demand was determined, and for the next periods they modelled demands as random variables with a normal distribution. Then, they reformulated the SMP into a deterministic manpower problem (DMP) by using a mean value model approach to reduce the complexity and to make it easier to solve. They further added a feedback mechanism to make sure the staff size was not overestimated in the far future.

Application to the airline manpower planning problem:
Cai et al. (2004) chose a stochastic integer programming approach to deal with uncertainty in manpower demand and then reformulated it into a deterministic model to reduce complexity and solution times. Manpower demand in airlines also has stochastic characteristics (e.g. because of delays and absence) but since the problem is already quite large and complex, it might be necessary to avoid stochastic formulations and shape them into deterministic ones.

Qi and Bard (2006) used a simulation model that integrated two staffing decision support systems for a large mail company. The first system, called *ESO*, was used to make optimal use of the automation equipment. The second system, called *SOS*, was used to determine the optimal number of employees while taking into account all labour requirements such as breaks and holidays. The latter system first solves a large-scale integer programme that determines the optimal number of shifts in a week. Post-processors then add lunch breaks and glue these shifts together into weekly schedules. It also takes into account the difference between part-time and full-time workers.

As can be seen in Figure 2.5, simulation was used twice to integrate both decision support systems. To do this, the mail company was modelled as a time-space network; the mail handlers were modelled using distance matrices and rules were imposed to govern their movement.

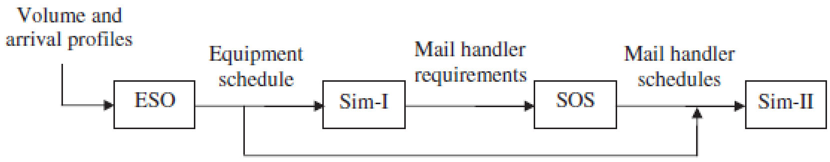

Figure 2.5: Overview of the integration of decision support systems and simulation as formulated by Qi and Bard (2006).

Application to the airline manpower planning problem:
Qi and Bard (2006) made use of simulation to integrate two models. The method of simulation was already used by Verbeek (1991) to model pilot's behaviour and can thus be used in an airline manpower planning context. However, in the article of Qi and Bard (2006) simulation is used in a very different way: rather than simulating people's behaviour, it is here used to model the flow and operations in a mail facility in order to come up with handler requirements, and in a later stage to validate the results. This seems less useful for the strategic airline manpower planning problem, since here the focus is on determining crew demand and supply, and not yet on the operational (time-space) phase.

Bard (2004) developed a mixed-integer linear programming model to find the optimal number of employees in a service organisation while taking the law and labour agreements into account. However, the focus in his article was not on promoting staff to higher functions, but on temporarily downgrading staff to satisfy demand in lower skill categories, similar to what Li et al. (2007) did for nurses. Bard (2004) assumed not only full-time and part-time workers, but also temporary employees to deal with exceptionally high demand (together with overtime). The article further stated that assuming more than two skill levels will likely make the model unmanageable.

Staff demand varied throughout the day and throughout the week and was defined per skill level. Although they mentioned that planners often have to deal with unknown spikes in demand, they chose to calculate average demand using historical data and to use that as an input to the model. The model's objective was to minimise the staffing cost and had a planning horizon of one week.

Since the problem took too long to solve using CPLEX, a sequential approach was chosen as solution method, for which they made use of greedy heuristic algorithms and post-processors.

In later similar research, Bard et al. (2007) took into account fluctuations in demand by developing a stochastic integer programming model, consisting of two stages. The first stage determined the optimal number of full-time and part-time permanent workers under unknown demand. In the second stage, the demand was known and shifts were formed. As in Bard (2004), temporary workers and overtime were used if the permanent workforce was too small to cover the demand. Historical data was analysed to find the demand distribution.

Application to the airline manpower planning problem:
The model described in Bard (2004) took too long to solve in CPLEX and therefore they used a sequential approach and post-processors. They further warned that using more than two skill levels might render the model unmanageable. Since the airline manpower problem is large and complex, the combination of different solution methods might be needed in order to still solve the problem.

In their later research the stochastic nature of crew demand is introduced, similar to Cai et al. (2004). Bard et al. (2007) uses a two-stage stochastic integer programming model, and this seems perfectly applicable to the strategic airline manpower planning problem, especially since the uncertainty in this phase is rather high.

Finally, De la Torre et al. (2016) developed a MILP model to strategically plan the staff size and composition in universities. The model is not only used for planning purposes, but also helps the university to assess different manpower strategies in terms of costs and structure.

For the numerical example they used a planning horizon of ten periods spread over eight years. Each period, hiring, dismissing and promotions were used to reach the target staffing level. These target levels were obtained from a forecasted number of yearly required teaching hours for each department. Furthermore, also forecasted attrition was used as an input. It was not mentioned how these forecasts were calculated. In contrast to most other articles on manpower planning, staff demand was assumed to be constant throughout the planning horizon for each unit and department.

The MILP formulation itself is rather straightforward and reminds of the crew transition problem for airlines (Holm, 2008; Hooijen, 2019). The objective was to minimise staffing cost and the gap between planned staff composition and the target level. De la Torre et al. (2016) chose to formulate this goal into one objective, rather than using multi-objective optimisation. This was done by assigning a cost to the aforementioned gap between planned and target staffing levels.

Application to the airline manpower planning problem:
Instead of choosing a multi-objective programming approach, De la Torre et al. (2016) combined two goals into one by assigning a cost value to the second goal. Although the model and goal are further rather straightforward, this approach could be used in the airline industry by for instance assigning a cost to crew satisfaction and days off in order to implement these into one objective.

2.3. Chapter Discussion

The goal of this chapter was to give an overview of the literature on the manpower planning problem, with a focus on strategic (long-term) planning. Since strategic manpower planning is the least researched phase in the airline industry, also other industries were looked at.

Four articles were discussed that investigated the strategic manpower planning for airlines. The objective always was to minimise costs and for this a MILP formulation was used (Verbeek, 1991; Holm, 2008; Morén, 2012), except for Hooijen (2019) who used a heuristic planning model. All articles focused on transitioning pilots to close the gap between supply and demand, with Morén (2012) also allowing for pilots to fly below rank and to fly on multiple aircraft types.

It was found that crew demand was always used as a given input and no article investigated the best strategy to determine long-term crew demand. Crew demand could be derived from the flight schedule (Hooijen, 2019) or from the expected crew utilisation (Holm, 2008) but no article explored the influence of the demand data's origin and quality on the outcome of the model.

It was also found that research is paying more attention to planning robustness by taking into account possible disruptions already in the planning phase. This can result in higher planned costs but lower actual costs. Schedule robustness can be measured by calculating the propagated delay in case of disruptions and can be achieved by using multiple objectives (Ehrgott and Ryan, 2002) or by using a stochastic formulation (Yen and Birge, 2006). No articles were found that take into account planning robustness already in the strategic phase.

In other industries, additional methods were used next to MILP formulations to determine the optimal staff size. In most articles, multiple objectives were identified and thus different methods were used to deal with this: goal programming (Trivedi, 1981; Horn et al., 2016), formulating it as one objective (Li et al., 2007; De la Torre et al., 2016), using chance constraints (Ganguly et al., 2014) or using a combination of techniques (Weigel and Wilcox, 1993). Multiple articles were found that allowed staff downgrading, similar to a pilot flying below rank (Li et al., 2007; Trivedi, 1981; Ganguly et al., 2014; Bard, 2004; Cai et al., 2004) in order to temporarily fill the gap between supply and demand.

Here too, staff demand was always used as an input and either based on historical data, or just a parameter without more details given. Li et al. (2007) however developed a separate demand forecasting module but did not specify how exactly it was calculated or what forecasting techniques were used. In the next chapter, various forecasting methods will be discussed.

3
Forecasting Methods for Manpower Supply and Demand

As already explained, strategic manpower planning has to deal with uncertainty. On the demand side, airlines still want to predict future crew demand without having a flight schedule ready. On the supply side, crew absence is the biggest uncertainty. In this chapter, various methods to forecast supply and demand will be discussed. In Section 3.1, current methods used by airlines to estimate future crew supply and demand will be explained. In Section 3.2, articles and research on possible forecasting methods will be analysed, both general mathematical models, and models for the airline industry and for other industries. Although not all presented models and methods are related to manpower planning, they could still be used to forecast for instance crew absence.

3.1. Current Methods Used in Airlines

Airlines usually do not like to disclose their manpower forecasting methods and therefore not a lot of literature has been written on this topic. Ciriani et al. (2013) identify two possibilities to determine crew demand: the airline's long-term fleet plan when looking more than a year ahead, and the flight schedules as soon as they are known (usually a year before the day of operation). Holm (2008) also identifies the possibility of estimating demand based on crew utilisation. All three possibilities will be explained below.

When estimating crew demand based on the airline's fleet plan, a so-called 'crew factor' can be used. This is a way for the airline to estimate how many crew members are needed per position to operate one aircraft of a certain type. For instance, if the airline estimates that they need six captains to operate one B737, and their long-term fleet plan states that they will be operating two B737s in two years, the estimated demand in two years is 12 captains for the B737. This crew factor depends on the type of airline and thus varies throughout the industry, although estimates fluctuate around 12 pilots per aircraft (CAE, 2017; Lovelace and Higgins, 2012). The crew factor can also be varied during the year to account for a busier or calmer schedule.

When using the initial flight schedule to estimate crew demand, 'production days' can be used. Production days are the number of working days in a crew pairing and are another way to express how many crew members are needed per position. Extra production days are used to account for training, vacation, sickness, etc. This way, a good demand estimate can be formed (Holm, 2008), but of course under the assumption that the flight schedule is already available in the strategic planning phase. An example calculation can be found in Table 3.1.

Lastly, also expected crew utilisation can be used in combination with an aircraft's expected number of block hours, based on either the (preliminary) flight schedule or on historical data (Ciriani et al., 2013; Yu et al., 2003). Block hours are the time between an aircraft leaving the departure gate and arriving at the destination gate and are usually calculated per month when estimating crew demand. The number of block hours per aircraft type is then divided by the target crew utilisation (i.e. the number of hours a crew member is supposed to fly per month) in order to get an estimate of the crew demand. The target

Table 3.1: Example calculation to determine the number of crew needed by using pairings and production days. The calculation is done for a single long-haul aircraft type (LH) and two positions (captain and first officer), with the assumption that this flight schedule is flown the entire year, and with the assumption that one captain and two first officers are needed per flight. All values are fictitious.

LH							**Production days**	
Monday	**Tuesday**	**Wednesday**	**Thursday**	**Friday**	**Saturday**	**Sunday**	**Captain**	**First Officer**
AMS-CUR	-	-	-	-	CUR-AMS		6	12
	AMS-PUJ	PUJ-AMS					2	4
Total per week:							8	16
Total per year:							416	832
Expected production days per person per year							220	220
Production days for training, days off, absence, ... per person per year							40	40
Production days left, per person per year							180	180
Crew needed							**416/180 = 2.3**	**832/16 = 4.6**
							⇒ 3	⇒ 5

crew utilisation varies per position and can be altered to account for peak and low season.

When it comes to determining future crew supply, airlines usually have a comprehensive knowledge about their staff. Everything starts with the number of employees, and then expected retirements are subtracted, together with expected sickness absence, parental leave, etc. The expected sickness absence can be a fixed number, or a dynamic one, based on historical data or recent trends. This will be explained more in the next section.

3.2. Forecasting Methods in Literature

Long-term manpower demand forecasting is probably the most difficult part of airline manpower planning, since the uncertainty is very high and the future demand depends on a lot of factors (Holm, 2008). Manpower supply forecasting is a bit easier, although there the uncertainty of employee absence plays an important role. Nonetheless, some methods could be used to take into account these uncertainties, as will be explained in the next sections. In Section 3.2.1 an overview of possible mathematical forecasting methods is given; Section 3.2.2 describes which forecasting methods are used in articles on manpower planning for airlines. Finally, Section 3.2.3 does the same for other industries.

3.2.1. Overview of Possible Methods

Safarishahrbijari (2018) has analysed 275 papers on workforce forecasting models in all kinds of industries and the article gave a good overview of possible forecasting methods. Seven major approaches were identified: qualitative models, optimisation models, generic mathematical models, statistics and regression, analytical stock-and-flow models, simulation models and time series models. An overview of these different approaches and their strengths and flaws will be given below, based on the article of Safarishahrbijari (2018).

Qualitative models are used in case of a lack of high-quality data and depend on the input of experts. Therefore this approach can sometimes be regarded as subjective and inaccurate, but can suffice in case the problem at hand is rather small.

Optimisation models have been used extensively already in the previously mentioned articles. Their goal always is to maximise or minimise (a) certain objective function(s), and several approaches exist, such as linear programming and goal programming. Optimisation models are very popular, but are mostly formulated linearly. In reality however, systems are rather dynamic and non-linear, and it can be hard and time-consuming to catch this with optimisation models.

The **generic mathematical methods** mentioned in Safarishahrbijari (2018) range from the use of linear differential equations to the labour multiplier approach to forecast future demand. Purely mathematical formulations can capture a the dynamics of a system accurately, but are impractical when it comes to characterising historical data.

Therefore, **statistical and regression models** have been developed. These models take into account the stochastic character of certain parameters and use historical data and a fitting process to predict future values. Although time series models also do this, statistical and regression models do not always have time as a parameter. For instance, regression analysis can determine "the causal effect of independent variables on dependent variables" (Safarishahrbijari, 2018). Again, this approach can be inefficient when modelling a dynamic system, since they usually consider only one output.

Analytical stock-and-flow models in manpower planning depict the employee cohorts as stocks and movement between them (such as promotions) as flows. This way of modelling makes the problem easy to visualise and understand, but it cannot be done for every problem (for instance in an airline the movement between crew positions can depend on bids, seniority, etc. and is thus not straightforward (Hooijen, 2018)).

The sixth method mentioned in Safarishahrbijari (2018) are **simulation models**. Simulation is a good way to investigate systems that cannot be solved analytically. This approach is preferred to answer what-if questions of complex problems, but are time-consuming to develop.

Finally, **time series models** are widely used to predict future values with the use of historical data. Several forecasting techniques exist, but the most used ones when it comes to manpower planning are Box-Jenkins, exponential smoothing and Markov modelling (Safarishahrbijari, 2018).

The Box-Jenkins method uses an autoregressive integrated moving average (ARIMA) to detect the best fit of a time series model to past values, and combines an autoregressive process with a moving average process (Commandeur and Koopman, 2007). Exponential smoothing on the other hand, uses the weighted moving average. In this case, exponentially decreasing weights are assigned to values that are further in the past (Verhagen, 2019). Finally, Markov modelling is based on transition probabilities between different states, and does only depend on the current state. However, Markov models are unable to interact with their environment and therefore are hard to use in the airline manpower planning problem, since crew transitions, seniority rules and crew bids all play a role there, similar to stock-and-flow models (Hooijen, 2018).

Catching seasonality and trends in time series is often difficult and cannot be done by every technique. Therefore, Qi and Zhang (2008) investigated what the best modelling approach would be when using artificial neural networks (NN) to model trend time series. They used a standard three-layer NN with one output node (the basics of neural networks and the use of it as a solution method for the manpower planning problem are further explained in Section 4.4). Monte Carlo simulations were then run in order to determine the best forecasting strategy out of a pool of five different data generating processes and four modelling approaches (all methods used NN). They concluded that modelling with differenced data was the best modelling approach when using NN for time series forecasting. This method detrends data by differencing.

Tseng et al. (2001) also tried to forecast seasonal time series, but they made use of grey models (GM). The theory of grey models was developed by Deng (1982) and is particularly good at dealing with poor information or small samples, and tries to look for patterns in this data. However this comes at a price: grey models are known to be bad at detecting seasonality in data. Tseng et al. (2001) wanted to overcome this by partially deseasonalising the data with the ratio-to-moving-average approach. The deseasonalised data was then used in combination with a standard GM(1,1) model. Mean squared error (MSE), mean absolute error (MAE) and mean absolute percentage error (MAPE) were used to evaluate the performance of the model. The authors concluded that their approach of combining a GM(1,1) model with deseasonalised data performed better than the neural network model, the SARIMA model and the GM(1,N) model.

3.2.2. Forecasting Methods in the Airline Industry

Rather than focusing on forecasting methods that are currently being used in airlines (Section 3.1), this section will focus on research and articles in this field. Within the airline industry, no articles were found that predict crew demand other than with the methods mentioned above in Section 3.1. Most research seems to focus on predicting crew supply, and more specifically predicting crew absence and modelling the use of reserve crew.

Homaie-Shandizi et al. (2016) developed a supervised learning method to forecast monthly pilot reserve hours. Although the developed model was meant to predict reserve hours of the next month after the new pairings were published (hence it is not strategic planning), the method they used might be applicable to long-term planning. Their model used historical data and monthly schedule characteristics, such as time away from base, number of landings, flying hours, etc. together with an iterative algorithm to predict how many reserve pilots would be needed in the next month. A modified version of classification and regression trees was used for the prediction model. This technique is based on Morgan and Sonquist (1963) and further explained in Breiman (2017). Their model resulted in a 15% improvement in predictions for the airline.

Van Drongelen et al. (2014) examined a possible correlation between different flight types and flight crew absence of more than seven days, in order to better predict and prevent it. For this, they used univariate and multivariate logistic regression on flight and absence data of more than 8000 employees. Logistic regression is a way to relate a binary result (like being absent or not) to one or more variables and is well explained in Kleinbaum et al. (2002). The univariate analysis of Van Drongelen et al. (2014) showed that multiple time zone crossings and a higher number of medium-haul flights resulted in a higher chance of crew absence. The multivariate analysis showed no significant correlation.

Something similar was done by Boot et al. (2017), but they used socio-demographic data, such as age, marital status and number of children, to predict long-term sickness for KLM employees. They too used logistic regression to find possible correlations. They were able to demonstrate a connection between long-term sickness leave and, for instance, higher age and a recent pregnancy, however the explained variance was low. The explained variance (or Nagelkerke's R^2) is a way of evaluating a model's goodness of fit and the closer the value is to zero, the worse the prediction capabilities of the model (Nagelkerke et al., 1991).

Over the years, Christopher Bayliss has worked on several methods to deal with the uncertainty of crew absence. In Bayliss et al. (2012), they developed a probabilistic model that scheduled reserve crew teams in such a way that the probability of crew unavailability was minimised. For their numerical example however, instead of taking absence probabilities from historical data, their input absence probabilities were uniformly distributed, resulting in an abnormal amount of disruptions. Several objective functions, such as minimising the sum of probabilities or the standard deviation, and heuristic solution methods, such as genetic algorithms and tabu search, were investigated. They concluded that the best objective function was to minimise the sum of probabilities of crew unavailability, and that the tabu search resulted in values close to optimality.

In Bayliss et al. (2019) they further improved this model by taking into account reserve crew induced delays and allowing multiple crew members to be absent on a single flight. This time, instead of using a uniform distribution for absence probability, it was assumed each crew member had a 1% chance of being absent. This number was based on historical data, but the model could also handle individual absence probabilities per crew member. The probabilistic model appeared to be much faster than classic simulation methods, such as Monte Carlo simulation.

3.2.3. Forecasting Methods in Other Industries

Since not many articles were found that investigate manpower forecasting methods in the strategic airline planning phase, this section will discuss methods that are used in other industries.

Ng et al. (2008) developed six different approaches to find optimal staffing levels before the employees' attendance rates are known (as is the case for strategic manpower planning). The goal was to find the right balance between satisfying the manpower demand and minimising staffing costs, since both under- and overstaffing can result in high costs.

The first approach replaced the unknown attendance rate by the average attendance rate; the second approach assumed the minimum attendance rate for all workers. These two approaches were deterministic and rather simple, but served as a good starting point for the other approaches.

The third approach used historical data to construct a single worst case attendance rate per worker type. The fourth approach first solved the problem with "different historical realisations of attendance rates" (Ng et al., 2008). Then, the maximum staffing level per worker type was used as a decision variable. The fifth approach used a two-stage stochastic LP programming and sample approximation. Lastly, the sixth approach used robust optimisation with ellipsoidal uncertainty sets, based on Ben-Tal and Nemirovski (1999). This method uses the assumption that the attendance rates are enclosed in an N-dimensional ellipsoid based on historical data. Realisations of the worst-case frontier were then also used in the fifth approach.

They concluded that stochastic programming (approach 5) and robust optimisation (approach 6) offered the best results in terms of staffing costs, however the fifth approach used considerably longer computation time.

Ho (2010) used grey models to forecast construction manpower demand using a limited amount of data, with a time horizon of one quarter. Ho (2010) proved that a single variable, first order grey model could still be used to predict manpower demand. He found that the optimal sample size was 5, which resulted in a mean absolute percentage error (MAPE) of only 3.21%, which can be considered very good. This MAPE also depended on the sample number: a higher sample number results in a lower MAPE if the data is very random; if the data is rather smooth, a lower sample number is better. It is thus important to vary the sample number in order to find the lowest MAPE.

In conclusion, Ho (2010) stated that more sophisticated grey models might perform even better than the basic GM(1,1) model used in his article, for instance the remnant GM(1,1) model, the GM(1,N) model, the GM(2,1) model or the Verhulst model. These models are explained further in Liu and Lin (2010).

3.3. Chapter Discussion

The objective of this chapter was to list and discuss several forecasting methods that can be used in the strategic planning phase of an airline to predict crew supply and demand.

Seven mathematical forecasting methods were identified, and their strengths and weaknesses were discussed. For the final approach, time series, several methods were further identified: Box-Jenkins, exponential smoothing, and Markov modelling. Since it can be hard to catch seasonality or trends with time-series, the use of neural networks and grey models were explored. Grey models are known to be bad at detecting seasonality, but this was overcome by deseasonalising the data first (Tseng et al., 2001).

Means to evaluate a model's goodness of fit were also identified: Nagelkerke's R^2 (Boot et al., 2017), mean squared error, mean absolute error and mean absolute percentage error (Tseng et al., 2001; Ho, 2010).

Three methods that are currently used by airlines to predict crew demand were identified: using the long-term fleet plan, using the flight schedule and preliminary crew pairings, or using the expected crew utilisation in combination with expected block-hours. No articles have been found that compare the effectiveness of these methods.

On the other hand, some articles were found that developed models to predict crew supply, and more specifically crew absence and reserve crew. This can be done by using supervised learning to find a relation between the flight schedule and reserve crew usage (Homaie-Shandizi et al., 2016), or by finding correlations between crew absence and operational factors (Van Drongelen et al., 2014) or socio-demographic factors (Boot et al., 2017). Another method to deal with the uncertainty is to introduce absence probabilities when making a crew plan. These absence probabilities are taken from historical data and could be calculated per position or even per crew member (Bayliss et al., 2012, 2019).

Other industries naturally also deal with employee absence. Out of six approaches, Ng et al. (2008) concluded that robust optimisation with ellipsoidal uncertainty sets was the best method to find optimal staffing levels before the attendance rate is known. Besides this, also a stochastic decision model (Cai et al., 2004) or grey models (Ho, 2010) could be used to deal with uncertain demand.

4

Solution Methods for the Manpower Planning Problem

This chapter will give an overview of possible solution methods for the strategic manpower planning problem. In general, most models described in this literature review were either solved by using a commercial LP solver or a heuristic method. Commercial LP solvers offer a classic and exact approach to find the solution to an optimisation problem, but usually have high computation times, especially for large problems such as the manpower planning problem. To deal with these high computation times, it is most common in research to either make adjustments to the exact solution method (such as branch and bound methods to reduce the solution space) or to use a heuristic method to find satisfactory (but sub-optimal) results. These methods can also be combined with simulations, and recently machine learning as well. All of these possible methods are explained in detail in sections 4.1, 4.2, 4.3 and 4.4 respectively.

4.1. Exact Methods

The easiest way to solve an optimisation problem formulated as a (mixed-integer) linear programming problem, is to programme the model into a **commercial solver** and let it find the optimal solution. Trivedi (1981) programmed his mixed integer goal programming model on nurse budgeting into a computer which solved it in about 60 seconds. The model was rather simple and assumed a constant nurse demand throughout the year, and this might explain the low computation time.

De la Torre et al. (2016) used the commercial optimisation software CPLEX Optimizer to solve their model to optimise the long-term academic staff size, with a planning period of eight years. Since the most complex scenario took so long to solve, they had to impose a maximum computing time of 10,000 seconds. In this case, the gap to the optimal solution was approximately 2%, which was considered very good given the time horizon of the planning.

Ganguly et al. (2014) also used CPLEX to solve their staff planning model for a medical emergency department. Optimal schedules were found within one second, although the model was heavily simplified in that case. When they allowed for different care providers to work as a team, the computation time increased to over five hours with an optimality gap of 3%. They concluded that it would be better to use techniques such as branch and price or column generation to decrease computation times in the future.

Although CPLEX now uses **branch and cut** by default when solving MIP problems (IBM, 2020), Júdice et al. (2005) specifically states the branch and cut version of the MIP solver in CPLEX was selected for the shift scheduling problem for a mail processing centre. The branch and cut method is a combination of two common solution methods: branch and bound, and cutting planes. It was designed in the 1980s to deal with growing problem sizes (more than 100 variables). Nowadays, the branch and cut method is suited to solve problems with many thousand variables (under some conditions) (Hillier and Lieberman, 2015). Although this method reduces computation time, Júdice et al. (2005) did not find an optimal solution after 12 hours.

Morén (2012) designed a new **branch and bound** method to solve the cockpit crew transitioning problem with a planning horizon of one to four years. His new method, called implication branching, was combined with the existing methods of reliability branching and distribution node selection. Implication branching uses knowledge that is specific to the problem at hand and uses the constraints to estimate how important certain variables are, and how many variables are influenced by a decision. By using this method, the same objective value was reached up to five times faster than with conventional branch and bound methods, although the solution time depended on the characteristics of the airline. For some airline configurations used in the numerical example, computation times of 144,000 seconds (40 hours) were still not enough to find the optimal value, although the method was still faster than conventional methods.

Finally, the concept of **column generation** to solve large integer programmes (Barnhart et al., 1998) has also proven its value in the airline manpower planning problem, especially in the scheduling phase (Kasirzadeh et al., 2015). Column generation only considers a feasible subset of variables in the beginning, and then only columns (variables) are added if they improve the objective function. Column generation was, among others, used by Gamache et al. (1999) to solve a large airline's rostering problem. In their case, computation times were reduced by a factor of a thousand.

4.2. Heuristic Methods

As Morén (2012) illustrated, the computation times to solve a realistic airline manpower planning problem using exact methods are extremely high. His model only included cockpit crew transitions, hence solving an even larger problem (by for instance including cabin crew, or extending the time horizon) borders on the impossible. For this reason, researchers often resort to heuristic methods. These provide sub-optimal, but acceptable results against lower computation time. Several methods and algorithms exist, as will be shown below.

One of the heuristic methods to solve large-scale integer problems is the use of **genetic algorithms** (GA). Genetic algorithms are based on Darwin's evolution theory and generally consist of five key concepts to mimic the process of natural selection: a population, a fitness function, selection, crossover and mutation. The fittest solution will be selected for reproduction and this way, after many generations, the optimal solution is found (Sivanandam and Deepa, 2008; Mallawaarachchi, 2017).

Cai and Li (2000) designed their own genetic algorithm to solve the scheduling problem of staff with mixed skills under multiple criteria. The goal was to integrate this model into an integrated scheduling system. Their proposed algorithm used a multi-point crossover, a parent selection based on the three objectives of the model, and an extra heuristic to solve infeasible crossovers. They concluded that the algorithm was only successful if the set of feasible schedules were identified *before* the GA was applied. This is because the algorithm is slowed down considerably if the problem at hand is large.

Škraba et al. (2016) also used genetic algorithms and combined it with stochastic local search to solve their manpower model for large hierarchical organisations and to determine the optimal manpower strategy in terms of transitions, recruitment, etc. Their fitness scores were calculated by taking the average number of transitions a person needed to reach his goal position.

Another method often used to solve large problems like these, is **tabu search**. This method uses a local search procedure and starts by only accepting improved solutions at each iteration to find a local optimum. However, the main characteristic of a tabu search is that it also accepts deteriorated solutions to avoid getting stuck in that local optimum. It then uses a memory structure to avoid returning to the previous solution in the next iteration. A tabu list keeps track of the forbidden moves (Hillier and Lieberman, 2015; Hooijen, 2018).

Thalén (2010) used tabu search to solve the cockpit crew planning problem. The planning horizon for his model was one year and he investigated both the staffing problem and the transition problem. The ultimate objective was to provide the right amount of pilots with the right qualifications against the lowest cost. The algorithm itself used two different neighbourhoods and best results were obtained when the tabu search was sequentially shifting between them. This method provided the same solutions as a commercial MIP solver (such as CPLEX), but was no less than 30 times faster. The problem at hand was quite large: it consisted of 1100 pilots, 13 positions and 3 bases. With a solution time of about 8000 seconds, a solution time that is 30 times faster makes a significant difference. The commercial solver only outdid the tabu search method if it was run for 2.5 days.

Tabu search was also applied to the cockpit crew transition problem by Hooijen (2019) to avoid awarding transitions between the same pilot positions constantly. These moves were added to the tabu list and thus the algorithm was forced to look for other, and eventually better solutions. The best transition option in the neighbourhood was then picked by a selection algorithm using a tree search, combined with either a naive selection algorithm, a greedy algorithm or Dijkstra's shortest path algorithm to decrease search space and computation time. The model and algorithm themselves were already explained in Section 2.1.3.

The lowest computation time was reached with the naive selection algorithm. This method however resulted in the worst solution quality and stability. He therefore concluded that the optimal configuration, in terms of both computation time and solution quality, was the shortest path algorithm based on Dijkstra's algorithm (Dijkstra et al., 1959).

A **greedy heuristic** was also used by Bard (2004) to solve the break assignment problem during staff scheduling since the problem took too long to solve using CPLEX. Different kinds of greedy algorithms exist, but in general a greedy algorithm chooses the best option at each iteration, hence it is a fast but short-sighted algorithm and it easily produces sub-optimal solutions. Dijkstra's algorithm, as used by Hooijen (2019), is an optimised version of a greedy algorithm, and looks more ahead (Gendreau et al., 2010).

In Bard (2004), a greedy heuristic was used to assign breaks to shifts. The algorithm went over the shifts chronologically and assigned available breaks. An assigned break was then removed from the list with available breaks. This was done until all shifts had breaks.

Horn et al. (2016) used a **fix-and-relax** heuristic to solve their mixed-integer linear programming model on strategic workforce planning in the army. The name *fix-and-relax* comes from first fixing, and then relaxing the model's variables. Since it took too long to solve with a commercial LP solver, they investigated different heuristic methods and chose to adopt the fix-and-relax method, based on Escudero and Salmeron (2005). Their heuristic method consisted of two stages: in stage 1 the goal was to meet the staffing targets; in stage 2 the goal was to minimise operational costs.

To clarify the concept of fix-and-relax, the algorithm of stage 1 will be explained. Integrality constraints were first relaxed to obtain an initial feasible solution. Then, all variables were fixed and the objective coefficients were set to zero. Then, per hierarchical staff position variables were relaxed and the model was solved (with integrality but still with objective coefficients of zero). In the end, all variables were relaxed, objective coefficients were released and the model was solved entirely.

The model, consisting of the two stages, and for an army with 41 trades, 7 ranks and a time horizon of 17 years, was solved in approximately 8000 seconds (2.2 hours). For stage 1 only, the commercial solver on the other hand needed more than 24 hours.

4.3. Simulation

Simulation is often used when the problem at hand is large and complex and when uncertainty is involved. Simulations may not have mathematical solutions and are good way to investigate systems or problems that cannot be solved analytically. Although they are an effective way to deal with uncertainties and to analyse different scenarios, they are often time-consuming to develop (Safarishahrbijari, 2018).

Verbeek (1991) made use of simulation in his strategic manpower decision support system. This was done to deal with uncertain pilot behaviour, but it is not further explained how exactly the simulation was set up. His model has a planning horizon of ten years and thus uncertainty plays a big role when planning this much in advance, hence it is a pity that no more details were given.

Qi and Bard (2006) also developed a simulation model for their research to determine the optimal staff size and composition for a postal service centre. Simulation was used to validate their results, and to determine whether the employee schedules meet the service standards. The outline of their model could already be seen in Figure 2.5 on page 40.

Atlason et al. (2004) used a combination of simulation and an iterative cutting plane method (an exact method) to minimise staffing costs in a service system, while satisfying a certain service level.

This service level was expressed by a service level function and it was assumed that this function was so complex that the algebraic form was unknown, hence simulation was the only way to evaluate its value. Furthermore, by combining integer programming with simulation, they only needed to simulate a small portion of the possible solutions.

Finally, Bayliss et al. (2017) used simulation to generate disruption scenarios for an airline. The goal of their model was to schedule reserve crew under crew absence uncertainty. To do this, they first simulated the airline without reserve crew to create disruption scenarios, with stochastic inputs such as flight times and crew absence. These scenarios were then used as input for the mixed integer programming model that determined which reserve crew schedules would have a positive impact on the simulated disruptions. The simulation could then be run again to validate these proposed reserve crew schedules.

4.4. Machine Learning

The technique of machine learning is used to categorise algorithms that improve their performance by experience. Machine learning is often used in cases where the algorithm itself is unknown, but the in- and output are known (e.g. distinguishing spam emails from normal ones, facial recognition, etc.). In this case, the computer is trained to generate the algorithm itself by using lots of in- and output data (Alpaydin, 2014). Machine learning can be combined with exact methods or heuristic search methods to improve the solution quality or lower the computation time. Although machine learning has many applications, it is not often used for the manpower planning problem.

Thathachar and Sastry (2002) described several **learning automata** (LA), which are a type of machine learning algorithms for finite action sets. The main characteristic of LA is that the probability distribution over the action set is updated at each iteration and depends on the previous actions and their results. The choice of which action to take is then based on this probability distribution. As the algorithm learns, the actions that are taken have an increasing chance of improving the objective function.

Beulen et al. (2020) used **neural networks** (NN) to evaluate airline crew's flight requests. Neural networks are a form of supervised machine learning and are based on biological neural networks in a brain and consist of four main components: an input layer, (a) hidden layer(s), an output layer and nodes, called neurons (Alpaydin, 2014).

In Beulen et al. (2020), the goal was to assess flight requests in a cost-efficient way. This was done by assigning a score to each request based on the cost impact this request may have in a later stage. They used two hidden layers and 100 neurons, and the algorithm was combined with a mixed integer linear programming model to optimise the rostering. The neural network algorithm was tested on data of a major European airline and granted 22% more requests while using nearly the same workforce size.

No articles have been found that use machine learning for the strategic manpower planning problem in airlines.

4.5. Chapter Discussion

The goal of this chapter was to give an overview of possible solution methods for the strategic manpower planning problem. Four major solution methodologies were identified: exact methods, heuristic methods, simulation and machine learning.

Exact methods mostly use a commercial solver to find the optimal objective value of a linear programming formulation. Multiple techniques exist to decrease the solution space and reduce computation time, such as column generation (Barnhart et al., 1998), branch and cut (Júdice et al., 2005) or branch and bound (Morén, 2012). However, when problems start to get very big, as is the case for airlines, computation times increase drastically, as Júdice et al. (2005) and Morén (2012) illustrated.

For this reason heuristic algorithms can be used that find sub-optimal but acceptable results against a lower computation time. Four algorithms were explained: genetic algorithms, tabu search, greedy heuristics and fix-and-relax heuristics. Cai and Li (2000) used genetic algorithms to solve a crew scheduling problem but concluded that it would be better to identify the sets of feasible schedules beforehand to decrease computation time. Tabu search is a popular method to solve large problems with promising results compared to commercial solvers. Thalén (2010) showed that acceptable results were reached with a tabu search after only a thirtieth of the time a commercial solver needed. Next, it was shown that simple greedy heuristics are fast but often produce far from optimal results. For this reason, it is favourable to use a more complex greedy heuristic such as Dijkstra's algorithm (Hooijen, 2019). Finally, Horn et al. (2016) demonstrated that by using a fix-and-relax heuristic large computation times can be reduced significantly.

Simulation can also be used to solve large and complex problems with uncertainty involved, or when certain parameters are impossible to formulate in an algebraic way (Atlason et al., 2004). Verbeek (1991) simulated pilot's behaviour but gave no further details. Simulation can also be used to validate results (Qi and Bard, 2006). Finally, simulation can also be used to generate different disruption scenarios and use these scenarios as an input to planning and scheduling models (Bayliss et al., 2017).

Ultimately the technique of machine learning, such as learning automata and neural networks was explained. Although no articles were found that use machine learning for the strategic manpower planning problem, Beulen et al. (2020) showed that it is absolutely possible to use machine learning in the airline industry and that it is capable of significantly improving results.

5

Research Gap and Opportunity

This chapter will describe the identified research gap based on the literature review. Next, the research opportunity will be presented in the form of a research question and several objectives.

5.1. Research Gap

The airline manpower planning problem has been called one of the most important ones in the airline industry, and labour costs are the second biggest expense for airlines (Sohoni et al., 2004; Belobaba et al., 2009). Yet, airlines generally start their manpower planning only one to one and a half year in advance, and most research focuses on the operational planning phase: developing optimal, cost-minimised crew schedules. Long-term, strategic manpower planning is a mostly unexplored area, especially forecasting crew demand and supply before the flight schedule is known.

Holm (2008) has called this crew demand forecasting the most difficult part of manpower planning. Methods used by airlines now depart either from the flight schedule, or from the fleet plan and then apply a crew factor, but no research has been conducted that analyses the effectiveness of different crew demand forecasting methods. Several mathematical forecasting models have been identified, such as time-series models, logistic regression, grey models and artificial neural networks.

Next to this, airlines have started to focus on planning robustness, i.e. making sure that the staffing levels can deal with disruptions or unexpected circumstances so that the impact on the operations is minimised. The idea is that higher planned costs may result in lower actual costs. Whereas earlier research focused mainly on cost minimisation, it now also focuses assuring a certain level of robustness in the planning. However, no articles were found that consider robustness already in the strategic manpower planning phase.

Finally, most articles on manpower planning, be it in the airline industry or beyond, make use of (MI)LP formulations and then solve it by either using a commercial solver or by using heuristic methods if the problem size becomes too large. Machine learning techniques, such as learning automata or neural networks, are barely used, despite having shown their value (Thathachar and Sastry, 2002; Beulen et al., 2020).

5.2. Research Opportunity

Based on the previous section the goal of the proposed research can be formulated in a research question:

> *How to model airline crew demand and determine the optimal staff size per crew position in the strategic planning phase before the flight schedule is known, while ensuring planning robustness?*

Research objectives:

1. Identify the contributing factors to and their effect on crew demand.
2. Analyse different crew demand forecasting methods and assumptions before the flight schedule is known.
3. Develop a strategic manpower planning tool that provides insight into future crew demand and the optimal staff size.
4. Explore options to measure and ensure planning robustness already in the strategic planning phase.

The proposed research question comprises two subproblems: modelling airline crew demand already in the strategic planning phase, before the flight schedule is known, and determining the ideal staff size while ensuring a certain robustness level. Both parts would need an output per crew position, and interact with each other in order to make it an integrated planning tool.

The objective of the tool would not be to determine the best strategy to close the gap between crew demand and supply, but rather to determine the optimal crew size in terms of cost and robustness already in the strategic planning phase before a flight schedule is known.

The first part would identify the various factors contributing to crew demand (flights, crew absence, etc.) and would investigate different forecasting methods and their effectiveness, such as time-series, neural networks, but also the methods currently used in the airline.

The second part would ensure planning robustness. For this, a measure of robustness in the strategic planning phase needs to be decided first. Articles found on planning robustness focused on the operational planning phase and used the chance of propagated delay as a measure of robustness. Clearly, this metric cannot be used for robustness in the strategic phase, since the flight schedule is not known yet, and it is all about planning robustness instead of scheduling robustness.

The goal of the integrated model would be to minimise costs but maximise robustness. These two objectives are in conflict but this can be solved by either using bi-criteria optimisation (Ehrgott and Ryan, 2002), by converting a multi-objective model into a single objective one (De la Torre et al., 2016; Li et al., 2007) or by assigning a robustness score to a crew size and improve it with machine learning techniques.

III

Supporting Work

A
Results with Different Numbers of Scenarios

This appendix presents the results of Case 1 and Case 2 with varying number of generated crew demand scenarios. The analysis is done for 8, 10, 12, 15, 20 and 50 generated scenarios per repetition. The results are the average over 20 repetitions.

In Table A.1 the results for Case 1 are shown. It can be seen that the average objective value increases with the number of scenarios. When generating more demand scenarios, the generated samples come from more extreme points in the demand distributions. This means that with increasing number of generated scenarios, there will be more scenarios with extreme demand values.

The model deals with these more extreme demand values in two ways. When there is much lower demand than supply, the model cannot do much: in Case 1, no dismissals are allowed and thus the model relies on the natural outflow to lower the total crew supply. This results in more crew than needed and thus higher costs. On the other hand, when there is much higher demand than supply, the model can hire temporary and permanent crew members. This also results in higher costs.

Table A.1: Results for Case 1 for different numbers of scenarios.

Scenarios (k)	Repetitions	Objective value [MU]			
		Min	Max	Avg	Δ(min,max) [%]
8	20	56445	56714	56572	0.48
10	20	56483	56804	56622	0.57
12	20	56563	56824	56674	0.46
15	20	56605	56857	56717	0.44
20	20	56683	57026	56793	0.60
50	20	56871	57201	57005	0.58

The same happens for Case 2, as can be seen in Table A.2. Here too, the objective value increases with an increasing number of generated scenarios. In Case 2, the model has the option to fire permanent crew members. Hence, when the crew demand is much lower than the supply, the model will choose to fire employees. However, this comes with a high cost. This means that again, both more extreme high and low demand scenarios result in higher costs.

Table A.2: Results for Case 2 for different numbers of scenarios.

Scenarios (k)	Repetitions	Objective value [MU]			
		Min	Max	Avg	Δ(min,max) [%]
8	20	53682	54149	53940	0.87
10	20	53828	54253	54067	0.78
12	20	54039	54515	54268	0.88
15	20	54147	54616	54390	0.86
20	20	54317	54690	54522	0.68
50	20	54791	55197	55053	0.74

B

Descriptive Sampling

The methodology used in this report adopts the principle of Latin hypercube sampling (LHS) to generate correlated crew demand scenarios. However, if the variables that need to be sampled are not correlated and independent, LHS is not needed. In this case, scenarios can be generated using the method of descriptive sampling, as used in Listes and Dekker (2005). This method divides the distribution in equal quantiles to generate scenarios. The different steps are explained below and are illustrated with examples obtained with the same data as used in our research.

1. Start with the monthly block hour demand distributions. These distributions are the convoluted daily Beta block hour distributions as presented in the report. They can be approximated by a normal distribution, i.e. $d_i \sim N(\mu_i, \sigma_i)$ with probability distribution function F_i. This can be seen in Figure B.1a.

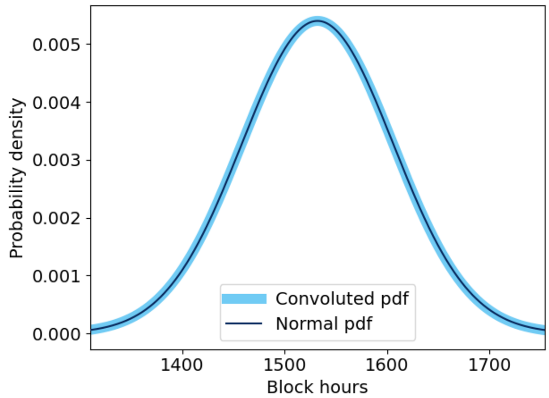

(a) Convoluted, normal block hour distribution.

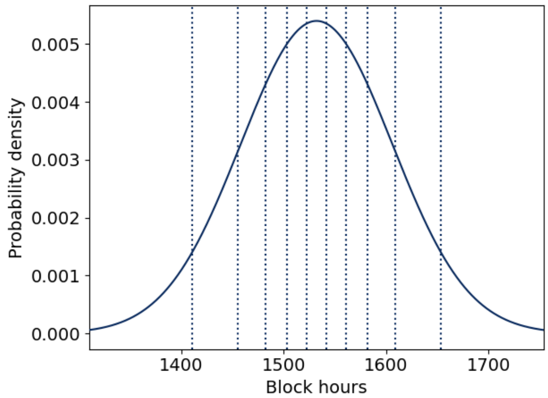

(b) Convoluted, normal block hour distribution divided in evenly spaced quantiles.

Figure B.1: Example of descriptive sampling.

2. Specify the number of scenarios (or *samples*) K that need to be generated. These scenarios d_{ij} will be sampled from distribution i at equally spaced quantiles according to:

$$d_{ij} = F_i^{-1}\left(\frac{j - 0.5}{K}\right) \quad j = 1, 2, ..., K \tag{B.1}$$

   Dividing the distribution into equally spaced quantiles makes sure the sampling occurs in a fair way. At lower densities, fewer demand values are sampled than at higher densities.

   This step is illustrated in Figure B.1b with $K = 10$. Here it can be seen that the quantiles are closer together at higher densities. The sample values are taken where the dashed quantile lines intersect the normal distribution. This results in a vector with the following 10 values: [1410, 1455, 1482, 1504, 1523, 1541, 1560, 1582, 1609, 1654]. Repeating this step for all N monthly block hour distributions hence results in N vectors.

3. Make a random permutation of the values d_{ij}, $j = 1, 2, ..., K$, for each $i = 1, 2, ..., N$ with N the number of distributions. This will result in K vectors (i.e. scenarios) of the form $[d_{1,j}, d_{2,j}, ..., d_{N,j}]$ with $j = 1, 2, ..., K$. Each scenario has the same probability $1/K$.

   To illustrate this, we can take the vector from Step 2 and assume another vector with sample values from another distribution, shown in Figure B.2. Then we have:

   Vector 1: [1410, 1455, 1482, 1504, 1523, 1541, 1560, 1582, 1609, 1654]

   Vector 2: [1210, 1272, 1309, 1338, 1365, 1390, 1417, 1446, 1483, 1545]

   Making a random permutation would then, for instance, result in the scenarios presented in Table B.1. Each scenario in the table has a probability of 0.1.

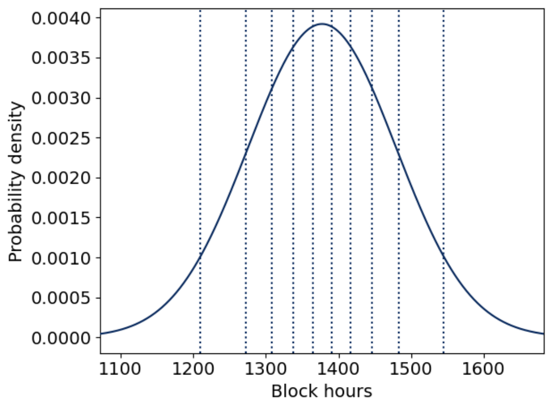

Figure B.2: Distribution used to generate samples for vector 2.

Table B.1: Example of random permutation of two distributions with 10 scenarios.

Scenario	d_1	d_2
1	1504	1365
2	1523	1272
3	1582	1446
4	1410	1309
5	1609	1338
6	1482	1390
7	1541	1417
8	1560	1545
9	1654	1210
10	1455	1483

C Implementation in Python

This appendix will explain how the convolution of the daily distributions and the Latin hypercube sampling method were implemented in Python.

C.1. Convolution of the Daily Distributions

The daily block hour distributions, who follow a Beta distribution, need to be convoluted to come to a monthly distribution. With n being the number of days in the respective month, the Beta distribution needs to be convoluted n times. This was implemented in Python using the following steps:

1. Discretize the Beta distribution's probability density function (pdf) into a probability mass function (pmf). For this the `pdf` and `pmf` functions from the `scipy.stats` library are used[1].

2. Create an array of size $n \times m$ with n the number of days in the respective month and m the number of discretization points.

3. Compute the one-dimensional discrete Fourier transform of the array. For this, the function `fft.fft` from the `Numpy` library is used[2]. This function uses the Fast Fourier Transform algorithm. This results again in an array of size $n \times m$

4. For each column in the resulting array, take the product of all elements. For this, the function `prod` from the built-in `math` module is used. This results in a one-dimensional array of length m.

5. Compute the one-dimensional inverse discrete Fourier transform of the resulting array. For this, the function `fft.ifft` from the `Numpy` library is used. This again results in a one-dimensional array of length m.

6. Access the real part of the output by using the `.real` command. These values form the pmf of the convoluted distribution.

C.2. Latin Hypercube Sampling

The Latin hypercube sampling (LHS) method was already described in the report. These steps are repeated here and for every step it is explained how it is implemented in Python.

1. Create a random $(k \times n)$ matrix $\mathbf{Z}^*$, with k the number of scenarios and n the number of crew positions. This matrix contains k Latin hypercube samples of size n from a standardised normal distribution. For this, the function `lhs` from the `PyDOE2` package[3] is used with the `center` criterion activated.

[1] `https://scipy.org/`
[2] `https://numpy.org/`
[3] `https://pypi.org/project/pyDOE2/`

The correlation matrix $\mathbf{I}^*$ of the samples and the identity matrix **I** do not coincide, hence they are not independent. To induce the desired correlation, the the $(k \times n)$ matrix **Z** is created by using lower triangular Cholesky decomposition (Golub and Van Loan, 1996):

$$\mathbf{I} = \mathbf{C} \cdot \mathbf{C}^{\mathsf{T}} \tag{C.1}$$

$$\mathbf{I}^* = \mathbf{E} \cdot \mathbf{E}^{\mathsf{T}} \tag{C.2}$$

$$\mathbf{Z} = \mathbf{Z}^* \cdot \mathbf{C} \cdot \mathbf{E}^{-1} \tag{C.3}$$

Z contains k independent samples of size n from a standardised normal distribution.

The Cholesky decomposition is implemented in Python by using the `cholesky` function from Numpy's `linalg` library. The correlation matrix $\mathbf{I}^*$ is obtained with the built-in `corr` function. The standardised normal distribution is obtained with Scipy's `norm.ppf` function.

2. Create a random $(k \times n)$ matrix **G**. This matrix contains k samples from a standardised normal distribution with the previously obtained correlation matrix **B** containing the correlation between crew positions. In other words, **B** contains the *desired* correlation. By applying lower triangular Cholesky decomposition, this desired correlation is induced in **Z**:

$$\mathbf{B} = \mathbf{P} \cdot \mathbf{P}^{\mathsf{T}} \cdot \mathbf{G} = \mathbf{Z} \cdot \mathbf{P} \cdot \mathbf{C}^{-1} \tag{C.4}$$

In Python, the Cholesky decomposition is again done with Numpy's `cholesky` function. Transposing and multiplying matrices can be done with Numpy's `transpose` and `dot` functions.

3. Apply the inverse transformation method to **G** to create $(k \times n)$ matrix **D**. The matrix **D** complies with the desired marginal distributions at each crew position. The inverse transformation method states that applying the inverse cumulative distribution function of any distribution F to a random variable with $U(0, 1)$ distribution results in a random variable whose distribution is exactly F (Ross, 2014). In this case, F is always a normal distribution, since we know that the monthly block hours are normally distributed. The correlation matrix of the generated samples is now equal to the desired correlation matrix **B**.

For this step, the function `cdf` from the `scipy.stats` library was used.

Bibliography

Alpaydin, E. (2014). *Introduction to machine learning*. MIT press, third edition.

Altenstedt, F., Thalén, B., and Sjögren, P. (2017). Solving the Airline Manpower Planning Problem. In *Proceedings of the 13th Workshop on Models and Algorithms for Planning and Scheduling Problems*.

Atlason, J., Epelman, M. A., and Henderson, S. G. (2004). Call center staffing with simulation and cutting plane methods. *Annals of Operations Research*, 127(1-4):333–358.

Bard, J. F. (2004). Staff scheduling in high volume service facilities with downgrading. *Iie Transactions*, 36(10):985–997.

Bard, J. F., Morton, D. P., and Wang, Y. M. (2007). Workforce planning at USPS mail processing and distribution centers using stochastic optimization. *Annals of Operations Research*, 155(1):51–78.

Barnhart, C., Belobaba, P., and Odoni, A. R. (2003). Applications of operations research in the air transport industry. *Transportation science*, 37(4):368–391.

Barnhart, C., Johnson, E., Nemhauser, G., Savelsbergh, M., and Vance, P. (1998). Branch-and-Price: Column Generation for Solving Huge Integer Programs. *Operations Research*, 46.

Bayliss, C., De Maere, G., Atkin, J. A. D., and Paelinck, M. (2017). A simulation scenario based mixed integer programming approach to airline reserve crew scheduling under uncertainty. *Annals of Operations Research*, 252(2):335–363.

Bayliss, C., Maere, G., Atkin, J., and Paelinck, M. (2019). Scheduling airline reserve crew using a probabilistic crew absence and recovery model. *Journal of the Operational Research Society*, pages 1–23.

Bayliss, C., Maere, G. D., Atkin, J., and Paelinck, M. (2012). Probabilistic Airline Reserve Crew Scheduling Model. In Delling, D. and Liberti, L., editors, *12th Workshop on Algorithmic Approaches for Transportation Modelling, Optimization, and Systems*, volume 25 of *OpenAccess Series in Informatics (OASIcs)*, pages 132–143, Dagstuhl, Germany. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik.

Belobaba, P., Odoni, A., and Barnhart, C. (2009). *The Global Airline Industry*. Aerospace Series. Wiley.

Ben-Tal, A. and Nemirovski, A. (1999). Robust solutions of uncertain linear programs. *Operations research letters*, 25(1):1–13.

Beulen, M., Scherp, L., and Santos, B. F. (2020). Dynamic Evaluation of Airline Crew's Flight Requests Using a Neural Network. *EURO Journal on Transportation and Logistics*, page 100018.

Boot, C. R. L., van Drongelen, A., Wolbers, I., Hlobil, H., van der Beek, A. J., and Smid, T. (2017). Prediction of long-term and frequent sickness absence using company data. *Occupational Medicine*, 67(3):176–181.

Breiman, L. (2017). *Classification and Regression Trees*. Routledge.

CAE (2017). Airline Pilot Demand Outlook, 10-year view. `https://www.cae.com/media/documents/Civil_Aviation/CAE-Airline-Pilot-Demand-Outlook-Spread.pdf`. Accessed in April 2020.

Cai, X. and Li, K. N. (2000). A genetic algorithm for scheduling staff of mixed skills under multi-criteria. *European Journal of Operational Research*, 125(2):359–369.

Cai, X., Li, Y., and Tu, F. (2004). Solving manpower planning problem with two types of jobs under uncertainty demand. *International Journal of Pure and Applied Mathematics*, 17(3):327–351.

Charnes, A. and Cooper, W. W. (1959). Chance-constrained programming. *Management science*, 6(1):73–79.

Ciriani, T. A., Fasano, G., Gliozzi, S., and Tadei, R., editors (2013). *Operations research in space and air*, volume 79, pages 413–414. Springer Science & Business Media.

Clarke, M. and Smith, B. (2004). Impact of operations research on the evolution of the airline industry. *Journal of aircraft*, 41(1):62–72.

Commandeur, J. J. F. and Koopman, S. J. (2007). *An introduction to state space time series analysis*. Oxford University Press.

De la Torre, R., Lusa, A., and Mateo, M. (2016). A MILP model for the long term academic staff size and composition planning in public universities. *Omega*, 63:1–11.

Deng, J.-L. (1982). Control problems of grey systems. *Systems & control letters*, 1(5):288–294.

Dijkstra, E. W. et al. (1959). A note on two problems in connexion with graphs. *Numerische mathematik*, 1(1):269–271.

EASA (2019). *Age Limitations for Commercial Air Transport Pilots*. Final report.

Ehrgott, M. and Ryan, D. (2002). Constructing robust crew schedules with bicriteria optimization. *Journal of Multi-Criteria Decision Analysis*, 11:139–150.

Escudero, L. F. and Salmeron, J. (2005). On a fix-and-relax framework for a class of project scheduling problems. *Annals of operations research*, 140(1):163.

Gamache, M., Soumis, F., Marquis, G., and Desrosiers, J. (1999). A column generation approach for large-scale aircrew rostering problems. *Operations research*, 47(2):247–263.

Ganguly, S., Lawrence, S., and Prather, M. (2014). Emergency department staff planning to improve patient care and reduce costs. *Decision Sciences*, 45(1):115–145.

Gendreau, M., Potvin, J.-Y., et al. (2010). *Handbook of metaheuristics*, volume 2. Springer.

Golub, G. H. and Van Loan, C. F. (1996). Matrix Computations, 3rd ed.

Grinold, R. C. and Marshall, K. T. (1977). *Manpower planning models*. North-Holland New York.

Hillier, F. S. and Lieberman, G. J. (2015). *Introduction to operations research, international edition*. McGraw-Hill Education, tenth edition.

Ho, P. H. K. (2010). Forecasting construction manpower demand by gray model. *Journal of Construction Engineering and Management*, 136(12):1299–1305.

Holm, Å. (2008). Manpower Planning in Airlines: Modeling and Optimization. Master's thesis, Linköpings Universitet.

Homaie-Shandizi, A.-H., Partovi Nia, V., Gamache, M., and Agard, B. (2016). Flight deck crew reserve: From data to forecasting. *Engineering Applications of Artificial Intelligence*, 50:106–114.

Hooijen, A. A. J. (2018). Literature Study: Strategic Airline Cockpit Crew Planning. Master's thesis, Delft University of Technology.

Hooijen, A. A. J. (2019). Cockpit crew transition planning optimisation. Master's thesis, Delft University of Technology.

Horn, M. E. T., Elgindy, T., and Gomez-Iglesias, A. (2016). Strategic workforce planning for the Australian Defence Force. *Journal of the Operational Research Society*, 67(4):664–675.

IBM (2020). Branch and cut in CPLEX. https://www.ibm.com/support/knowledgecenter/SS9UKU_12.10.0/com.ibm.cplex.zos.help/refcppcplex/html/branch.html. Accessed in August 2020.

Ives, C. D. (1992). Promotional Planning for Airline Technical Crew.

Júdice, J., Martins, P., and Nunes, J. (2005). Workforce planning in a lotsizing mail processing problem. *Computers & operations research*, 32(11):3031–3058.

Kasirzadeh, A., Saddoune, M., and Soumis, F. (2015). Airline crew scheduling: models, algorithms, and data sets. *EURO Journal on Transportation and Logistics*, 6.

Kleinbaum, D. G., Dietz, K., Gail, M., Klein, M., and Klein, M. (2002). *Logistic regression*. Springer.

Li, C. (2015). Chance-constraint method. https://optimization.mccormick.northwestern.edu/index.php/Chance-constraint_method. Accessed in June 2020.

Li, Y., Chen, J., and Cai, X. (2007). An integrated staff-sizing approach considering feasibility of scheduling decision. *Annals of Operations Research*, 155(1):361–390.

Listes, O. and Dekker, R. (2005). A scenario aggregation–based approach for determining a robust airline fleet composition for dynamic capacity allocation. *Transportation Science*, 39(3):367–382.

Liu, S. and Lin, Y. (2010). *Grey systems: theory and applications*. Springer Science & Business Media.

Lovelace, K. and Higgins, J. (2012). US pilot labor supply. In *US/Europe International Aviation Safety Conference, Cleveland Ohio USA June*, volume 12, page 14.

Mallawaarachchi, V. (2017). Introduction to Genetic Algorithms. Accesed in June 2020.

Medard, C. P. and Sawhney, N. (2007). Airline crew scheduling from planning to operations. *European Journal of Operational Research*, 183(3):1013–1027.

Morén, B. (2012). Utilizing problem specific structures in branch and bound methods for manpower planning. Master's thesis, Linköpings Universitet.

Morgan, J. N. and Sonquist, J. A. (1963). Problems in the analysis of survey data, and a proposal. *Journal of the American statistical association*, 58(302):415–434.

Nagelkerke, N. J. D. et al. (1991). A note on a general definition of the coefficient of determination. *Biometrika*, 78(3):691–692.

Ng, T. S., Huang, H. C., and Ng, J. Y. (2008). Human resource planning with worker attendance uncertainty. In *2008 IEEE International Conference on Industrial Engineering and Engineering Management*, pages 364–368. IEEE.

Qi, M. and Zhang, G. P. (2008). Trend time–series modeling and forecasting with neural networks. *IEEE Transactions on neural networks*, 19(5):808–816.

Qi, X. and Bard, J. F. (2006). Generating labor requirements and rosters for mail handlers using simulation and optimization. *Computers & Operations Research*, 33(9):2645–2666.

Ross, S. M. (2014). *Introduction to probability and statistics for engineers and scientists*. Academic Press, fifth edition.

Safarishahrbijari, A. (2018). Workforce forecasting models: A systematic review. *Journal of Forecasting*, 37(7):739–753.

Santos, B. F. (2018a). Lecture 1 - Introduction, Planning Framework and Demand Analysis. AE4423 lecture slides, Delft University of Technology.

Santos, B. F. (2018b). Lecture 8 - Scheduling Planning 2. AE4423 lecture slides, Delft University of Technology.

Sivanandam, S. N. and Deepa, S. N. (2008). *Introduction to Genetic Algorithms*. Springer.

Škraba, A., Stanovov, V., Semenkin, E., and Kofjač, D. (2016). Hybridization of stochastic local search and genetic algorithm for human resource planning management. *Organizacija*, 49(1):42–54.

Sohoni, M., Johnson, E., and Bailey, T. (2006). Operational airline reserve crew planning. *J. Scheduling*, 9:203–221.

Sohoni, M. G., Johnson, E. L., and Bailey, T. G. (2004). Long-Range Reserve Crew Manpower Planning. *Management Science*, 50(6):724–739.

Thalén, B. (2010). Manpower planning for airline pilots: A tabu search approach. Master's thesis, Linköpings Universitet.

Thathachar, M. A. L. and Sastry, P. S. (2002). Varieties of learning automata: an overview. *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)*, 32(6):711–722.

Trivedi, V. M. (1981). A mixed-integer goal programming model for nursing service budgeting. *Operations Research*, 29(5):1019–1034.

Tseng, F.-M., Yu, H.-C., and Tzeng, G.-H. (2001). Applied hybrid grey model to forecast seasonal time series. *Technological Forecasting and Social Change*, 67(2-3):291–302.

Van Drongelen, A., Van der Beek, A., Penders, G., Hlobil, H., Smid, T., and Boot, C. (2014). Sickness absence and flight type exposure in flight crew members. *Occupational medicine (Oxford, England)*, 65.

Verbeek, P. J. (1991). Decision support systems—an application in strategic manpower planning of airline pilots. *European Journal of Operational Research*, 55(3):368–381.

Verhagen, W. (2019). Lecture 7: Supportability engineering: basic concepts and assumptions; forecasting spare parts demand. AE4465 lecture slides, Delft University of Technology.

Weigel, H. S. and Wilcox, S. P. (1993). The Army's personnel decision support system. *Decision Support Systems*, 9(3):281–306.

Winston, W. L. and Goldberg, J. B. (2004). *Operations research: applications and algorithms*, volume 3. Thomson/Brooks/Cole Belmont^ eCalif Calif.

Yen, J. W. and Birge, J. R. (2006). A Stochastic Programming Approach to the Airline Crew Scheduling Problem. *Transportation Science*, 40:3–14.

Yu, G., Dugan, S., and Argüello, M. (1998). Moving toward an integrated decision support system for manpower planning at Continental Airlines: Optimization of pilot training assignments. In *Industrial Applications of Combinatorial Optimization*, pages 1–24. Springer.

Yu, G., Pachon, J., and Thengvall, B. (2003). Optimization-based integrated manpower management for airlines. In *Operations Research in Space and Air*, pages 407–434. Springer.

Yu, G., Pachon, J., Thengvall, B., Chandler, D., and Wilson, A. (2004). Optimizing pilot planning and training for continental airlines. *Interfaces*, 34(4):253–264.