Improve workload prediction in the field of Cargo Operations

Author:
Bob Bokern

Supervisors:
dr. J. B. P. Mulder
J. C. de Jong
dr. R. Bekker

A thesis submitted in fulfillment of the requirements for the degree of Master of Science Business Analytics

Faculty of Sciences, VU University Amsterdam

Wednesday 16th December, 2015

Preface

This thesis was written to obtain the Master's degree in Business Analytics, at the Faculty of Science of the VU University Amsterdam. The research focusses on the short-term workload forecast at the Schiphol Cargo Hub. Different forecasting techniques have been used to create a historical flight demand forecast as well as a forecast based on actual booking data. The research was commissioned by the Air France KLM Operations Research department in association with the PRI (Planning, Roostering en Indeling) department of KLM Cargo.

I would like to thank several people for supporting me during my entire graduation period. First of all, I would like to thank my supervisors at the VU University: René Bekker and Fetsje Bijma. They have guided me during my graduation, especially on the theoretical part. Next, I want to thank Joke de Jong, my supervisor from the PRI department, who always was there when I had business related questions. In addition, I want to thank her for the great time I had at the KLM Cargo division. Finally, special thanks goes to Jeroen Mulder. Jeroen was my supervisor from the KLM Operations Research department and not only guided me, but also challenged me a lot when I was facing difficulties. His willingness to help and his ability to motivate people has made me rise above myself during this internship.

Bob Bokern, December 2015

Executive Summary

Predicting the expected workload is a common problem for a lot of industries. For KLM Cargo, the question of whether the expected workload can be predicted, shows up regularly. Insight in the expected workload at the day of operation in advance, enables KLM Cargo to adapt and therefore improve the resource planning for that day. As a result, direct costs (i.e. unnecessary paid salaries when having too many resources) and indirect costs (i.e. loss in revenues due to cargo that is not processed in time when having insufficient resources) can be reduced. Therefore the research question answered in this thesis is:

Is it possible to create a reliable forecast of the expected cargo workload (within an error of 10%), split by freight building?

To answer this question is: Yes! This thesis shows that we can create a forecast with an error of 2-3 %. This is within the bound of 10% stated by the business. The forecast that has been created, has a forecast horizon of 10 days ahead and is based on two different data sources: post departure (historical) *flight* data and pre departure *booking data*. The flight data, obtained from the Altea booking system, contains the total amount of transported cargo per flight and per departure date. The booking data, obtained from the AFLS cargo system, contains all information related to a booking. A booking can be made only 10 days in advance, the time at which the booking window opens.

To create a forecast on the historical flight data, two different forecasting techniques have been used: Autoregressive Moving Average (ARMA) models and Exponential Smoothing (ES) models. To be able to determine which forecasting model would perform best, single flight forecasts have been made on a period from the past and a Weighted Absolute Percentage Error (WAPE) measure has been used to measure the forecast accuracy from those historical forecasts. At the end, those single flight forecasts are aggregated to create a forecast on total day and freight building level. From the final results, we concluded that this way of forecasting gives accurate forecasts (with WAPE values below the 5% for the forecasts on freight building level). Unfortunately, because of missing European flights in the flight data that has been used, the forecast for freight building 1 could not be made. Furthermore, the booking and flight forecast can not directly compared with each other, because there is a difference between the total weight of transported cargo. This discrepancy between the two data sources should be investigated.

The booking forecasts have been created using linear regression models. Two different forecasting approaches have been examined: a *direct* approach and an *indirect* forecasting approach. Because there were no clear business rules for assigning bookings to a specific freight building, only a forecast split by incoming and outgoing flights has been made.

The direct approach predicts the expected amount of cargo directly from the sum of booked weights at a certain day before departure. First results, based on forecasts made on a historical period, showed that within 2 days before departure a forecast of the expected workload can be made with a WAPE value of below the 5%.

The indirect approach consists of two steps. First, an estimation is made of how the booking that are made up to a certain day before departure will materialize at the day of operations. Next, this materialization value is used to predict the total amount of cargo (weight of bookings that are made plus the weight of the bookings still to be made) that will arrive at the day of operations. The results showed that overall performance of this indirect approach performs slightly worse compared to the direct approach.

During this internship we encountered several issues. To cope with those issues, recommendations were made for further research (see Chapter 9). The recommendations were grouped into different categories: *business*, *model* and *data*.

Contents

Chapter 1

Introduction

This research was commissioned by the Operations Research (OR) department of Air France KLM (AFKL) in association with the PRI department of KLM Cargo (in Dutch: Planning, Roostering en Indeling). Looking at the subject of this master thesis, the research focusses on the AFKL Cargo business at Amsterdam Schiphol Airport. In order to meet the vision of AFKL - to be more efficient in terms of increasing their profitability and productivity (by 4% annually) in the coming 5 years - KLM Cargo wants to revise their resource planning.

Currently, KLM Cargo has the feeling that at the day of operations it is often the case that there are too many or too few resources available to process the workload. Too many resources implies unnecessary direct costs (unnecessary paid salaries), whereas a lack of resources will lead to delays (i.e. cargo which is not processed on time) and late deliveries of shipments, so indirect costs and low service levels. To reduce these costs, KLM Cargo wants to get more insight in the amount of cargo that has to be processed at the day of operations. To measure the service level, the Flight as Planned (FAP) index has been introduced. This KPI measures the percentage of cargo packages that is on time for their original flight plan. The target for the FAP is 86.3% and for 2015 so far, a FAP has been realized of 86.4% (so just above target).

The aim of the research is to have a forecast of the expected amount of cargo that has to be processed by Amsterdam Schiphol Airport at the day of operations. Rationale for this is that a better indication of the expected workload will lead to a reduction in direct costs (unused resources) without increasing indirect costs (late delivery and low service levels). Therefore, a 10 days ahead forecast will be created to estimate the expected workload that has to be handled during the day of operations. Because of the huge amount of data (several gigabytes), the forecasting code has been developed in the programming language Python in combination with the usage of Spark (open source computing framework for Big Data computing). The code itself is running on a Hadoop cluster.

To be able to measure the impact of the final solution, resource planners will use the forecast in practice to see how it performs with respect to the current way of planning resources. This validation process will be done after this thesis is finished. The reason for this approach of benchmarking is because in practice it seems to be hard to quantify over/under capacity. This can be due to the fact that people will always work on something, whether there is too much or not enough work to do.

The outline of the report has been set up in a way that it presents the research that has been done step by step. First, in the *Process description* section (see Chapter 2), a general introduction will be given of Air France KLM Cargo and KLM Cargo in particular. Furthermore, an overview of the current processes at KLM Cargo will be outlined. Next, the problem that was given by KLM will be described in the *Problem description* (see Chapter 3).The *Literature review* section (see Chapter 4) will focus on literature study on the subject of this research. This includes, among others, insight in previous research done by former KLM interns and work that has been done by our Operational Research colleagues from Air France. The final design of the solution that helps to solve the problem can be found in the *Model* section (see Chapter 5). In the *Analysis* section (see Chapter 6), a description of the data is given and some comprehensive analysis of this data. In addition to the report, also the implementation code of this solution will be handed over to the business. To get insight into the performance of the designed solution, results of the model and resulting conclusions from those results will be shown in the sections *Results* (see Chapter 7) and *Conclusions* (see Chapter 8), respectively. Finally, some discussion and recommendation for further research will be given in the section *Discussion and Recommendations* (see Chapter 9).

Chapter 2

Process description

In this chapter a brief insight into the Air Cargo market will be given, together with some background information of Air France KLM and the Cargo division of AFKL including their operating hubs. Furthermore, the current processes of handling cargo and estimating the expected demand will be described. The descriptions of the processes are based on personal interviews with employees of the KLM Cargo department. This chapter is primarily intended for people who have no knowledge about the basic Cargo processes.

2.1 Air Cargo Market

Next to the passenger market there is second market in aviation: the air cargo market. This market comprises all the transport of air cargo goods carried in an aircraft. Considering the different type of air cargo goods, three categories can be distinguished: air freight, air express (specialties) and airmail.

2.1.1 Historical overview

The history of the air cargo market shows that air cargo is a growing business. Until 2008, the year that the global economic crisis started, the total market was more or less steadily growing in terms of freight tonne-kilometers and volume. Due to this economic crisis the market collapsed at that time, but finally recovered in 2010 with a growth of more than 19% over the depressing results of 2009. Since 2010 the air cargo market stagnated till early 2013 (see Figure 2.1). The prolonged period of weak growth can be attributed to two factors: a weak world economy and slack trade growth. On a positive side, world air cargo began to grow again in the second quarter of 2013. By July 2014, traffic had grown 4.4% compared to the first seven months of 2013. Forecasts for even better economic and trade growth should lead to sustained air cargo traffic growth in 2015 and 2016 (Boeing, 2014).

Figure 2.2 zooms in on the Air France KLM cargo market. Just as in the previous figure, a drop is observed at the end of 2008. However, when the global market started to recover, the AFKL cargo market remained (more or less) constant. It took until the second quarter of 2013 until the market recovered to the level of the beginning of 2008. Though, over the past two years, we can see a downward trend, which we do not observe considering the world cargo market. This negative trend represents the tough times for KLM Cargo nowadays.

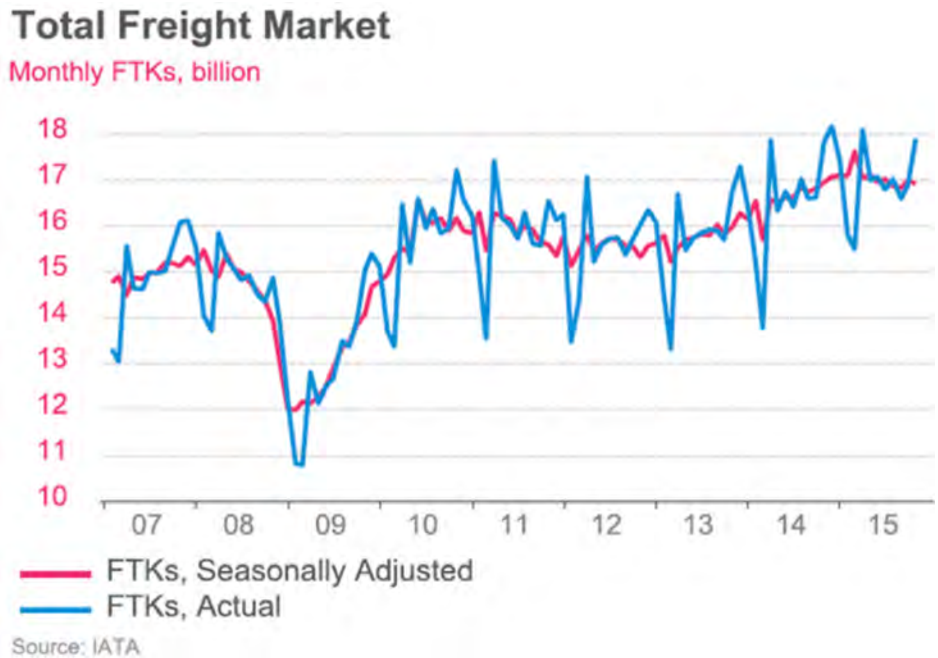

Figure 2.1: Freight tonne-kilometers over time (IATA, 2015)

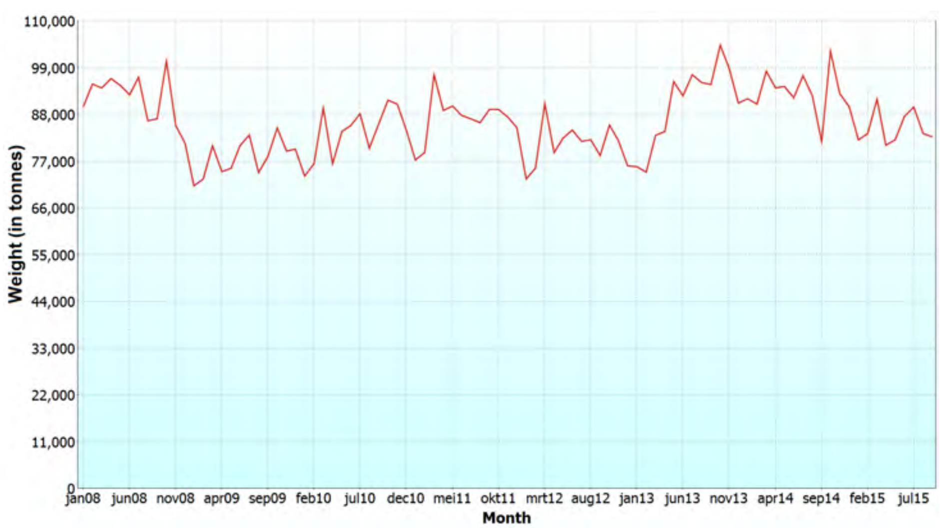

Figure 2.2: Transported cargo weight by Air France KLM over time (Versnel et al., 2015)

2.1.2 Cargo booking process

When transporting a package of cargo by air, a reservation of space must be made for this specific flight. This is called a booking. Typically there are two kinds of bookings:

- Bookings on allotment
- Free sales (spot) bookings

A booking on allotment is a continuous reservation for a certain amount of space at one or more flight-date combinations. Free sales (spot) bookings are almost the same as allotment bookings, except that they only apply to a single flight-date combination. Furthermore, allotment bookings are made longer in advance than free sales bookings. The reason for that is that free sales bookings can only be made after the *booking window* of a flight has started, in contrast to allotment bookings. The booking window of a flight is the time interval in which bookings can be made for a specific flight. Roughly speaking, the start of the booking window is more or less 10 days before departure.

Although cargo is booked on a flight in advance, customers do not always stick to the booking made. In other words, the amount of booked cargo is not always equal to the amount of cargo that will show up (discussed in more detail in Section 4.2.2). This phenomenon makes workload estimation (more) difficult for airline companies. To prevent that too much/not enough space is available in the aircraft for all bookings, airlines steer on their cargo capacity (Popescu et al., 2006). This strategy of steering on the capacity is often based on historical data, which gives insight in the show up rate of the booked amount of cargo.

2.1.3 Cargo transportation process

There is a general pattern in the way how cargo is transported from origin to destination. Often, the flow of a shipment starts with an agent who wants to transport a package of cargo. Most of the time, the package will be picked up by a cargo carrier and loaded into a truck. The truck will then transport the package to the airport where the package will be loaded into an aircraft for air transportation. When the package arrives at its destination airport, the package is offloaded from the aircraft and again loaded into a truck of a cargo carrier who finally delivers the package to its final destination. The whole flow as described above can be presented graphically (see Figure 2.3).

Figure 2.3: General cargo flow

2.2 Air France KLM Cargo

In May 2004, Koninklijke Luchtvaart Maatschappij (KLM) joined forces with Air France (AF) and together they became the largest European airline group. The combined network of both airlines ensured that nowadays Air France KLM carries more than 77 million passengers per year to more than 240 destinations in 103 countries. Although both airlines are one group now, each of the airlines retained their own identity, trade name and brand. Looking at the joined AFKL group, there are three core businesses that can be distinguished (AFKL, 2015a):

- Passenger Business
- Engineering & Maintenance
- Cargo

Looking at the Cargo business, Air France KLM Cargo is part of the Air France-KLM-Martinair Cargo group. This is the dedicated air cargo business of the AIR FRANCE KLM Group, generating a combined turnover of three billion euros yearly.

Historically, Air France Cargo, KLM Cargo and Martinair Cargo have always been at the forefront of transporting and handling a variety of cargo. Building on this experience, Air France-KLM-Martinair Cargo offers a wide range of air transport services in the market, providing seamless connections throughout the world, with more than 250 destinations in 116 countries. Looking at the mission of Air France-KLM-Martinair Cargo, the main focus is on becoming customers' preferred carrier.

2.2.1 Dual Hubs

Although KLM and Air France are part of the same group, both airlines run their own operations at their own mains hub: Amsterdam-Schiphol (KLM) and Paris-Charles de Gaulle (Air France). The two hubs function in tandem, and by combining both networks, Air France KLM is better able to meet the desires of the customers. The main cargo activities at Paris-Charles de Gaulle take place in Gare n1 Extra Large (G1XL). The G1XL airfreight terminal is one of a kind in terms of its high-tech facilities and flexibility in serving customers. It forms the beating heart of the company's global cargo network. All the cargo activities at Amsterdam Schiphol Airport take place at the Schiphol Cargo Hub. The hub consists of three freight buildings which offers numerous facilities for handling the customers' special products. These facilities include an Animal Hotel for 'live' animals, a Conditioning Competence Centre for conditioning and cooling of shipments and a vault for valuables (AFKL, 2015b).

2.2.2 Air France KLM fleet

To be able to realize the mission of Air France-KLM-Martinair Cargo, the process of handling and transporting of cargo has to be well designed. Therefore, Air France-KLM-Martinair uses different types of aircrafts and trucks. First of all, there is the regular Passenger aircraft PAX, where the lower deck (so called 'belly') is filled with freight and the main deck is filled with passengers. Then there is the combined aircraft, where the main deck is filled with passengers as well as with freight (in the back) and also the lower deck is filled with freight. Finally, KLM has a few full-freighter aircrafts that are totally filled with freight (main and lower deck). In Figure 2.4 (Wever, 2011) a representation of each of the different aircraft types is shown. The Air France KLM fleet itself is described in more detail in Table 2.1.

In contrast with the passenger and combined aircrafts, only the management of the full-freighter is in control of the KLM Cargo department. The management of the PAX aircrafts and the combined aircrafts are in the hands of the KLM Passage department. This means that the Cargo Network department, except for the network of the full-freighters, fully depends on the decisions made on Passenger Network side. Looking at the network design of carriers that only transport cargo, this is a big disadvantage. Such carriers use an opportunity driven approach to design their network. This means that those carriers enable themselves to

Figure 2.4: Comparison of the different aircraft types

Aircraft Type	Aircraft Type Number	Cargo Capacity (Tonnes)	# KLM Aircrafts	# Air France Aircrafts
Full-Freighter	Boeing 747 - 400	112	3	-
	Boeing 777 - 200	102	-	2
Passenger	Airbus 318 - 100	1	-	18
	Airbus 319 - 100	1 - 2	-	38
	Airbus 320 - 200	1 - 2	-	45
	Airbus 321 - 100	1 - 2	-	5
	Airbus 321 - 200	1 - 2	-	16
	Airbus 330 - 200	10.9	12	15
	Airbus 330 - 300	17.4	5	-
	Airbus 340 - 300	13.9	-	13
	Airbus 380 - 800	10 - 12	-	10
	Boeing 737 - 700	1 - 2	18	-
	Boeing 737 - 800	1 - 2	25	-
	Boeing 737 - 900	1 - 2	5	-
	Boeing 747 - 400	12.5	5	4
	Boeing 777 - 200	12.7	15	25
	Boeing 777 - 300	23.7	10	39
	Embraer 190	1	28	10
	Fokker 70	1	19	-
Combined	Boeing 747 - 400	35	17	-

Table 2.1: Air France-KLM fleet (AFKL, 2015b), (Spotters, 2015)

adapt their cargo network to recent developments (for example when a new potential market arises) in the air cargo market to increase their revenues.

2.3 KLM Cargo

KLM Cargo manages everything that is related to the air cargo transported via Amsterdam Schiphol Airport. Besides the actual transport of cargo, also all cargo handling processes are in control of KLM. For instance, this involves loading and unloading of cargo and the preparation of cargo for further transportation. To ensure that all cargo processes run smoothly, a good understanding of the general flow of cargo and all the sub processes (handle, prepare and transport of cargo) is necessary.

To help to smooth the process of handling and transporting cargo, KLM Cargo makes a distinction between several 'flows' of cargo. Roughly speaking, those different flows are characterized by the ingoing and outgoing transport mode (aircraft or truck). The idea behind this distinction is that each of the different flows has its own specific cargo handling process. Therefore, at the Amsterdam Schiphol Cargo Hub there are different freight buildings which all take care of the handling process of a certain flow of cargo. This will be described in more detail in Section 2.3.1.

2.3.1 Schiphol Cargo Hub

The cargo hub itself consists of three buildings: freight buildings 1, 2 and 3. As mentioned in the previous section, each freight building handles its own flow of cargo. Those different 'types' of cargo flows are shown in Figure 2.5. In this figure, also the relation between the different freight buildings is included (based on (d'Engelbronner, 2012)).

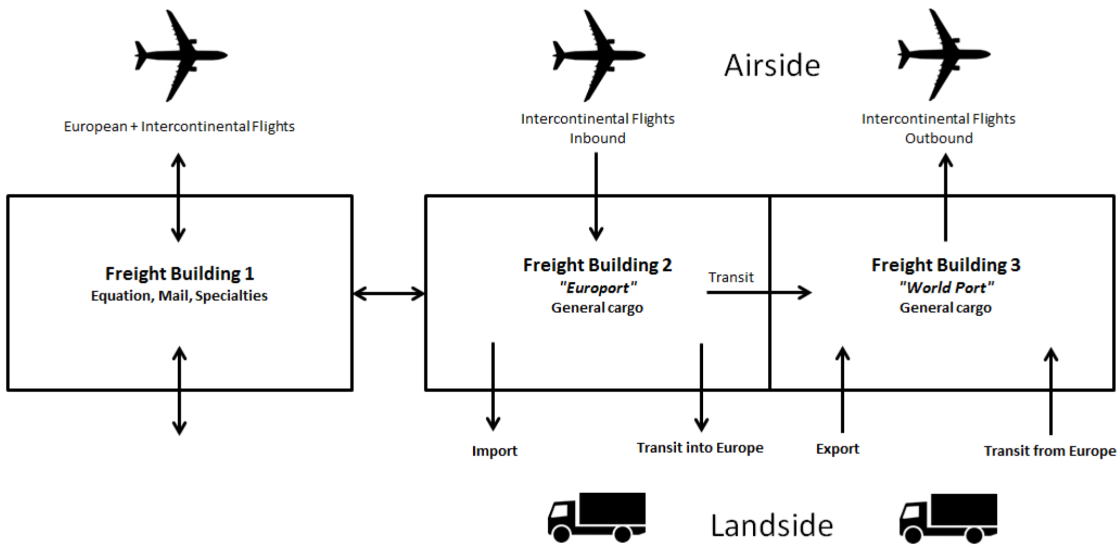

Figure 2.5: Overview of the Schiphol Cargo Hub

Freight building 1, commenced in 1967, takes care of all the cargo that belongs to one of the following categories: Equation, Mail and Specialties. These categories of cargo include mail, small packages (Equation), special shipments (for example urgent shipments), pharmacy and 'live' animals. The cargo can be delivered by trucks and aircrafts (both ICA and EUR).

Freight building 2, also known as Europort, became operational in 1982 and is responsible for all incoming (general) cargo that arrives by aircraft. When the cargo is offloaded from the inbound flight, it has to be processed for further transportation by truck (transit into Europe) or aircraft (transit to intercontinental destinations). If the cargo is at its final destination (import), it can be picked up by the owner.

In 1991 the last of the three freight buildings, also known as World Port, opened its doors. This building handles all of the (general) cargo that is destined for the (outbound) ICA flights. This cargo can be truck transit cargo from Europe, air transit cargo from intercontinental flights or export from the Netherlands itself.

Some more details about the freight buildings are summarized in Table 2.2. In this table, the number of resources only represents the employees who are actually handling the air cargo (so only the 'executive' employees).

Freight Building	**Floor area (m^3)**	**Resources (executive employees)**
1	34.000	*unknown*
2 ('Europort')	32.000	KLM: 100 Flex: 25
3 ('World Port')	34.000	KLM: 180 Flex: 55

Table 2.2: Key figures freight buildings (de Jong, 2015)

2.4 Cargo handling process

When the cargo arrives at an airport, it often is not at its final destination. Before the cargo can be transported further it has to be processed. The reason for this is that cargo destined for multiple destinations is packed together on a transportation device. One of the transportation devices that is used is the Unit Load Device (ULD), see Section 2.4.1. Because nowadays the workload of freight building 2 has been determined in terms of ULDs, this transportation device has been described in more detail in Section 2.4.1.

The process of preparing cargo for further transportation is very common. First of all, the cargo has to be offloaded from the aircraft or truck. After it is offloaded, the cargo is transported to a place at the airport at which the cargo will be processed. This preparation process consists of two phases: a break down and build up phase. During the break down phase, cargo destined for multiple destinations, but that has been stacked onto the same transportation device, will be separated from each other. Next, during the build up phase, the cargo is collected per destination and built up on a transportation device again. Finally, all the cargo is transported to the right truck or aircraft for further transportation towards its final destination.

2.4.1 Unit Load Device (ULD)

A *unit load device (ULD)* is a container or pallet for cargo transportation. Three types of ULDs can be distinguished. A *through unit load device (T-ULD)* contains cargo that has the same destination. For this kind of ULD there is no need to do additional processing at the freight building. More often arriving ULDs contain cargo for multiple destinations. In that case, some of the shipments have to be loaded into aircraft A and some into aircraft B. This kind of ULD is called a *mixed unit load device (M-ULD)*. When a ULD is considered to be a M-ULD, it needs some additional processing to make sure all the cargo will be loaded into the right aircraft/truck. Third, there is a *bijbouwer unit load device (BB ULD)*, which is actually a variation on the T-ULD. In fact, a BB ULD is a T-ULD that does have space left for some more shipments with the same destination. This kind of ULDs can be convenient when some space has to be saved, for example when some bookings show up with more cargo than their booked amount.

2.4.2 Cargo handling process at the Schiphol Cargo Hub

The general process of handling cargo at the Schiphol Cargo Hub is independent of the freight building where it has to be processed. However, as mentioned before, freight building 1 takes care of the 'special' cargo, where freight buildings 2 and 3 are taking care of the more 'general' cargo. The cargo that is handled by freight building 1, often needs some extra care during the transport and is therefore packed differently than general cargo. There are multiple reasons why some cargo is transported differently from others, such as that the cargo must retain at a certain temperature (e.g. donor organs, pharmaceuticals) or that it has to be transported extra safely (e.g. dangerous goods, chemicals).

When the cargo arrives at Schiphol, different steps have to be taken before the cargo is ready for further transportation. As stated above freight building 1 has a slightly different process in comparison with the other two freight buildings. To take care of the 'special' cargo that is handled at this freight building, freight building 1 is equipped with special equipment like climate rooms and animal shelter. Also qualified staff is available to make sure that all cargo is handled in the right way. Those people are important to have, because each category of cargo that is part of the 'special' cargo has its own procedure to follow. That is also the reason why freight building 1 has its own logistics, because their processes are significantly different from the general processes that take place at the freight buildings 2 and 3.

When the cargo arrives at the freight building, first the cargo will be split according to the type of ULD. All T-ULDs are checked on different aspects: correctness of the ULD construction, the weight and height of the ULD compared to the specified values and the correctness of the pallet-tag. When one or more of those aspects does not seem to be correct, the ULD will be transported to the 'pallet buildup' process area. Furthermore, the status of the ULD will be changed to BB ULD. Otherwise, when everything seems to be correct, the T-ULD will be immediately stored into the PCHS (see section 2.4.3).

For the M-ULDs another process takes place. Those ULDs are send to the 'pallet breakdown' process area. There the M-ULD will be broken down and all the shipments will be identified for further processing. Which ULD will be processed at what moment depends on its deadline (the departure time); the closer to the deadline, the earlier the ULD will be processed. After it has been broken down, the corresponding shipments are loaded onto a 'new' ULD at the 'pallet buildup' process area. Then it is stored in the PCHS until the truck/aircraft is ready for departure.

2.4.3 Pallet Container Handling System (PCHS)

When ULDs do not have to be processed immediately after arrival, they will be stored in the Pallet Container Handling System (PCHS). This is a handling system for the fully automated handling of pallets and containers, from delivery to loading and offers 1,200 storage positions (AFKL, 2015b). From there, these ULDs are placed in a queue for further transportation.

The moment at which the ULD will be transported to the workplace is determined by the PCHS itself and is based on a user setting x, denoting the number of hours until the cargo has to be ready for further transportation, that is deadline oriented. When a ULD has to be ready for transportation to the aircraft or truck within this x hours, the PCHS will automatically bring this ULD to the work floor. To make sure that cargo is indeed on time for further transportation, KLM Cargo holds on to a threshold of two hours before departure at which all the cargo for a specific flight/truck has to be ready. Of course, when multiple ULDs have to leave within this time interval of x hours, it is not always the case that there is enough space available on the work floor to store all the ULDs. Therefore, the PCHS makes an ordered list of ULDs according to the Earliest Deadline First (EDF) principle. The deadline of a ULD is specified as the deadline of the colli with the earliest deadline that is on that ULD. When there is no ULD to process, the deadline setting in the PCHS system will be manually adjusted to a larger time interval to guarantee a continuous work flow at the work floor.

2.5 Workload estimation process

To be certain that all the cargo that arrives at Schiphol Cargo Hub is ready on time for further transportation, KLM Cargo has to make sure that there are resources available to handle the expected arriving cargo. The department of Planning, Roostering en Indeling (PRI) creates a resource planning long in advance to cover the expected workload.

2.5.1 Planning, Rostering and Allocation process

According to the PRI department a proper resource planning tries to find a balance between customer concerns, business concerns and the concerns of the employees. Moreover, the PRI department emphasizes that especially in service providing organizations with many employees, like KLM Cargo, a good resource planning is a (key) success factor. Therefore, the PRI process has been developed which deals with all the requirements that has to be met during the process of planning, rostering and the allocation of the resources (such as taking into account collective labor agreements, personal preferences, etcetera).

To come to resource planning that meets all those requirements, the PRI process consists of five steps: Forecasting, planning, rostering, allocation, evaluating/correcting (de Jong, 2014). Each step has been performed at a certain moment in time. After the last step, the process starts over again. A graphical representation of the PRI process is shown in Figure 2.6. This research will focus on the *allocation* phase.

In the first step, a long term *forecast* of the expected workload is made. This forecast can actually be divided in two parts: a strategic and a tactical forecast. The strategic forecast is a long-term forecast (made 1-5 years in advance) intended for getting a high level indication of how your workload will change over a long period. This forecast is based on looking at possible relationships between industrial trends/development of macro-economic variables and the workload. The tactical forecast is a more short-term forecast (made <1 year in advance) and is meant to get insight into the expected production for the coming period.

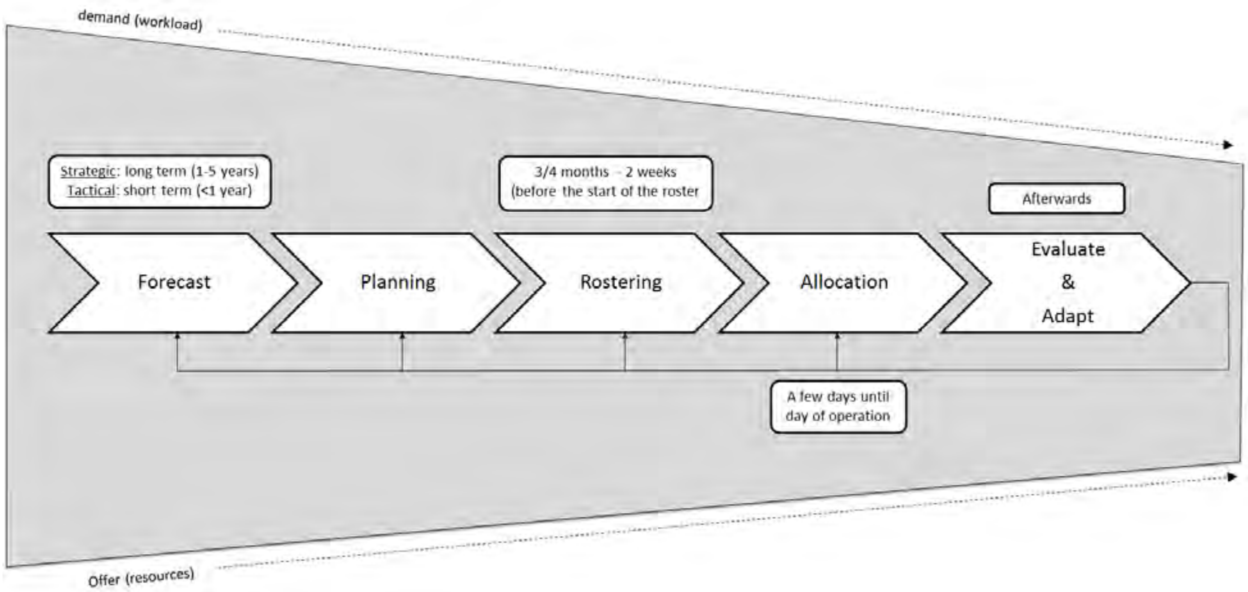

Figure 2.6: Planning, Roostering en Indeling (PRI) process

For the second step, i.e. *planning*, the occupancy of resources are determined. For this, the expected amount of work, coverage choices (over/under capacity, quality, business risk), the new (Cargo) timetable and the percentage of attendance are taken into account.

After the planning phase, the basic and personal *rosters* are created (= roster planning). This phase will take place 3-4 months before the start date of the new roster. Both rosters are based on the occupancy, rostering rules and comply with the collective labor agreement requirements. The basic roster comprises the default number of resources needed on a specific work location for each work shift for each weekday. The personal rosters concern the actual allocation of individuals to work locations at certain shifts through the week, such that the basic roster is filled. In general, there are three kinds of work shifts, as shown in Table 2.3:

Shift	**Start**	**End**	**Time interval**
Morning	06:00	14:30	06:00 - 14:30
Evening	14:00	22:30	14:00 - 22:30
Night	22:00	06:30	22:00 - 06:30

Table 2.3: Work shifts

When all previous phases are completed, the *allocation* phase takes place. Although, this process of allocation differs per freight building, a general outline of the process will now be given. This general description is based on interviews with KLM Cargo resource/roster planners. Furthermore, this phase of the PRI process is the part where the research focuses on, since a better workload estimation should lead to a better allocation of resources. To avoid that there are too many/few resources for the expected amount of work, there is an alignment of staff. When the day of operation approaches, adjustments to the high-level planning are made by the resource planner (only for freight building 3) and/or the head operations flow (HOF) managers. A meeting between the HOFs and the resource planner has been set up each morning to check the status of the resource planning. During these meetings, adaptations to the resource planning are made based on new information such as illness of employees and changes in the expectation of the amount of work that will arrive (provided by RM, see Section 2.5.2). In case of insufficient resources, extra flexworkers are hired to handle the expected amount of work. In case of overcapacity there are multiple options, such as stimulate leave or assign personnel to other work locations (Veldman, 2015). The process of allocation takes place on the day of operation itself. At this day, the actual assignment of resources to the different activities (unloading trucks, build up ULDs, break down ULDs,etc.) is done.

The fifth step is that the whole PRI process is been evaluated by looking at all the phases above and see how the results of those phases correspond with reality. If it turns out that some steps of the process do not seem to perform well, necessary adaptions are made.

2.5.2 Revenue Management estimation

When the day of operation is approaching, more and more information will be available that gives an indication of the amount of work that can be expected for the day of operation itself. To make it possible for the resource planner and HOF to adapt their resource planning to the most recent updates of the expected amount of work, the Revenue Management (RM) department of KLM Cargo provides them with the current booking information that is available. However, in general most of the bookings will arrive only a few days (or even a few hours!) before the day of operation; thus solely looking at the bookings that are already have been made will not give a proper indication of the real amount of work. Therefore, a comparison between the current bookings and the booking for a similar day in the past is made, to see how the current booking behavior for this day is related to the behavior in the past. Based on this comparison, a rough estimation can be made whether or not the booking behavior is running behind or ahead compared to the booking behavior as observed in the past. Looking at the aim of this research, the forecast made should make this rough estimation more accurate.

Chapter 3

Problem description

3.1 Background

In order to be more efficient, Air France KLM (AFKL) wants to increase their productivity. High level the objective of AFKL is to increase their productivity with 4% yearly for the coming five years. As a result, the KLM Cargo division is looking for quick wins that can help to achieve this objective. Within the Cargo division, the Cargo Operations department is responsible for the handling and further transportation of arriving cargo at Schiphol airport. To handle all the arriving cargo, KLM employees and flex workers work together in teams to make sure that all the cargo is processed in time. The ratio of KLM employees and flexworkers depends on the number of KLM employees that are available during a certain day (on average 20% flex (de Jong, 2015)). When there are too few KLM employees available to handle the estimated workload, extra flex workers will be hired. More detailed background information can be found in the previous chapter (Chapter 2)

Currently, within the department of PRI (Planning, Rostering and Allocation), which creates the rosters for Cargo Operations employees, there is a feeling that it happens too often that too many/not enough resources are scheduled on the day of operations with respect to the workload. Obviously, overcapacity of resources is associated with direct costs (unnecessary paid salaries). The other way around, when there are not enough resources available, there could be a loss in revenues due to cargo that is not processed in time and late deliveries of shipments. Looking at the current results of KLM Cargo, around 91.6% of the cargo is delivered as promised and, even a bit lower, around 86.4% of the cargo is flown as planned (just above target: 86.3%). The Cargo department thinks that the objective, of increasing productivity, can partially be achieved by optimizing their resource planning. This planning comprises the number and type of resources (KLM employees and flexworkers) that is needed to handle the expected cargo workload that arrives at a certain day of week during a specific work shift.

Looking at the problem of optimizing the resource planning, the problem can be split in two parts: a *forecasting* and a *planning* part. The forecasting part focuses on what will be the expected amount of cargo that has to be processed at the Schiphol Cargo Hub on the day of operations. The planning part will focus on this estimated workload with respect to intermediate process steps (what cargo will be processed at what time). Because no accurate information about the expected workload is available, this research will focus on the forecasting part. Rationale for this is that a more accurate forecast will lead to a reduction in direct costs (unused resources) without increasing indirect costs (late delivery and low service levels). With this more accurate estimation of the expected cargo traffic, the KLM Cargo department is able to estimate more accurately the resources needed (split into KLM employees and flexworkers).

Although the 3 freight buildings are working together as one, they are all managed individually. Therefore an estimation of the workload has to be made for each of them, based on the expected cargo that will be handled by that specific freight building. The department of PRI furthermore demanded for an estimation error of less than 10%. This, results in the following research question:

Is it possible to create a reliable forecast of the expected cargo workload (within an error of 10%), split by freight building?

To accomplish this objective, a 10 days ahead forecast will be made based on different data sources (related to Section 2.5.1): a short-term forecast based on historical flight data and a very short-term forecast based on pre-departure booking data.

First of all, a forecast for the *short term* based on historical flight data has been created. This data is post departure data and therefore very accurate. The forecast will be created with the use of different forecasting techniques and are chosen based on previous studies. Rationale for this is that the flows have their own specific characteristics (for example more random pattern versus a high seasonal pattern) and therefore the specific techniques could improve the accuracy when predicting different flows.

Secondly, a forecast for the *very short term* has been created based on booking data. This forecast can be seen as a more tactical forecast and an expansion on the flight forecast. The booking data that is used, is pre departure data and therefore can still be modified. This is due to booking updates and cancellations. A reason to also create a forecast based on booking data is that not only flows show different behavior, but also on a more detailed level (such as product level or agent level there is different behavior observable when looking at booking characteristics.

3.2 Deliverables

The final goal of KLM Cargo is to get insight in the amount of cargo that has to be processed at a specific day and shift. To accomplish this, a forecast is needed on flight and booking level. With these forecasts the Cargo Operations department should be able to better estimate the number of resources needed to cover the workload. At the end of the research an implementation in Python of the forecasting techniques has been given. Next steps, out of the scope of this research, will be to design a tool, where insight is provided in the expected amount of arriving/departing cargo per time shift (or time interval).

Chapter 4

Preliminary forecasting studies

In this chapter, three studies in the field of cargo demand forecasting will be discussed. The outcomes of those studies have been a first inspiration for the solution approach to the problem.

In comparison with the Air Passenger business, where much more research has been done on the subject of predicting expected number of passengers, there is very little literature available on workload estimation for Cargo business. Although it seems that the problem is the same for both of the businesses, there are some important differences. One important difference is the uncertainty of Air Cargo supply. The volatility of the air cargo market is very high. This means that the air cargo demand is highly sensitive for changes in external factors, such as bad weather conditions (which can result in bad harvest) and consumer trust. Compared to the number of seats on a flight, the available cargo capacity of a flight is uncertain and cannot be determined in advance (Huang and Hsu, 2005). This is, among others, due to the fact that for AFKL cargo business is most of the time subordinate (except for the full freighter aircrafts) to the passenger business (as mentioned before in Section 2.2.2). Therefore, the available cargo capacity depends on the (expected) number of sold seat tickets. Furthermore, the expected workload for cargo business is two dimensional; as well weight as volume has to be taken into account. In addition to this multi-dimensionality property, also the cargo booking process is an extra difficulty because of the uncertain actual amount of cargo that will show up for transportation (see Section 2.1.2).

Looking at the problem of predicting the expected workload, it seems that this is a common problem for a lot of industries. For KLM Cargo, the question of how the expected workload can be predicted more accurately shows up regularly. To tackle the problem of estimating the amount of cargo that will arrive at the Amsterdam Schiphol Cargo Hub, some previous studies have been carried out. The conclusions of those studies are used as a starting point. The methods that will be used for the rest of this research will be described from a theoretical point of view in *Model* (see Chapter 5).

Three studies, all related to forecasting cargo demand, have been investigated. Regarding the forecast level of each study, the three studies can be subdivided into two groups: forecast on booking level ((d'Engelbronner, 2012) and (Benfatma and Garcia, 2014)) versus flight level ((Garcia, 2014)).

4.1 Forecast on flight level

Until 2014, Air France did not have any kind of forecast to support the Revenue Management Cargo department on the monitoring of flights. Therefore, they wanted to have a forecast of the daily demand at flight level. At the AFKL Operations Research department, multiple demand forecasting tools were developed.

A long-term (12 months ahead) forecast was created based on monthly air freight data by Versnel. Multiple forecast methods were tested:

- Autoregressive Moving Average (ARMA) model: single time series model
- Vector Autoregressive Model (VAR) model: a forecasting model which takes into account multiple time series (competitor effects)
- Factor-Augmented Vector Autoregressive Model (FAVAR) model: forecasting model which takes into account both competitor effects and effects related to economic variables

To determine which method performs best for different flows, all methods were tested on the last 12 months and the best method was selected to forecast. The results of this forecasting study have been implemented in an internal KLM tool called CADEAU (Versnel et al., 2015). Although the results of this study were very promising, the outcomes have not been used for this research because of the difference in forecast horizons.

In addition to the long-term forecast, Garcia created a short-term (14 days ahead) forecast of this daily demand and volume at flight level. Different forecasting models were compared to find the best model:

- Exponential Smoothing (Holt-Winters) model
- Autoregressive Integrated Moving Average (ARIMA) model
- Robust Calendar model: recent trend adjusted for calendar effects (day of week, special events, freighters)
- Volatile DoW model: recent trend adjusted for day of week effects

First results showed that the Robust Calendar Based model performed best, with an average error rate of 45% (at day level). However, Garcia also observed that the accuracy of the forecast on flight level was affected by capacity changes of a flight due to aircraft type and schedule changes. To be able to take care of this capacity effect and schedule changes, a new forecast was created based on cargo load factor. The cargo load factor represents the percentage of the available cargo space in an aircraft that is used (also RM Cargo steers on load factor). In the end, the forecasted load factor is translated into a forecasted weight and volume by multiplying this forecasted load factor with the capacity of the aircraft. From the outcomes of this new forecasting approach, Garcia concluded that an exponential smoothing model performed best (Garcia, 2014). Looking at the accuracy of the forecast, the average error rate at day level dropped by almost 20% in comparison with the first results to 26% on average. Furthermore, almost 40% of the forecasted flows had an average error rate lower than 15% after.

Since, during his forecasting study, Garcia concluded that exponential smoothing models performed best, those models have been used for this research. Furthermore, from the long-term demand forecast, the ARMA model has been selected because this was the only single time series model (and for this research we do not take into account relations between multiple time series or exogenous effects).

4.2 Forecast on booking level

Creating a forecast on booking level is interesting to examine. The main reason for this, is that booking data contains detailed pre-departure information about the cargo that belongs to that booking. The booking data comprises all the characteristics of bookings that causes workload at the airport terminals such as: weight, volume, number of packages, ULD type, etcetera. In comparison to the historical flight data, which does not contain such detailed information about the transported cargo, those characteristics could improve the accuracy of the forecast.

4.2.1 Aggregated booking forecast

d'Engelbronner did research on creating a forecast based on booking data. He showed that forecasting the total final weight of a single booking and even of all bookings per single flight is inaccurate. Therefore, he proposed to make a forecast on day level by aggregating all the bookings per day. When looking at the different types of forecasts (as mentioned in Figure 2.6), this type of forecast can be seen as a short-term forecast. Based on current bookings made, an estimation of the expected materialization of those bookings can be determined. To improve the accuracy of the forecast, different forecasting techniques were examined by d'Engelbronner to forecast the total weight of all bookings:

- Time series models: predict total weight based on historical manifest data
  - Moving Average
  - Simple exponential smoothing
- Regression models: predict total weight based on booking levels per Days Before Departure (DBD)
  - Simple Regression
  - Multiple Regression
- Hybrid models: combine Time Series and Regression models
  - Classical pick-up model
  - Advanced pick-up model

Next, d'Engelbronner did an analysis on the forecast performance to see which technique performs best. We compared the overall forecasting accuracy of the different techniques regarding the forecasts errors made at all DBD in the booking window. From the outcome of this analysis, d'EngelBronner showed that the overall error rate (looking at the Mean Absolute Percentage Error (MAPE)) was the lowest for the simple regression model. Hence, the conclusion was that creating a simple regression model per DBD, looking at the aggregated booking and manifest (observed values) data, is the most adequate for providing accurate and reliable predictions of air cargo demand per day.

4.2.2 Single booking forecast

In the aviation industry, airlines steer on their capacity. This strategy is based on the assumption that bookings arrive with an amount of cargo that differs from the booked value. Looking at the realization of a booking k, B_k, and its booked weight w_k and volume v_k, there are four different scenarios possible for the final weight wf_k and volume vf_k:

- No-show: wf_k and vf_k are equal to 0. This can be due to cancellations or a booking that does not show-up on time for the flight departure.
- Low-show: $0 < wf_k < w_k$ or $0 < vf_k < v_k$. The actual booking has less weight/volume than booked.
- Show: $wf_k = w_k$, $vf_k = v_k$.
- High-show: $w_k < wf_k$ or $v_k < vf_k$. The booking shows up with more weight/volume than booked.

Because there is no direct information available in advance about how a booking will materialize, steering on flight capacities is often very hard. Therefore, French colleagues (Benfatma and Garcia) from the OR department at Air France, developed a method that forecasts the *materialization* of a single booking: Opportunity Assessment (OA).

To be able to do proper steering on the capacity of an aircraft, KLM Cargo wants to know what will be the amount of freight, intended for that specific aircraft, expected to arrive at the day of departure. In fact, a prediction is needed of the show up rate of the booked cargo. This show up rate gives an indication about the percentage of the booked weight/volume that shows up when the booking arrives at the warehouse. The assumption behind forecasting such a show up rate is that a booking exhibits some characteristics which tells something about how the booking arrives at the warehouse (Benfatma and Garcia, 2014).

Benfatma and Garcia found out that the *agentURN*, a unique id corresponding to the agent who creates the booking, is an important indicator which provides information about the materialization of a booking. The reason for the importance of this feature could be explained by the particular booking behavior that an agent often has. This booking behavior could help to give a first indication of how a booking will materialize. As an example, consider agent X, who always shows up with a booking which has a higher weight value than the initial booked weight. When this is the case, KLM Cargo could take into account this high-show and better steer on the remaining capacity for the flight corresponding to that booking. To detect such relationships between booking characteristics and the no-show rate of a booking, Benfatma and Garcia developed a decision tree algorithm (trained on 400 days of historical booking data). With the help of this decision tree, the show up rate of all incoming bookings are forecasted. The resulting forecast are calculated by calculating the show up part of the booking and subtracting this from the original booked value. In the end, those materialization forecasts of all bookings will help the RM Cargo Department to better steer on flight capacity.

Chapter 5

Model

In this chapter the forecasting models that have been used for this internship are described from a technical point of view. Furthermore, the solution approach to the problem is presented. This solution consists of two parts: a forecast on *flight* and on *booking* data.

5.1 Forecast on flight data

To create a proper forecast, different forecasting techniques have been used and compared with each other. Furthermore, we looked at forecasts on different levels (booking, flight) and different scopes (forecasts per flight, day). From the experience of (Garcia, 2014) we have chosen to look at exponential smoothing models for forecasting on flight level. In addition, an Autoregressive Moving Average (ARMA) model has been added to the list of forecasting models to investigate. To create a proper forecast with both forecasting methods, the forecasted time series has been *deseasonalised* and made *stationary* upfront (see Section 5.1.2). In the end, the forecasts will be corrected for the flight schedule (based on OAG, see Section 6.1.1). The reason for this is that when a flight is not scheduled for a certain date, no forecast has to be made of the expected amount of cargo that will arrive and therefore the 'forecast' can be set to zero.

When performing the forecast, there are multiple settings that can influence the predicted values for the future. Therefore, multiple combinations of different settings will be used to see what would be the best combination of settings for a specific flight. Such a combination of settings is called a *scenario* and consists of the following parameters:

- Last date: the last date that may occur in the time series to forecast
- Amount of history: the number of months that will be taken into account (looking back from the last date in the time series). By default 3, 6, 9 and 12 months will be considered.
- Initial guess values for parameters: initial guess values (by default 0.3, 0.5 and 0.8) for the exponential smoothing models to start with when searching for optimal parameter settings.
- Forecast error KPI: the key performance indicator which will be used to maximize the accuracy of the forecast when searching for the best parameter settings for the exponential smoothing models. The KPIs that are implemented are MAPE, WAPE, RMSD and SMPE (see Section 5.1.1).

5.1.1 Exponential smoothing

Exponential smoothing is a technique to produce a smoothed time series. The original work was done by Robert Goodell Brown (1956) and expanded by Charles C. Holt (1957). The technique tries to forecast the future by weighting past observations using exponentially decreasing weights. Furthermore, those exponential smoothing models make no assumptions about any correlation between consecutive observations of the time series (Coghlan, 2010). To take into account the different dynamics that might be present in the data, such as trend and seasonality effects, different variants of the exponential smoothing are used. In general, these variants can be divided into three categories: single, double and triple exponential smoothing (Gardner, 1985). Because of the many variants of the exponential smoothing methods, only the ones with additive trend and additive seasonal components are described in this section. In Appendix B a complete list of all standard exponential smoothing methods can be found.

Single Exponential Smoothing

The single exponential smoothing method is a method that is suitable for time series which do not have a trend or seasonal component. This method attaches more weight to more recent observations when those observations have a higher predictive value in comparison with observations from the distant past.

Looking at the formula of the single exponential smoothing method, the m steps ahead forecast can be mathematically formulated as:

$$
\begin{gathered}
S_t = \alpha X_t + (1-\alpha) S_{t-1} \\
\hat{X}_t(m) = S_t \\
0 \leq \alpha \leq 1,
\end{gathered}
$$

where S_t is the smoothed level of the series (estimation of the next value X in the time series), X_t the observed value at time t, $\hat{X}_t(m)$ the m steps ahead forecast and α a smoothing parameter.

The smoothing parameter α controls the rate at which the weight decreases. The general expression for the decrease of the weight at a certain time t can be derived from the formula of the single exponential smoothing method by using successive substitution:

$$
\begin{aligned}
S_t &= \alpha X_t + (1-\alpha) S_{t-1} \\
&= \alpha X_t + (1-\alpha)(\alpha X_{t-1} + (1-\alpha) S_{t-2}) \\
&= \alpha X_t + \alpha(1-\alpha) X_{t-1} + (1-\alpha)^2 (\alpha X_{t-2} + (1-\alpha) S_{t-3}) \\
&= \alpha X_t + \alpha(1-\alpha) X_{t-1} + \alpha(1-\alpha)^2 X_{t-2} + ... + \alpha(1-\alpha)^n X_{t-n} + ... + (1-\alpha)^t X_0 \\
&= \sum_{i=0}^{t-1} \alpha (1-\alpha)^i X_{t-i} + (1-\alpha)^t X_0
\end{aligned}
$$

From this last equation the following formula for the weights for each observation at time t-i can be obtained: α * (1 - α)^i. It is obvious that this is an exponential function (hence the name of the method).

The smoothing parameter α can take a value between 0 and 1, values closer to 1 will attach more weight to more recent observations. When the value of α is closer to 0, more weight is assigned to less recent observations. To show the impact of the value of the smoothing parameter α on the weights for a certain time t, an example is given in Table 5.1 (Hyndman and Athanasopoulos, 2014).

Double Exponential Smoothing

The single exponential smoothing method is suitable for time series without a trend (and seasonal) component. When there is a trend in the data, this trend must be taken into account in the forecast. Therefore, Holt (Kalekar, 2004) extended the simple exponential smoothing method to allow forecasting of time series with a trend. The concept behind the model is to introduce a term that takes into account the possibility

Observation	α**=0.2**	α**=0.4**	α**=0.6**	α**=0.8**
$\mathbf{y}_t$	0.2	0.4	0.6	0.8
$\mathbf{y}_{t-1}$	0.16	0.24	0.24	0.16
$\mathbf{y}_{t-2}$	0.128	0.144	0.096	0.032
$\mathbf{y}_{t-3}$	0.1024	0.0864	0.0384	0.0064
$\mathbf{y}_{t-4}$	$(0.2)(0.8)^4$	$(0.4)(0.6)^4$	$(0.6)(0.4)^4$	$(0.8)(0.2)^4$
$\mathbf{y}_{t-5}$	$(0.2)(0.8)^5$	$(0.4)(0.6)^5$	$(0.6)(0.4)^5$	$(0.8)(0.2)^5$

Table 5.1: Impact of smoothing parameter α on the weights of the different observations y_{t-i}.

that a series could contain a trend. This term is the smoothing parameter for the trend, known as the parameter γ, and is also updated with the help of exponential smoothing. In addition to a regular trend component, also exponential smoothing models have been used with a so called *damped trend*. This damped trend is used to prevent over forecasting, when a constant/exponential (increasing/decreasing) trend is estimated indefinitely into the future. To prevent the resulting forecast to 'explode', because of an indefinitely increasing or decreasing estimated trend, the trend will be damped/weakened (hence the name of the specific trend component). Especially when the forecast horizon becomes very long, a damped trend can prevent the model from over forecasting and improve the forecast accuracy (SBRANA, 2012). Looking at the mathematical representation, the m steps ahead forecast at time t ($\hat{X}_t(m)$) for the double exponential smoothing method (with additive trend) is given by:

$$
\begin{gathered}
S_t = \alpha X_t + (1-\alpha)(S_{t-1} + T_{t-1}) \\
T_t = \gamma(S_t - S_{t-1}) + (1-\gamma)T_{t-1} \\
\hat{X}_t(m) = S_t + mT_t \\
0 \leq \alpha \leq 1 \\
0 \leq \gamma \leq 1,
\end{gathered}
$$

where $S_t - S_{t-1}$ is the trend and T_t is the smoothed additive trend at the end of period t.

Triple Exponential Smoothing

Triple exponential smoothing methods are able to handle data which has both a trend as well as a seasonal component, where the double exponential smoothing was just able to detect a trend component. Seasonality is defined to be the tendency of time-series data to exhibit behavior that repeats itself every L periods. "The term season is used to represent the period of time before behavior begins to repeat itself" (Kalekar, 2004). L is therefore the season length in periods. There are two types of seasonality: *additive* and *multiplicative*. "The additive method is preferred when the seasonal variations are roughly constant through the series, while the multiplicative method is preferred when the seasonal variations are changing proportional to the level of the series" (Hyndman and Athanasopoulos, 2014). An example of additive seasonality is when every Tuesday 100 tonnes of cargo more than on Wednesday is transported. In this case the seasonality is additive in nature and can be represented as an absolute amount. Looking at the same example if an airline transports 10% of cargo more, the seasonality is multiplicative in nature and can be represented as a constant factor. In order to enable the triple exponential smoothing method to detect seasonal patterns, a rule of thumb is used which requires that at least two full seasonal cycles have to be present in the data.

Looking at the formula of the triple exponential smoothing method (with additive trend and additive seasonal component), the m steps ahead forecast ($\hat{X}_t$(m)) is given by:

$$
\begin{gathered}
S_t = \alpha(X_t - I_{t-p}) + (1-\alpha)(S_{t-1} + T_{t-1}) \\
T_t = \gamma(S_t - S_{t-1}) + (1-\gamma)T_{t-1} \\
I_t = \delta(X_t - S_t) + (1-\delta)I_{t-p}
\end{gathered}
$$

$$\hat{X}_t(m) = S_t + mT_t + I_{t-p+m}$$
$$0 \leq \alpha \leq 1$$
$$0 \leq \gamma \leq 1$$
$$0 \leq \delta \leq 1,$$

where I_t is the smoothed seasonal index at the end of period t. This smoothed seasonal index is, among others, estimated by the smoothing parameter for seasonal indices δ and the number of periods in the seasonal cycle p.

Parameter estimation
When forecasting with exponential smoothing methods, one should keep in mind that the accuracy of the forecast can be tremendously improved by choosing the 'right' values for all the parameters. Therefore, an optimization method has been implemented to find the 'optimal' values for the different smoothing parameters α, γ and δ (depending on the exponential smoothing method that is used). To be able to interpret the estimated values for all the different parameters, first the impact of the different possible values for each parameter will be discussed.

As mentioned before, the smoothing parameter α will attach more weight to more recent observations when its value gets closer to 1. When the value of α is closer to 0, more weight is given to observations from the more distant past. The trend smoothing parameter γ affects the trend component of the exponential smoothing model. When the value of this parameter γ is equal to 0, then the trend component is constant for all values of the time series. If the parameter value is getting closer to 1, then the trend component will be modified at every step by the corresponding forecast error. Roughly speaking, considering the value of the seasonal smoothing parameter δ, a similar effect as with the trend smoothing parameter can be observed. When the value of the parameter δ gets closer to 0, the seasonal component for a certain point in time will be identical to the predicted value for the same time observation during the previous seasonal cycle(s). When the value of δ is equal to 1, then the seasonal component is modified at every step by the corresponding forecast error (Zahedi, 2008).

To optimize the parameter values for the exponential smoothing models, an optimization method is used that consists of a function solver and a Key Performance Indicator (KPI) which tries to find a value for the parameters such that the value of this KPI is minimized. To see whether the use of different KPIs will result in better forecasts, several KPIs are tested: Mean Absolute Percentage Error (MAPE), Weighted Absolute Percentage Error (WAPE), Root Mean Square Deviation (RMSD) and the Symmetric Mean Percentage Error (SMPE) (variant of the SMAPE). All of these are formulated below. Within these formulas, y_t is the observed value at time t, S_t the forecasted (smoothed) value at time t and n the number of observations.

$$MAPE = \frac{1}{n}\sum_{t=1}^{n} |\frac{y_t - S_t}{y_t}|$$
$$WAPE = \frac{\sum_{t=1}^{n} |y_t - S_t|}{\sum_{t=1}^{n} y_t}$$
$$RMSD = \sqrt{\frac{\sum_{t=1}^{n} (S_t - y_t)^2}{n}}$$
$$SMPE = \frac{\sum_{t=1}^{n} S_t - y_t}{\sum_{t=1}^{n} y_t + S_t}$$

Depending on the exponential smoothing method and the KPI that are used, the optimal value(s) for the parameter(s) will be determined. Therefore, a function solver tries to minimize the error between the observed value y_t and the smoothed value S_t by tuning the parameter values within the allowed interval (only values between 0 and 1 are allowed). The resulting minimization problem can be formulated as:

$$\begin{array}{rccccc} \text{Minimize} & KPI & & & & \\ \text{s.t.} & 0 & \leq & \alpha & \leq & 1 \\ & 0 & \leq & \gamma & \leq & 1 \\ & 0 & \leq & \delta & \leq & 1 \end{array}$$

To solve this minimization function, a Sequential Least SQuares Programming (SLSQP) solver with a step size of 0.05 used for the numerical approximation from the Scipy package of Python has been used.

5.1.2 Autoregressive Moving Average (ARMA)

An Autoregressive Moving Average (ARMA) time series model is used to forecast a single time series and consists of two parts: an *autoregressive* and a *moving average* part.

An autoregressive (AR) model expresses a time series as a linear function of its past values. As an input for these kind of models, the only parameter to provide is the number of lagged past values p that are included. Looking at the formula of an autoregressive model of order p, i.e. an AR(p) model, this can be generally formulated (by recursive substitution) as:

$$X_t = \phi_1 X_{t-1} + \phi_2 X_{t-2} + ... + \phi_p X_{t-p} + \epsilon_t, \epsilon_t \sim \ WN(0, \sigma_\epsilon^2),$$

where X_t is the value of the observation of the time series at time t, ϕ_p the autoregressive coefficient of lag p and ϵ_t the error term at time t. The term WN stands for a white noise distribution (whose characteristics will not be discussed in this research).

The second part, the moving average (MA) model, is a variant of an ARMA model in which the time series is regarded as a moving average (unevenly weighted) of a random shock series ϵ_t. The only input parameter for moving average models is the number of the time lag q on the error terms (also called the residuals) (Meko, 2009). Considering a moving average model of order q, a MA(q) model, the model can be generally formulated as:

$$X_t = \epsilon_t + \Theta_1 \epsilon_{t-1} + ... + \Theta_q \epsilon_{t-q}, \epsilon_t \sim \ WN(0, \sigma_\epsilon^2),$$

where X_t is the value of the observation of the time series at time t, Θ_q the moving average coefficient on lag q and ϵ_t the error term at time t.

From the models above, a combined model can be created by combining the lagged terms (on the time series itself (AR) and on the error terms or residuals (MA)). The resulting model is called the Autoregressive Moving Average (ARMA) model. The generic formula of an ARMA(p,q) model can be formulated as:

$$X_t = \phi_1 X_{t-1} + \phi_2 X_{t-2} + ... + \phi_p X_{t-p} + \epsilon_t + \Theta \epsilon_{t-1} + ... + \Theta_q \epsilon_{t-q}, \epsilon_t \sim \ WN(0, \sigma_\epsilon^2),$$

, where ϕ_1,...,ϕ_p, Θ_1,..., Θ_q are the autoregressive and the moving average coefficients of lag $1, ..., p$ and $1, ..., q$ respectively.

Stationarity
Many time series show trends and/or seasonal behavior. This trend and seasonality components affect the value of the time series at different moments. This means that the mean and variance of the time series are time dependent. In other words, the joint probability function of the time series process changes over time. Time series that showing such behavior are called *non-stationairy*.

Most time series models, including ARMA models, assumes that this property is not in the time series that will be forecasted (Coghlan, 2010). In other words, those time series models assumes the input time series to be a *stationary* process. Looking at the definition, a time series process is said to be (weakly) stationary if the first and second moment are time invariant (Stigler, 2014):

$$E(X_t) = \mu, \forall t$$

$$Var(X_t) = \gamma_0 < \infty, \forall t$$

$$Cov(X_t, X_{t-k}) = \gamma_k, \forall t, \forall k$$

To check those conditions, a so called 'unit root' test can be used (which will not be explained in this report, for more info see (Hyndman and Athanasopoulos, 2014)). This kind of test is designed for determining whether a time series is stationary or not. In this research the *Augmented Dickey Fuller* test is used with a 5% significance threshold. When the outcome of the test shows that the time series is non-stationary (so one of the conditions above does not hold), the stationary condition often is enforced by *differencing* the time series. The differenced series is the change between consecutive observations in the original series, and can be written as (Hyndman and Athanasopoulos, 2014):

$$X_t' = X_t - X_{t-1},$$

where X_t is the observed value at time t and X'_t the differenced value at time t.

After the time series is differenced, once again the stationary conditions have to be checked to see whether the resulting time series is a stationary process. If this is not the case, then the time series has to be differenced once again etcetera.

Model estimation
When forecasting with ARMA models, first the time series should be stationary. When this is the case, the parameters p and q of the ARMA(p,q) model have to be estimated. To do this, two measures, called Information Criterion (IC), are used: the Akaike Information Criterion (AIC) and the Bayesian Information Criterion (BIC). Those information criteria can give an indication of the goodness-of-fit of the model with respect to the data. When looking at the output of both information criteria, the model with the lowest AIC or BIC will be chosen. Both criteria charge a penalty for the use of too many parameters p. The measures are defined as:

$$AIC(p) = n \log \hat{\sigma}_p^2 + 2p$$

$$BIC(p) = n \log \hat{\sigma}_p^2 + p \log n,$$

where n is the sample size, p the total number of parameters estimated and $\hat{\sigma}_p^2 = RSS/(n-p)$ (where RSS is the residual sum of squares). The residuals, $\hat{\sigma_p}$, are calculated by taking the difference between the observed value at time t and the fitted value. Then, to calculate the RSS, the sum of the squared residuals will be calculated: $\sum \hat{\epsilon}_t^2$. When looking at those both criteria, these goodness-of-fit criteria are 'like' the $\hat{R}^2$ or minimum $\hat{\sigma}^2$-type criterion (Maddala, 1992).

5.1.3 Decomposition of the time series

Although the ARMA and exponential smoothing models can detect trend and seasonality (only exponential smoothing models), also a decomposition method has been considered to see whether this could improve the forecasts. For the seasonality component that is detected, a weekly pattern is considered. For ARMA models, this method can help to improve the forecast, because of the method itself is *not* designed for detecting seasonality patterns. Considering the exponential smoothing models, the decomposition method does not fully take over the detection of a seasonal component. Since the decomposition method, as we use it, only detects weekly seasonality, (seasonal) exponential smoothing models could still be useful for detecting month or year seasonality.

The decomposition method, hence the name of the method, decomposes the time series into three components: trend, seasonality and an error component. An example of such a decomposition of a time series can be found in Figure 5.1. The motivation for the use of this decomposition method is that experience from Bokern and Mulder, this method estimates the seasonal effect well (Bokern and Mulder, 2015).

First the trend of the time series will be determined by taking a moving average over the time series. This moving average considers a symmetric window with the length of the predefined frequency of the data. The frequency of the data, that is used for this research, is equal to 7 (daily data). So the moving average at a certain day x, will be the average of the values at times $x-3$, $x-2$,..., $x+2$, $x+3$.

After the trend component of the time series has been determined, it will be removed from the time series. Next, from the resulting time series the seasonal component will be determined. This will be done by taking an average for each specific time unit in the frequency period of the data over all periods. For example, considering the Monday, the value of the seasonal component for the Mondays will be the average over all Monday values.

Finally, the error component is determined. For this the determined trend and seasonal component will be subtracted from the original 'raw' time series. The resulting series is considered to be the error component.

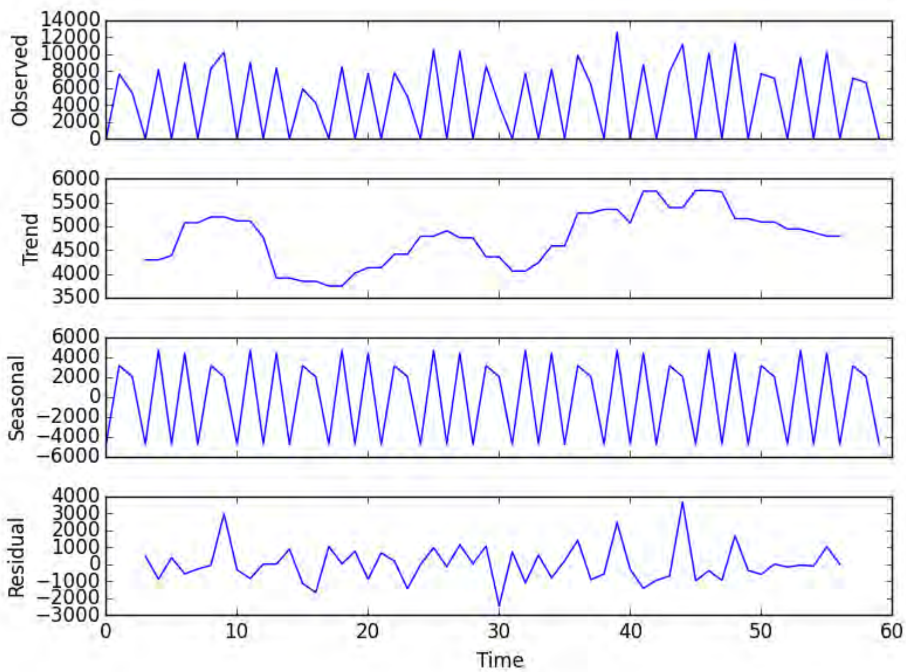

Figure 5.1: Decomposition of the time series of Flight KL0409 Amsterdam (Netherlands) - Almaty (Kazakhstan)

5.1.4 Forecast run design

To be able to determine what set of parameters performs best for a particular flight, a scenario analysis will be performed. In fact, during this analysis a grid search is performed over all possible scenarios (all combinations of predefined parameter values). Considering the parameters last date and amount of history, forecasts of different moments in time can be created by cutting off the time series. The initial guess parameter gives a start solution for searching for the optimal parameter settings when estimating an exponential smoothing model. Taking different initial guess values for those exponential smoothing parameters, can prevent the model ending up in a local optimum. Finally, looking at different forecast error key performance indicators might give different forecasting results, hence each of those methods determines the accuracy of the forecast in another way. It turns out that it is not the case that a particular forecasting method will always perform better than another, when comparing forecasts for different flights.

To be able to determine in advance the set of starting parameter values that performs best for a specific flight at a specific moment in time, all (user defined) scenarios will be used to forecast the past 10 days. Next, the WAPE KPI is used to score the different methods. Therefore, the forecast for the past 10 days are compared with the actual flown cargo of the last 10 days. In that way, the 'best' performing forecasting model and scenario combination can be determined for each individual flight. Those 'best' settings will then be used to forecast the cargo workload for the coming 10 days.

Considering the implementation of the forecast run, the forecasting code has been developed on an Hadoop cluster in combination with the use of Spark (open source engine for processing big data) and the programming language Python. To be more specific, the generic forecasting component has been developed in Python and the Hadoop cluster and Spark are used to do the actual computation. This computation of the forecasts for all single flights has been done in parallel. In this way, the computation power of the cluster will be fully exploit and the run time will be much lower.

5.2 Forecast on booking data

In addition to the forecast based on historical flight data, also a forecast based on booking data has been created. The advantage of using booking data compared to the historical flight data is that booking data contains actual and more detailed information about the workload that can be expected at the day of operations. To create such a forecast, a simple linear regression model (see Section 5.2.1) has been used to estimate the final weight that will be delivered at the Schiphol Cargo Hub of all bookings made up to a specific Days Before Departure (DBD) x. In other words, the model tries to estimate the *materialization* of the bookings at DBD 0. To estimate this materialization of all bookings made up to a certain DBD x (called Bookings at hand), the model searches for a linear relationship between the booked and the actual weights of the bookings made based on bookings from the past. With the resulting relationship, the materialization of future bookings will be predicted.

When 'knowing' the materialization of bookings at a certain DBD, one should also keep in mind that there are still bookings to be made. So, in addition to forecasting the final materialization of a booking, also an estimation has to be made of total weight of future bookings. For this estimation also a linear regression model will be used. In this case the model tries to find a relationship between the bookings that are made up to DBD x and the overall total weight of all bookings (also incorporating the bookings that are made after DBD x) that will depart at DBD 0.

Next to the booked weight that is in the booking data, a more accurate estimate of the materialization of the booking is already available before the booking arrives: Opportunity Assessment (OA). This is a tool, made by Air France colleagues, which predicts the materialization of a booking based on specific booking characteristics. The underlying forecasting technique is a decision tree algorithm that is trained on one year of historical booking data. This forecast could help to improve the forecasting accuracy of the booking forecasts. To see whether this is indeed the case and to see whether OA might be an enrichment to the model, a comparison is made between the booked weight and the forecasted materialization by OA. However, if it turns out that the OA data improves the model, the final booking forecast (for the moment) will still be based on booking data. This has to do with the limited amount of OA data that is available (only two months versus almost a year of booking data).

5.2.1 Simple linear regression

A simple linear regression model tries to determine a (linear) relationship between an *explanatory* or *predictor* variable x and an *explained* or *response* variable y. This relationship can be formulated as

$$y = \alpha + \beta x + u,$$

where u is called an *error* term, α the intercept and β the slope coefficient of the regression line.

To estimate the regression coefficients α and β, a least squares method is used. This method tries to minimize the sum of squares between the observed value y and a predicted value x. In other words, the method of least squares tries to find a value for $\hat{\alpha}$ and $\hat{\beta}$ (estimators for α and β), such that the sum of squares Q will be minimized. This sum of squares Q is defined as:

$$Q = \sum_{i=1}^{n} (y_i - \hat{\alpha} - \hat{\beta} x_i)^2,$$

where n is the number of observations.

In Figure 5.2, an example of a sample of observations and an estimated regression line is given. The combination of the text above together with this figure provides an easy understanding of how the simple linear regression method works. Also the different parameters are visually explained in the figure.

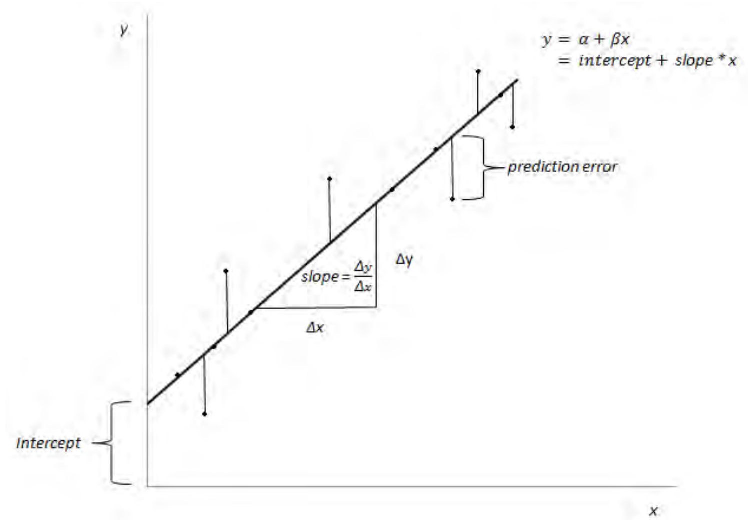

Figure 5.2: A sample of observations and an estimated simple linear regression line

Chapter 6

Data

In this chapter the data that is used will be described. Then the data analysis will be presented, to get more insights into the data that is used.

6.1 Data Description

For this study we used two main data sources: *Flight* data and *Booking* data. The flight data is obtained from the French part of the Operations Research Department of AFKL. The booking data has been retrieved from the RM Cargo department of AFKL. The origin and content of both data sources will be described below.

6.1.1 Flight data

The *Flight* data contains details about all flights from the beginning of 2010 until September 2015. Looking at this historical flight data, the data includes fields like: flight number, aircraft type, origin and destination of the flight and the total amount of cargo (in weight and volume) that has been transported. For all fields that are in the data, see Appendix A, Table A.1. Because the flight data is post departure data, it is considered to be accurate. We emphasize that this data source contains ICA flights (transporting cargo intended for freight building 2 and 3), but no EUR flights transporting cargo intended for freight building 1.

From the flight data source more insight can be retrieved into the cargo demand fluctuations throughout the different seasons (high and low season) and even throughout the week. Furthermore, the historical flight data gives information about the expected workload that can be expected in the future. It is obtained from the RM Cargo department (not directly from a data system). The RM department itself obtains information about cargo that has been transported from a database system (called Altea). With this information RM creates a flight data file which incorporates this data enriched with some calculations.

OAG data

Official Airline Guide (OAG) provides up-to-date and highly accurate flight schedule information of more than 900 airlines all over the world. This information can be used to get insight into the flight schedules of other carriers. With respect to the forecasting part of this study, flight schedule information provided by OAG provides high valuable information when creating flight forecasts. For example, when it is known that a flight X does only fly on Mondays and Wednesdays, cargo demand forecasts created for all other weekdays can be set to 0. This will lead to an overall improvement of the forecast performance. An overview of all the data fields that are in the OAG scheduling data files can be found in Appendix A, Table A.2.

6.1.2 Booking data

The *booking* data contains all information of the incoming bookings that have been made after December 2014. Booking information contains information about the amount of cargo (in weight and volume), the origin and destination of the cargo, the product that will be transported, etc. All the fields in the booking data are summarized in Appendix A, Table A.3. This booking data can be valuable for KLM Cargo helping them estimating the expected workload at the Schiphol Cargo Hub at a certain day in the near future. In contrast to flight data, where only historical data is available, booking data also provide the analysts with 'actual' information about future bookings. Furthermore, this booking data contains more detailed information such as product code and whether a booking has been canceled during the booking period. Before any analysis on the data has been done, it is important to make clear where the data originates from. In Figure 6.1 a schematic overview is given about the data warehouse used for bookings.

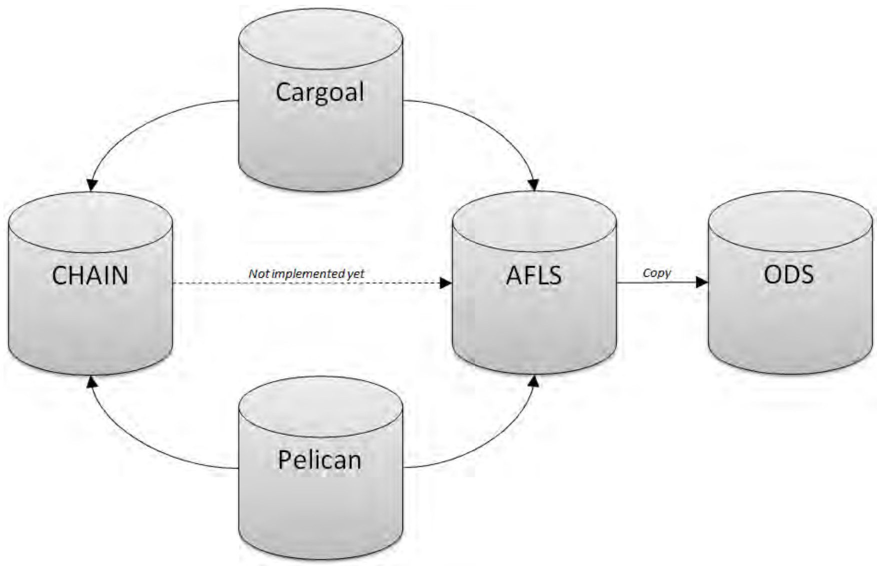

Figure 6.1: Origin booking data

First of all, there are the two cargo booking tools of AFKL. Air France Cargo bookings are made in 'Pelican', the French cargo booking tool. At KLM Cargo there is a similar booking tool called 'Cargoal'. Both systems were operating independent from each other until the second half of 2015, when Accenture Freight & Logistics Software (AFLS) integrated both systems. Ever since, both of the systems feed the AFLS system with their booking data. Beside feeding AFLS, also a system called 'CHAIN' is fed by both 'Pelican' as well as 'Cargoal'. The CHAIN system, is a system used by the Cargo Operations Department to track all the cargo from the moment the cargo arrives to the moment the cargo departs. In this 'CHAIN' system, updates on the cargo bookings will be registered when the actual cargo shipment differs from the original booking. At the moment, this update information is not available at the AFLS system. This part of the integration process will be finished in the first half of 2016.

The last remaining data source is the ODS database. This database is updated a couple of times a day and contains a copy of the AFLS database. The ODS database is also the data source that has been used for this study. In the future, when CHAIN is fully integrated with the AFLS system, the ODS database will contain more reliable and up-to-date information. In the end, this will result in a more reliable analysis in the future.

6.1.3 Data quality

Considering all data sources that have been used, we should keep in mind the reliability and the completeness of the data. First of all, the flight and the booking data sources used for the forecasts both have their advantages and disadvantages. Considering the *reliability* of data, the flight data is more reliable than the booking data source. The reason for this is that the flight data source contains post departure data and therefore this source is more accurate in comparison with the pre departure booking data source (where still updates on the data can be made). Because the booking data is pre-departure, it is not always the case that the final materialization of a booking is known. Therefore, the last known booked weight of a booking is considered to be the final materialization. This is also the value used for scoring the forecast accuracy.

Looking at the *completeness* of the data, it turns out that the booking data source is more complete. This can be explained by the fact that the flight data only contains ICA flights and therefore misses all European flights. Both the reliability as well as the completeness of those data sources should be kept in mind when looking at the final results.

6.2 Data Analysis

To be able to create a forecast based on the data sources described above, an initial analysis has been done to get some feeling for the data. For this first analysis only AF, KL and DL flights are considered, because those are the only carriers whose cargo will be handled by the Schiphol Cargo Hub. Furthermore, differences among the summer/high season (April to September) and the winter/low season (October to March) have been considered.

6.2.1 Flight data analysis

In Figures 6.2 and 6.3 a closer look has been taken at the number of flights (that contain cargo) that arrive and depart from Amsterdam Schiphol Airport and the amount of cargo that has been processed. From this figure, several conclusions can be drawn.

First, around new year we can observe a drop in the data, considering the total sum of cargo that departs from AMS. Furthermore, at the 17th of June an outlier in the data can be identified when looking at the AMS arrival graph. However, in contrast to the cargo weight break at the beginning of 2015, there is no clear explanation for this second observation. When looking at the number of flights, also an outlier is noticeable. Since no reasonable explanation could be found from the business perspective regarding those strange results, those observations can be seen as outliers in the data.

Secondly, looking at the differences between summer and winter season it seems that during summer season more aircrafts are coming to AMS. A reasonable explanation for this is that during the summer season more people travel by air due to the summer holidays. We can conclude this from the fact that the total amount of transported cargo stays more or less the same, where the number of aircrafts increase.

Because no analysis with respect to carrier level can be performed, we looked a bit more into detail at carrier level. From these results, insights can be obtained into the average number of aircrafts that arrive/depart from AMS during the different days of the week. Again, we have done this analysis for both summer as well as the winter season. From the resulting Figure 6.4, the overwhelming amount of KL flights is something that can be directly concluded. Furthermore, the number of AF flight can more or less be neglected. The low number of AF aircrafts can be explained by the fact that cargo destined for AMS, will (almost) always be transported by KL. Moreover, when cargo intended for AMS arrives at CDG, the cargo will most of the time be trucked to AMS instead of transported by aircraft.

Another interesting observation is the shift of the cargo peaks. Looking at the arrival peaks (Thursdays and Saturdays) and looking at the departure peaks (more on the Fridays and Sundays) we can notice a shift of one day. Roughly speaking, this means that the aircrafts that are coming in, leave the day after. Considering the overall average (blue dotted lines), the average number of aircrafts arriving and departing from AMS are more or less equal (what we would expect).

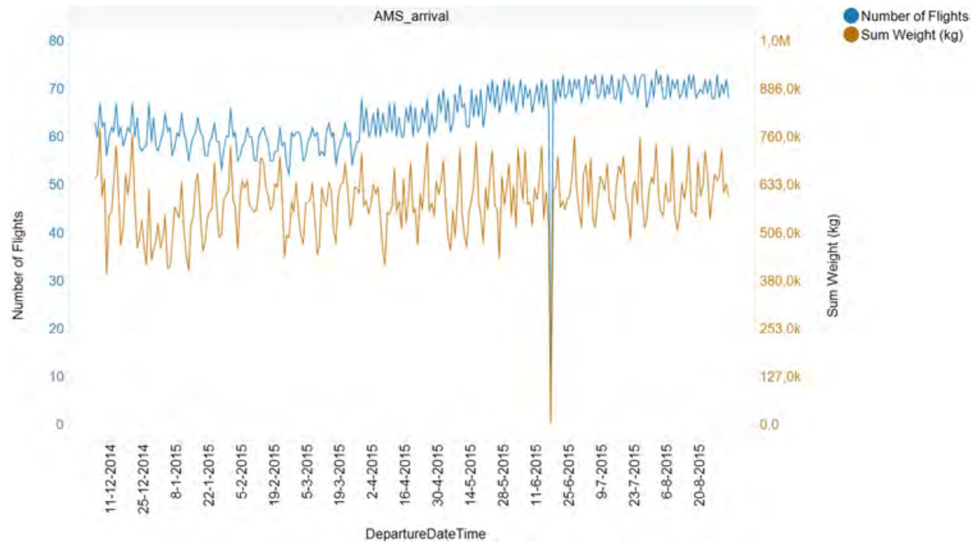

Figure 6.2: Number of aircrafts (AF, KL, DL) and amount of cargo that arrives at Amsterdam Schiphol Airport

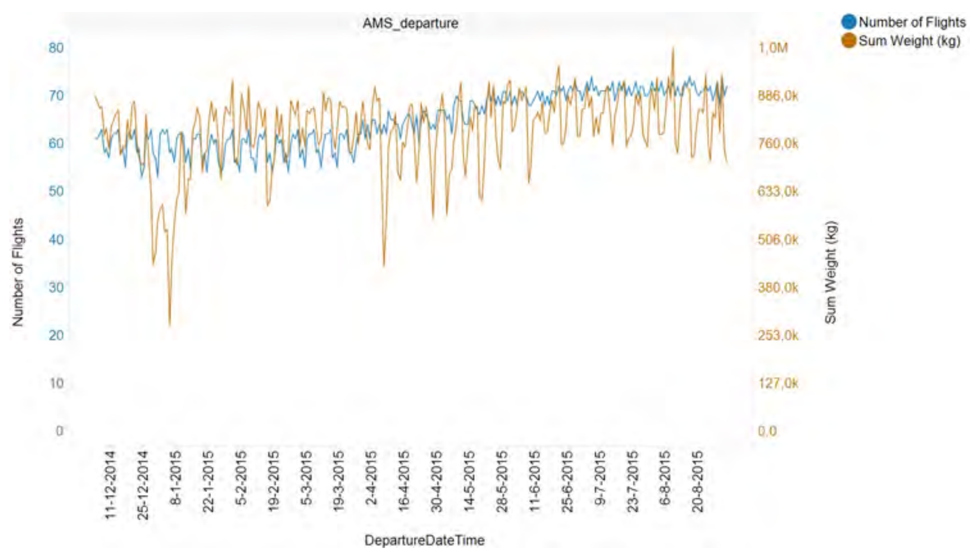

Figure 6.3: Number of aircrafts (AF, KL, DL) and amount of cargo that departs from Amsterdam Schiphol Airport

Considering the average amount of arriving and departing cargo from Amsterdam Schiphol Airport, an interesting observation based on Figure 6.5 is the relatively large difference between the incoming and outgoing amount of cargo. This can also be observed by looking at the averages (blue dotted lines) again. Looking at the amount of cargo that arrives at AMS, there is a slight peak noticeable on Saturday. Considering the outgoing cargo, the weight of cargo increases throughout the week and most of the cargo leaves AMS during the weekend. Between summer and winter season, the average amount of cargo transported per week seems to be higher during summer season. A last interesting thing that can be concluded from this figure is the difference of the arriving and departing cargo at AMS that has been transported by Delta Air Lines (DL) (airline which has entered into joint venture with AFKL). Looking at the incoming and outgoing amount

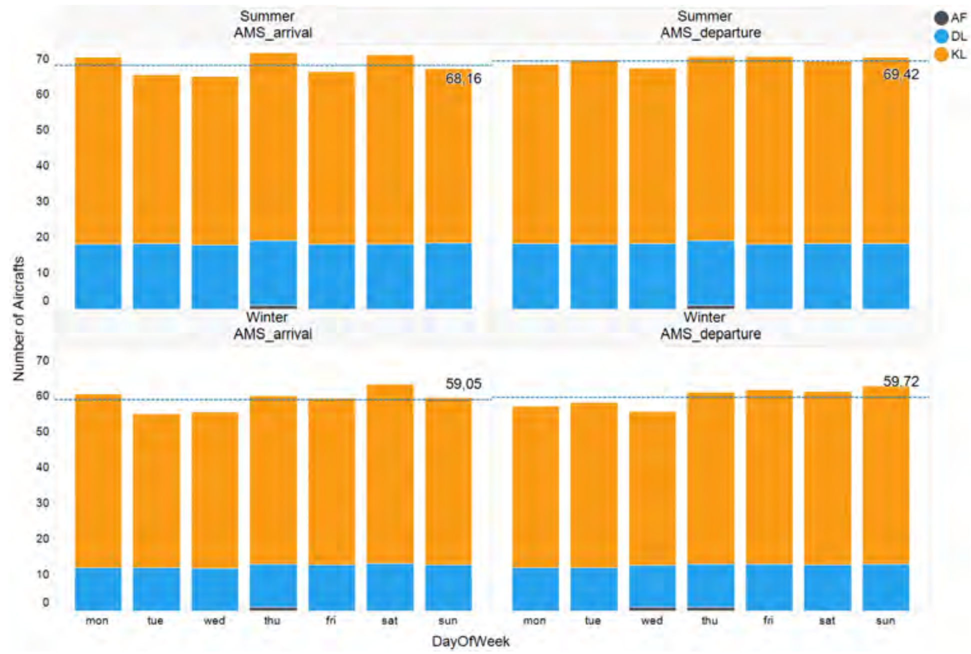

Figure 6.4: Average number of arriving and departing aircrafts from Amsterdam Schiphol Airport split by carrier and season

of cargo transported by DL, the incoming amount of cargo can almost be neglected in comparison with the outgoing cargo. An explanation for this could be that airlines often transport more cargo to their own hubs.

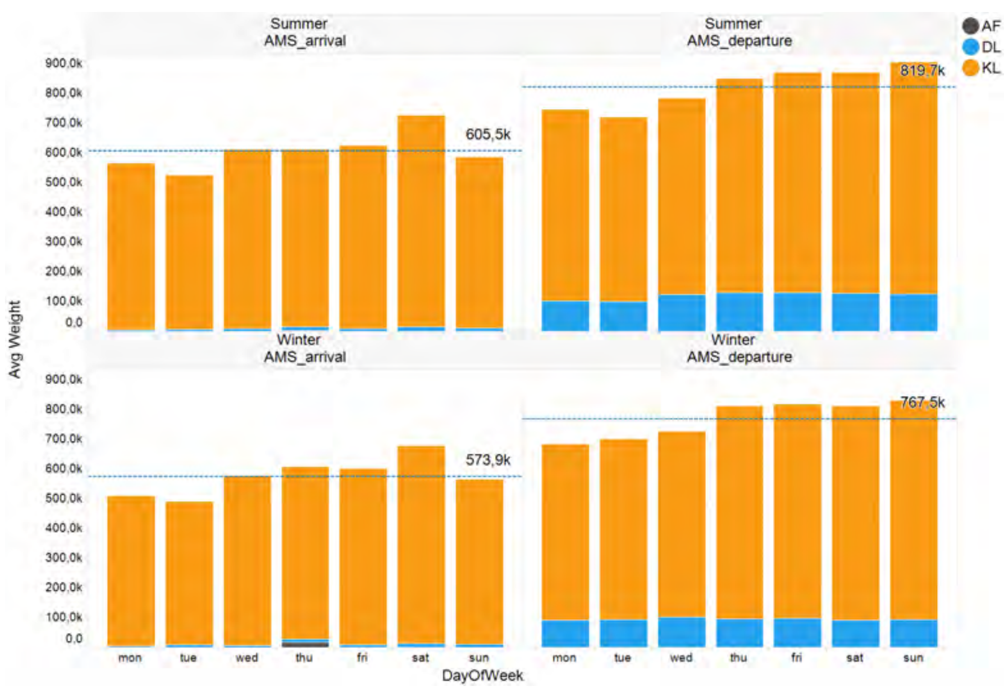

Figure 6.5: Average amount of arriving and departing cargo (kilograms) from Amsterdam Schiphol Airport split by carrier and season

6.2.2 Booking data analysis

Booking data give pre-departure information of the workload that can be expected for a certain departure date in the near future. From the moment that the booking window for a future date opens, cargo bookings come in. To get some feeling of how bookings show up at the warehouse, an analysis has been performed on historical booking data. In this analysis, the latest booking weight update is considered to be the actual (or *manifest*) value. First an single booking performance analysis has been performed to see how booking show up at the day of operations with respect to its booked weight value. Second, there has been looked at the booking behavior of people; at what days what portion of bookings are made for a specific day of departure. Finally, also an analysis has been performed on the transit time of the cargo bookings. The transit time is the time between the arrival of the booking at the warehouse and the time the booking will leave by flight/truck.

Show up rate

In Figure 6.6 a scatter plot is made of the manifest value of a single booking versus the booked weight value. Each data point represents a single booking made x days before departure (DBD) for $x = 0$, 1, 2 and 3. In this analysis only bookings that were made at least three days before the day of operations were considered, to make a fair comparison between the different DBDs. In Table 6.1 the results are shown in numbers.

Based on the results that can be obtained from the table, it seems that there is a lot of probability mass on the diagonal. On this diagonal are all bookings whose booked weight is equal to the actual weight. From Figure 6.6, we can observe that as the day of operations approaches more data point move to this diagonal. Therefore, we can conclude that the estimate of the bookings become more accurate, in comparison with the manifest weight, when the day of operations approaches. This can be due to bookings that are updated during the booking period (or even canceled). Considering the table we can see that when the day of operations approaches, the percentage *Show* increases and both the percentages *High* and *Low* shows decrease. These results emphasize the conclusions based on Figure 6.6.

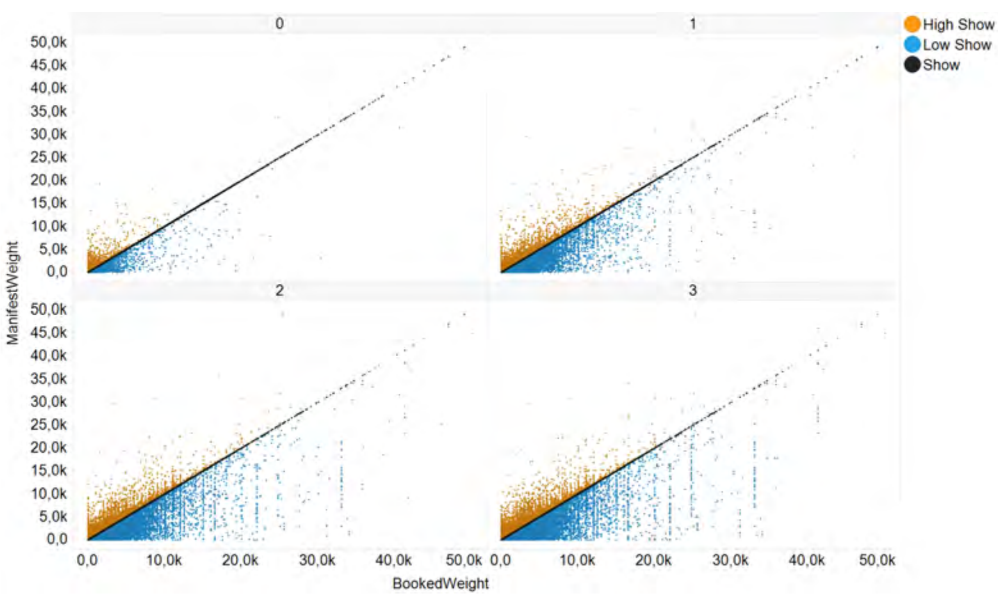

Figure 6.6: Booked versus Actual weight for 0, 1, 2 and 3 days before departure

To get a more clear overview of the differences between the booked weight and the actual weight of a booking, we performed an analysis on the weight error. The weight error is defined as the booked weight minus the actual weight. For this analysis the same data as before is used, so only the bookings that were made at least 3 days before the day of operations were considered. From all the weight errors, the distri-

Dmin	ShowUpStatus	Number of AWB	Percentage
0	High Show	1949	0,35 %
	Low Show	2974	0,54 %
	Show	544701	99,10 %
1	High Show	24113	4,39 %
	Low Show	52598	9,57 %
	Show	472913	86,04 %
2	High Show	54422	9,90 %
	Low Show	115387	20,99 %
	Show	379815	69,10 %
3	High Show	72701	13,23 %
	Low Show	151483	27,56 %
	Show	325440	59,21 %

Table 6.1: Summary show up status bookings

bution can be visualized in a histogram. The error distributions for the different days before departure are shown in Figure 6.7. Because some huge outliers were present in the data which blew up the histogram, we left them out of the figure. Within the histogram plots, the distinction between the different show up statuses (high show, low show, show) are made by coloring them. Furthermore, the average error as well as the origin are marked by a vertical line. Also the value of the average weight error is stated in the figure. From those results, one can observe that average weight error value is negative and gets closer to 0 as the day of operations approaches. This would imply that overall a higher percentage of the bookings will be a low show, because a negative weight error means that the booked weight is higher than the actual realized weight. Obviously, this is the same as was observed from Table 6.1. In addition to the average weight error that decreases when the day of operations approaches, also the shape of the distribution gets more concentrated around the perfect weight error of 0. This implies that the booking information gets more accurate over time (due to updates made on the bookings).

Booking behavior

To get a feeling for the overall booking behavior, a scatter plot has been created (see Figure 6.8) of the bookings that are made until day of week x for departure day of week y for all departure dates of 1st of January until 30th September. For example, considering departure day Monday, a scatter plot has been created of all bookings that are made for that specific Monday at the Tuesday, Wednesday,..., Sunday before. So looking at a specific data point on the horizontal axis the sum of all bookings made up to a specific booking day of week is stated (called the Bookings at hand), where on the vertical axis the total final amount of cargo is stated. This total final amount of cargo includes both the bookings made as well as the bookings to come.

Looking at the figure, a logical pattern can be recognized. When the day of departure approaches the bookings at hand gets closer to the total final amount of cargo. Furthermore some 'outliers' are observed which are distanced from the rest of the data, for example the data points at the 1.2M line in the Monday figure. Zooming in on those points showed that these outliers were all departure dates that occurred in the first week of January, so the low amount of transported cargo can be an aftereffect of New Year.

However, a more interesting thing that can be observed from the figure is that in general there is a little gap noticeable between DBD 1 and DBD 2. Except for the departure weekdays Sunday and Monday, where this effect is not there. This implies that at DBD 1 still a lot of bookings are made for the day of departure. This could be due to the fact that for departure days during the weekend (Sunday) or right after the weekend (Monday) bookings will be made before the start of the weekend (so longer in advance).

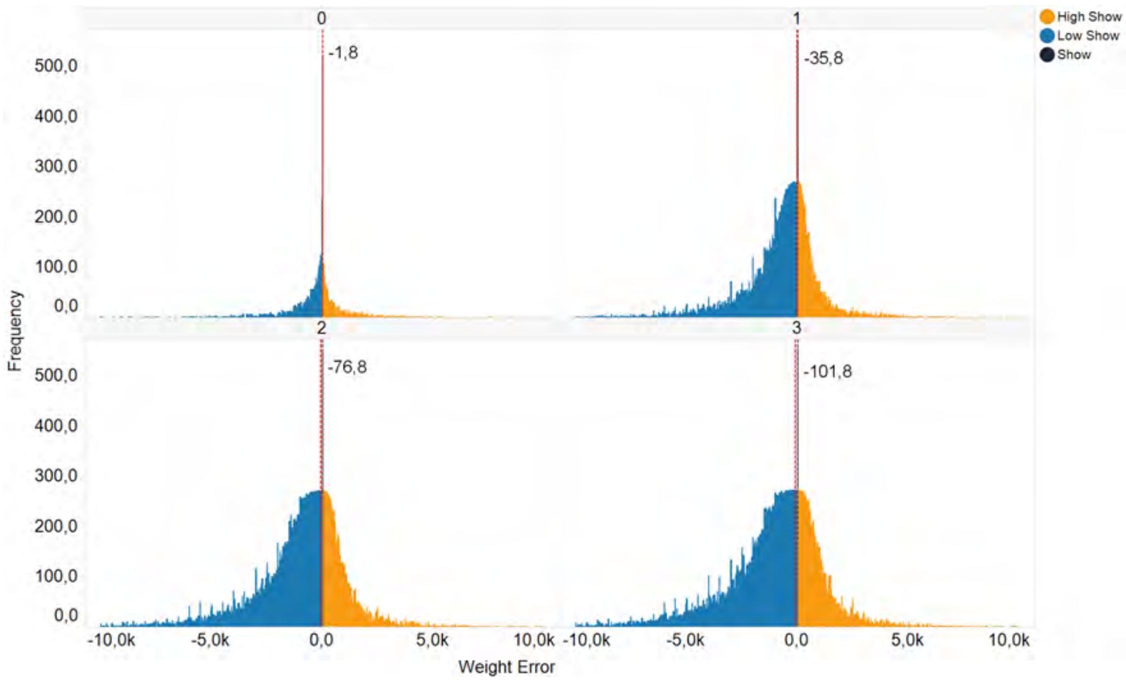

Figure 6.7: Weight error distribution for 0, 1, 2 and 3 days before departure

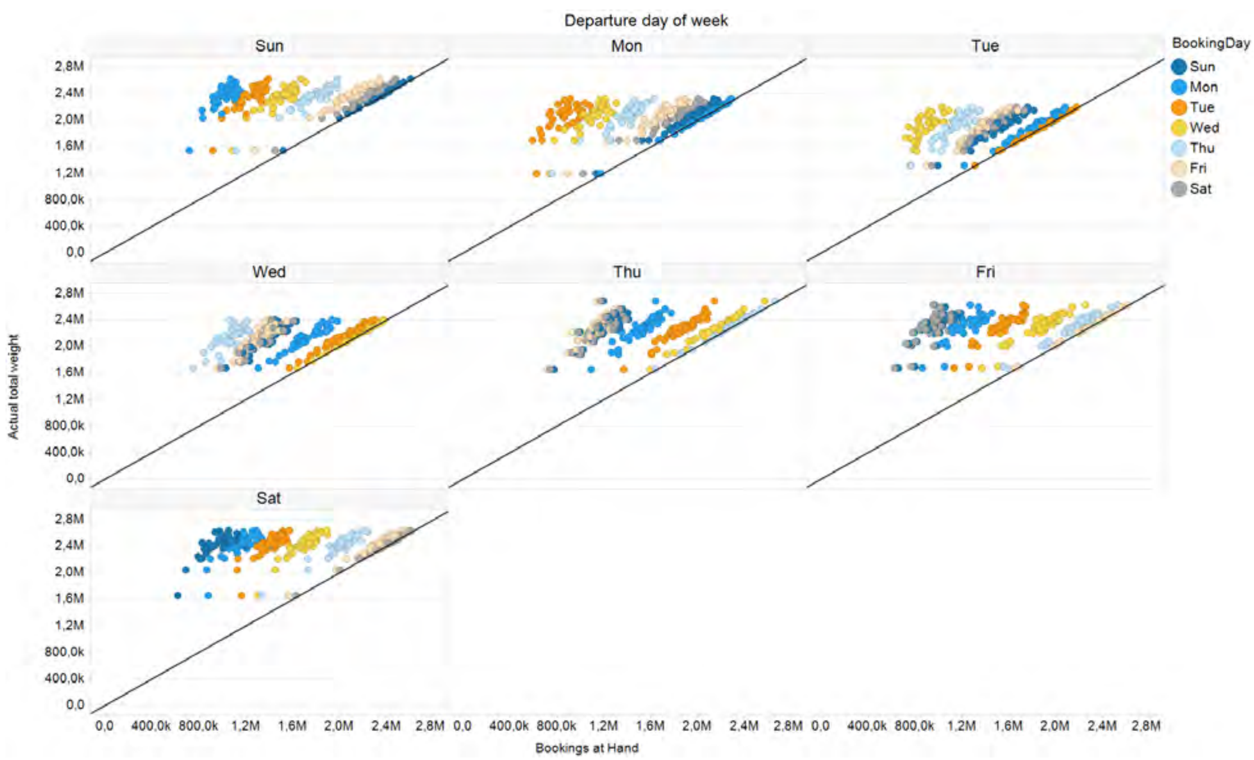

Figure 6.8: Booking behavior per departure day (colored by booking day of week)

Transit time

When approaching the day of operations, more and more bookings will arrive at the Schiphol Cargo Hub warehouse. Currently, KLM Cargo has the idea that the cargo that misses their connection arrives just before the deadline, such that there is not enough time to handle this cargo. With the help of the so called CHAIN universe, insight can be gained into the transit times of all bookings. This CHAIN universe is a database in the data warehouse containing data from the CHAIN application, a system that tracks the cargo from the moment it arrives at the Schiphol Cargo Hub until it leaves the warehouse (see Figure 6.1). Furthermore, based on this data, a distinction can be made between on time and out of time bookings. The results can be found in Figure 6.9.

The figure contains all bookings that were handled by the Schiphol Cargo Hub in March 2015 sorted by the booked day of departure. There are two vertically oriented grey lines plotted. Those lines represent the 5 and 24 hour moments before the departure time of the booking. In general, KLM Cargo uses an ultimate deadline for a booking to be on time if it arrives at least 5 hours before departure. The 5 hours are necessary to make sure the cargo will be on time for its connection (think of performing security checks, processing the cargo and transport and load the cargo into the trucks and aircrafts). Because of the short transit time, the cargo that arrives at the warehouse less than 24 hours before departure is called *Hot* cargo. On the opposite, if transit time is longer than 24 hours, the cargo is called *Cold*. In terms of planning, when processing what cargo, *Cold* cargo could be processed at a later stage compared to *Hot* cargo which has to be processed as soon as possible.

Looking at the bookings that are on time, the great majority of those bookings arrive at the warehouse a day before departure. Due to this late arrivals, the workload will be high for the Cargo Operations department, because all those bookings have to be processed right away. Only for the Mondays and Sundays, most of the bookings arrive more than one day in advance. Both conclusions are consistent with the results we obtained earlier from Figure 6.8. An explanation for the different behavior of the booking arrivals for the Mondays and Sundays, is that suppliers transporting their goods before the start of the weekend. This can be due to closed distribution centers of the suppliers during the weekend or more expensive storage costs at third parties over the weekend. This phenomenon is strengthened by the figures of the Saturday, Sunday and Monday. On the Saturday a high peak can be observed of bookings arriving within the last 24 hours before departure. For the Sunday there are a lot of bookings arriving at the last 48 hours and for the Monday a lot of bookings arriving even longer in advance.

Looking at the out of time bookings, a remarkable thing can be observed. It seems to be that most of the bookings that arrive too late for their connection, are already in the warehouse (or in the PCHS to be more precise) more than one day in advance. Based on this conclusion, it would be an interesting thing to see whether those bookings are out of time because of a shortage of resources or that this is more of a planning problem. This is out of the research scope of this thesis.

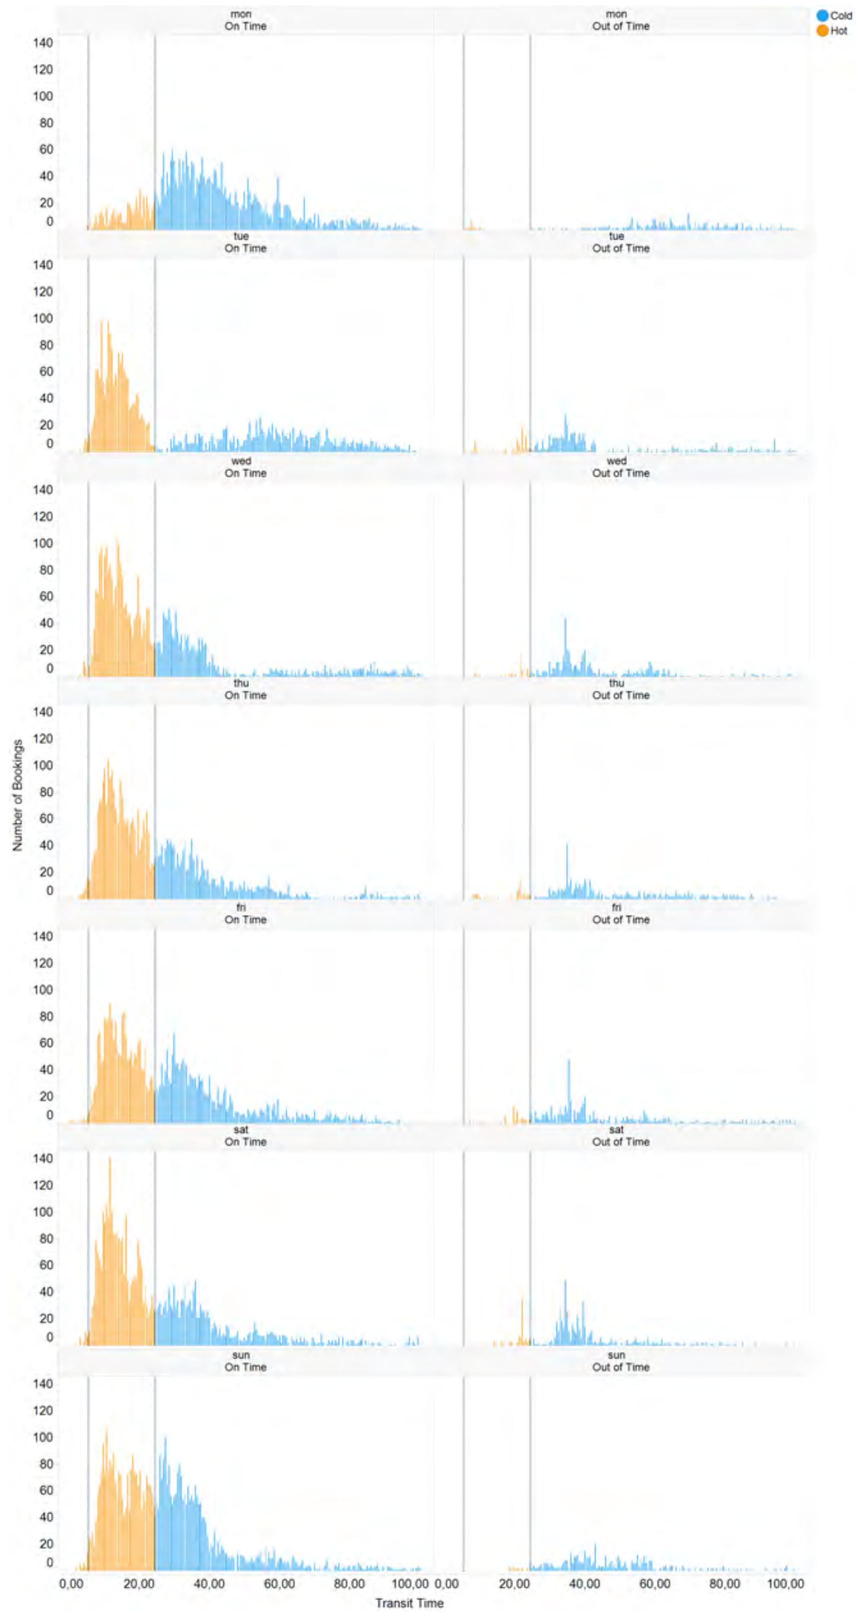

Figure 6.9: Transit time distribution (colored by hot (arriving ≤24 hours in advance) and cold (arriving ≥24 hours in advance) cargo

Chapter 7

Results

In this chapter the results from the forecasting models will be described. Furthermore, for the forecast based on flight level (see Section 7.1), a scenario analysis has been done to see what will be proper settings for the forecast model. To be able to analyze the performance of the forecast accuracy, a period from the past (third week of August, 2015) has been forecasted and compared to the actual observations.

For the booking forecast model (see Section 7.2), a comparison has been made of the differences between the two booking data sources: ODS and OA.

Looking at the results, we should keep in mind the completeness and reliability of both data sources (as described in Section 6.1.3). Care is required when comparing the final results. Since, at the time of this research, there were no clear business rules for splitting the total forecast into a forecast per freight building, there has been chosen (in consultation with the PRI department) to just create a forecast on board point and off point level.

7.1 Flight forecast

To take a look at the performance of each of the forecasting methods, a forecast run has been performed. Within this forecast run, a 10-days ahead forecast has been made for all flights that have flown after the 10th of August 2015 (taken into account 12 months of history and using an initial guess value of 0.5 for the exponential smoothing parameters). From those forecasts a comparison can be made between the different forecasting methods. The meaning of the method abbreviations can be found in Appendix B.

First, we take a look at an example of a single flight forecast (where just a subset of all forecast techniques used are plotted). This example is shown in Figure 7.1. From this example figure, it is easy to see at what days this specific operates. In this case, from the OAG data source, we know that that this flight is not operating on Mondays, Thursdays and Saturdays. Therefore, the forecast has been corrected to zero for those days.

Second, from the figure there can be observed that the forecasting methods all have a 'problem' with creating an accurate forecast for the Wednesdays. This could be due to schedule changes throughout the year (a similar example will be discussed in more detail in Section 7.1.1).

Next, an analysis of the best forecasting method based on the forecast accuracy can be done for all different flight forecasts. To calculate the forecast accuracy the WAPE has been used. The justification for using this criterion to score the accuracy of all forecasts, is that it can handle zero values in the actuals (MAPE cannot), is a relative measure (RMSD says something about the absolute deviation error) and is also relatively well protected against outliers (SMPE cannot handle outliers in the forecasts that cancel each others error). Therefore this error KPI will be used for computing the forecast accuracy. Considering the example, the WAPE value of the best performing forecasting method (exponential smoothing model with a damped multiplicative trend and no seasonal component) is 21,8%.

Although, for this example and with those specific parameter settings the forecasts perform well, there is no guarantee that this will also hold for other flight forecasts. Therefore, a scenario analysis will be done to see what the impact on the forecast accuracy will be of different sets of parameter settings. This will be discussed in Section 7.1.2.

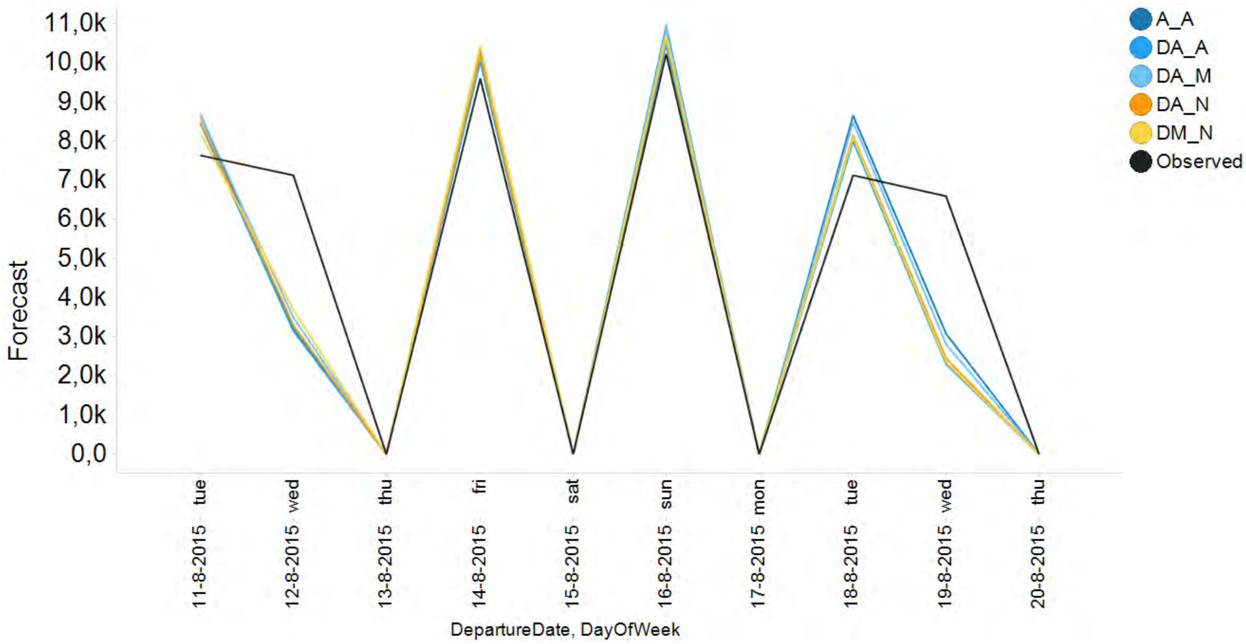

Figure 7.1: Flight KL0409 Amsterdam (Netherlands) - Almaty (Kazakhstan)

7.1.1 Raw versus decomposed time series

As mentioned in Chapter 5 (see Section 5.1.3) a time series will always be deseasonalised upfront, by using a decomposition method. To measure the impact of this deseasonalizing of the time series upfront, an analysis has been performed between forecasting *raw* and *decomposed* time series. As an example, flight KL0884 which flies from Xiamen to Amsterdam has been considered. The results of both forecasts are shown in Figures 7.2 and 7.3. Between brackets the best forecasting method for the flow has been stated. For this comparison we used a default scenario: 12 months of history, a guess value of 0.5 and the WAPE as error KPI.

Considering the operating days of flight KL0884 itself, it seems that this flow shows seasonal behavior. Therefore, both forecasts are corrected for the non-flying weekdays (Tuesday, Thursday, Saturday and Sunday). Looking at the overall forecast of both methods, it seems that the forecast for the Wednesday is most of the time too low. An analysis on the historical observations of all Wednesdays of the last 12 month shows that this particular flight only operates on Wednesdays during the summer season (see Figure C.1, Appendix C).

Looking at the accuracy of the forecasts, just by looking at the figure, it seems that the forecast based on the decomposed time series performs better. To compare the both approaches, the WAPE values of the best performing forecasting methods are compared. The results can be found in Table 7.1. In this example, the forecast is more accurate for the time series that first has been decomposed upfront. Because the same analysis for other flows with more variation in the amount of transported cargo showed similar results (see Figures C.2 and C.3, Appendix C), we decided to always decompose the time series upfront.

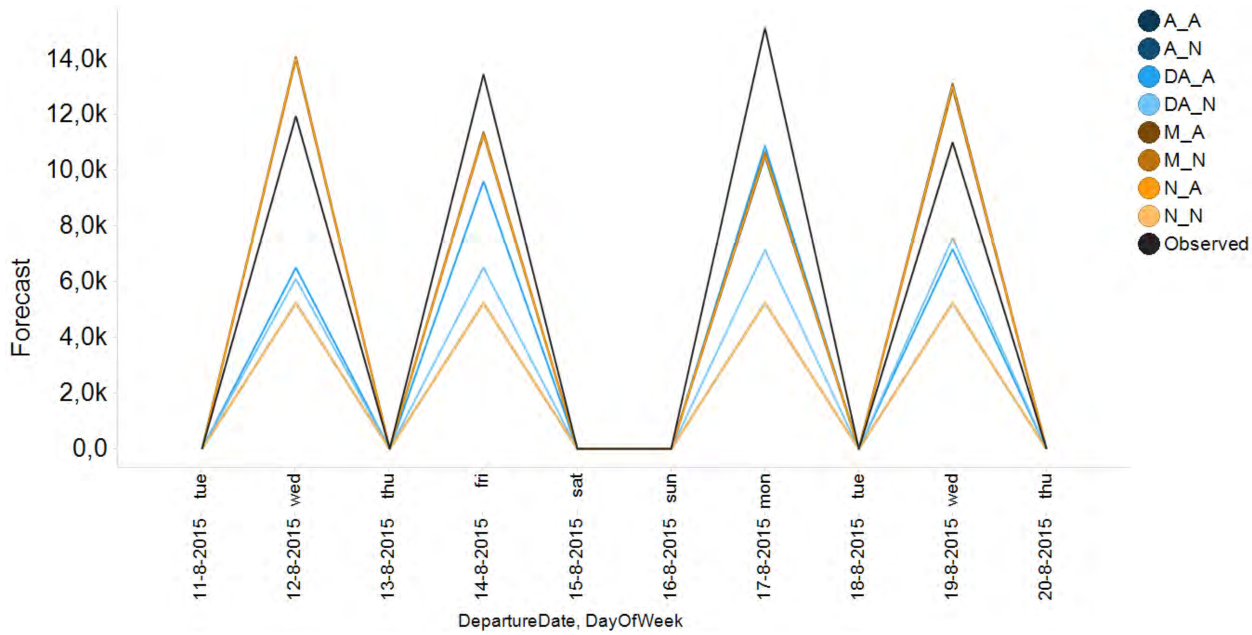

Figure 7.2: Raw (M_A): Flight KL0884 Xiamen (China) - Amsterdam (Netherlands)

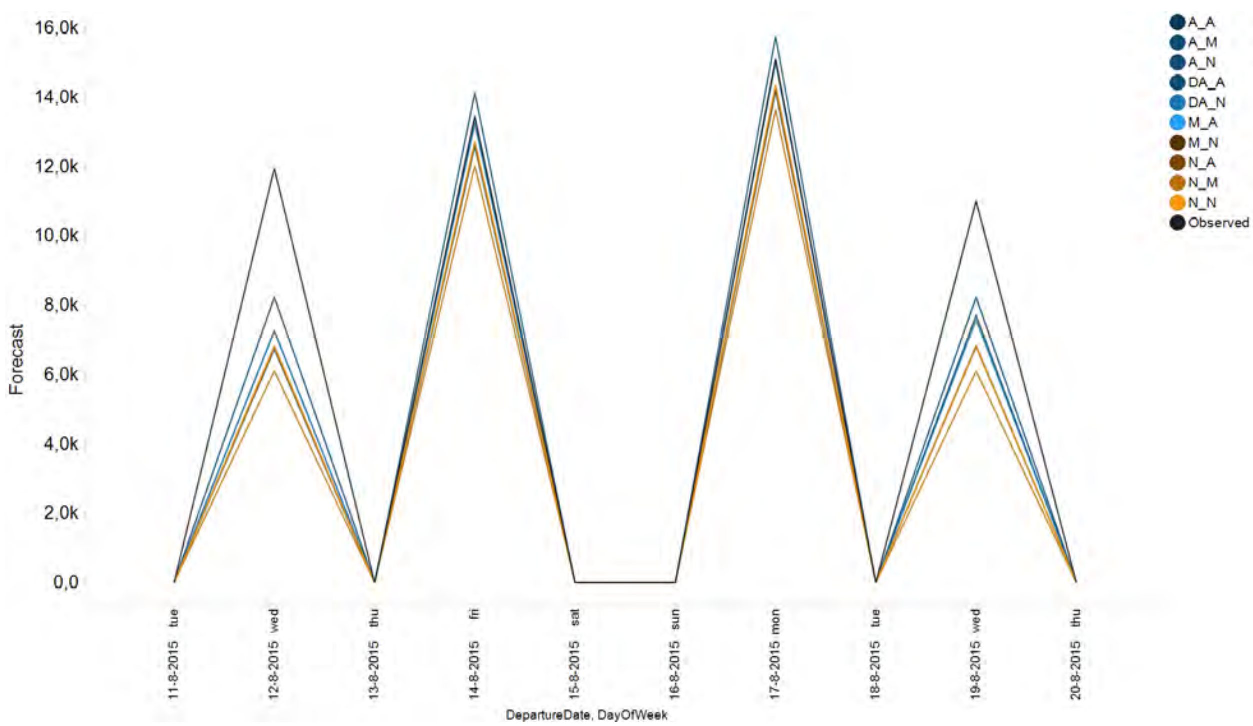

Figure 7.3: Decompose (A_M): Flight KL0884 Xiamen (China) - Amsterdam (Netherlands)

Time series	Best method	WAPE
Raw	M_A	20.36%
Decomposed	A_M	15.2%

Table 7.1: Raw versus decomposed time series

7.1.2 Scenario Analysis

To be able to determine which forecasting technique in combination with what *scenario* (see Section 5.1) performs best when forecasting a particular flight, different scenarios and forecasting methods have been tested on the period of 11 - 20 August, 2015. The resulting forecasts will be scored based on their forecast accuracy. Finally, the best scenario and forecasting model will be used to forecast the next 10 days. In Figure 7.4 the best forecast are shown for the for flight KL0597 which flies from Amsterdam (Netherlands) to Cape Town (South Africa). Furthermore, the actuals for that period are shown in the figure as well. The top 20 best results from this analysis, colored by error rate brackets of 0.1%, can be found in Table 7.2.

For this particular case, the best forecasting model is an exponential smoothing model with a damped additive trend and an additive seasonal component, considering 12 months of history, the RMSD as error KPI to estimate the exponential smoothing parameters starting with an initial guess value of 0.50. The WAPE corresponding to this best scenario/forecasting method combination can also be obtained from the table and has a value of 9.76%. Furthermore, the WAPE values of the top 20 best scoring methods are within a range of 0.5% (the error range runs from 9.76% - 10.18%).

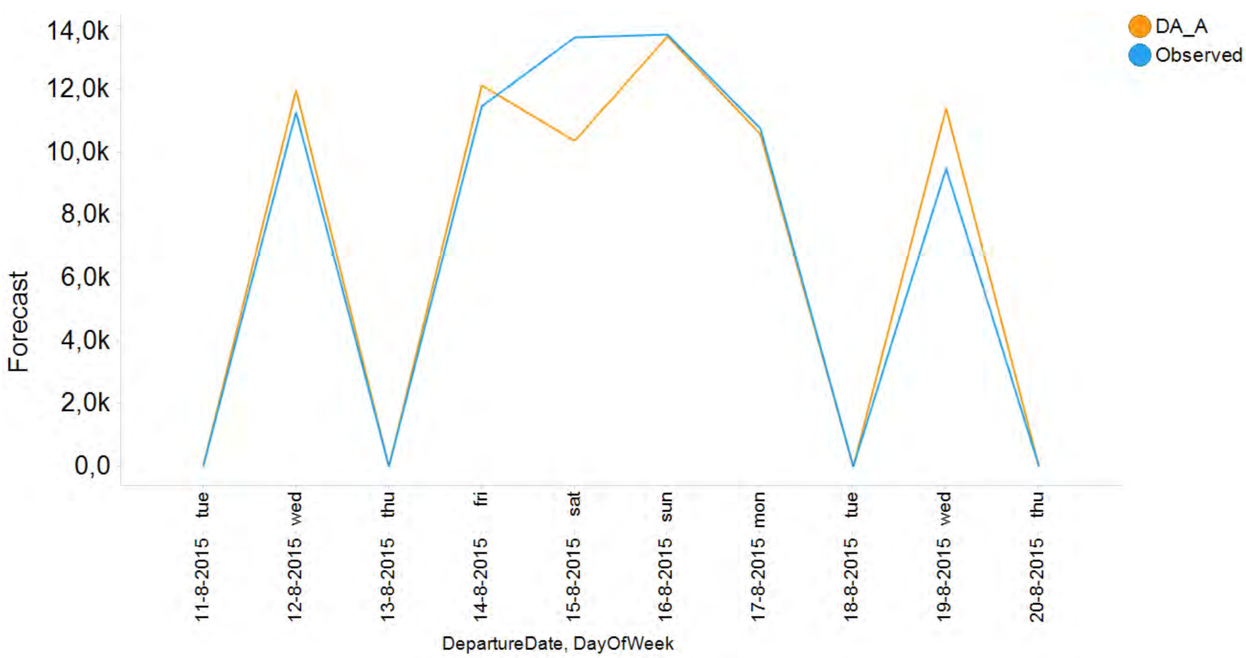

Figure 7.4: 'Best' Forecast vs Actuals - Flight KL0597: Amsterdam (Netherlands) - Cape Town (South Africa)

Best Method	Amount of History	Error KPI	Initial guess value	Error Rate (WAPE)
DA_A	12	RMSD	0.50	9,76%
DA_A	9	WAPE	0.30	9,76%
DA_A	12	WAPE	0.80	9,81%
A_N	12	SMAPE	0.30	9,87%
M_N	12	SMAPE	0.30	9,87%
N_N	12	SMAPE	0.30	9,87%
A_N	12	MAPE	0.80	9,89%
M_N	12	MAPE	0.50	9,89%
M_N	12	MAPE	0.80	9,89%
N_N	12	MAPE	0.30	9,89%
N_N	12	MAPE	0.50	9,89%
N_N	12	MAPE	0.80	9,89%
A_N	12	MAPE	0.30	9,91%
A_N	12	MAPE	0.50	9,93%
M_N	12	MAPE	0.30	9,93%
DA_A	12	RMSD	0.30	9,96%
DA_A	12	RMSD	0.80	10,01%
M_A	12	RMSD	0.80	10,03%
N_N	12	RMSD	0.50	10,17%
N_N	12	WAPE	0.30	10,18%

Table 7.2: Scenario analysis KL0597: Amsterdam (Netherlands) - Cape Town (South Africa)

Case studies

To get a feeling for the impact of using different forecasting methods, we performed some case studies to see which method would perform best and in which situation. This analysis should give an indication of which forecast model would perform best for different types of time series (highly seasonal and more variable time series). Three different cases are identified below. For each case the corresponding time series has been plotted (orange line represents the realized weight for the forecasting period that has been considered). Also, a subset of the best forecast methods are plotted against the observed value.

Case 1: 'new' flows - DL0230: Boston (United States) - Amsterdam (Netherlands)

Considering the flight history of flight DL0230, we can observe that this flight just started flying from the beginning of July 2015. Looking at the behavior of the time series, there is no direct seasonal pattern observable (see Figure 7.5). Therefore we expect that the best performing model will be a model considering at most 3 months of history and incorporating no seasonal component. From the WAPE values of the different forecasting methods, we can conclude that the best performing model is an exponential smoothing model with a multiplicative trend and (as expected) no seasonal component. The corresponding WAPE value is 34,71% (with a top 20 error range running from 34.71% - 38.09%).

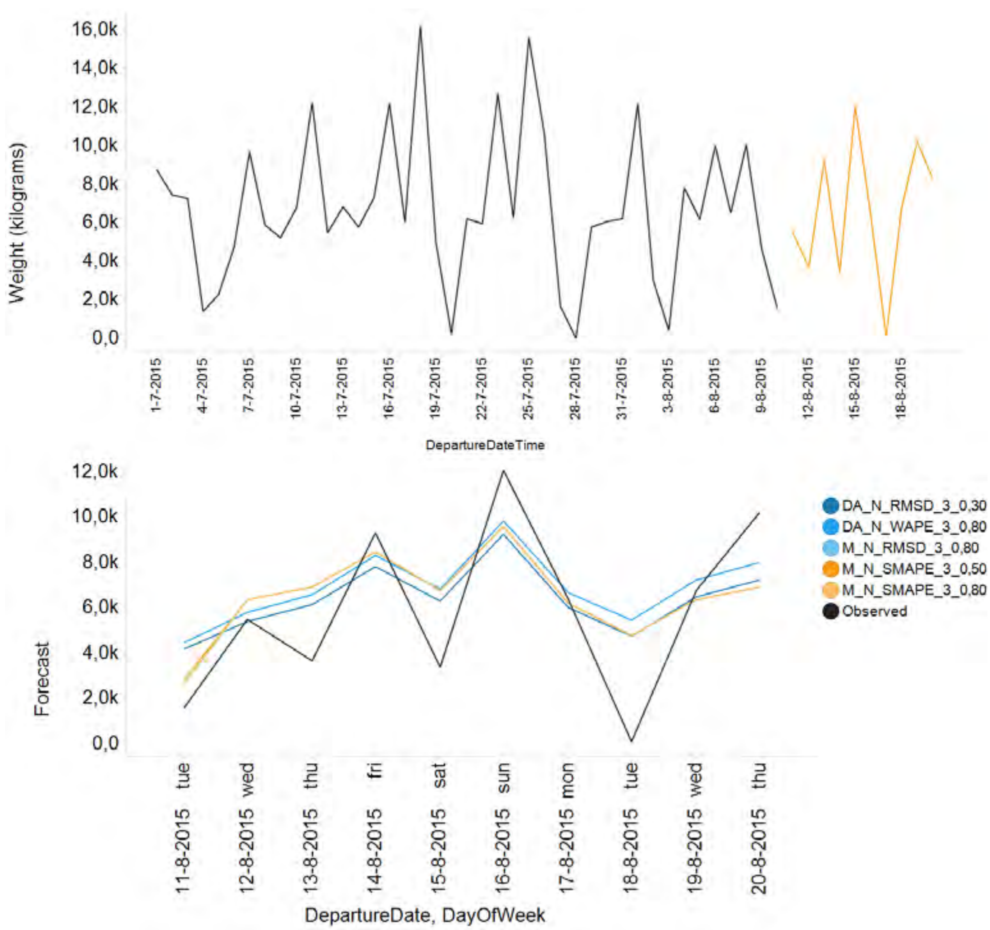

Figure 7.5: Flight DL0230 Boston (United States) - Amsterdam (Netherlands), M_N - WAPE of 34.71%

Case 2: seasonal flow - KL0574 Dar es Salaam (Tanzania) - Amsterdam (Netherlands)

Looking at flight KL0574, the historical flight data contains different patterns during the year. Until March, the time series shows a kind of random behavior. From March until the second half of June, it seems that a more seasonal pattern shows up. Towards the end of the time series, the distances between the peaks from this seasonal pattern are shrinking. An explanation for this could be that the frequency of the flight changes and more flights have been carried out. Because the pattern, that seems to change during the year, we would expect that the best performing forecasting method should take into account a small amount of history and a seasonal component. This expectation has been confirmed by the forecasting figure, in combination with the WAPE scores, from which can be obtained that an exponential smoothing model with both a multiplicative trend as well as a multiplicative seasonal component performs best with a WAPE value of 18.85% (with a top 20 error range running from 18.85% - 38.16%).

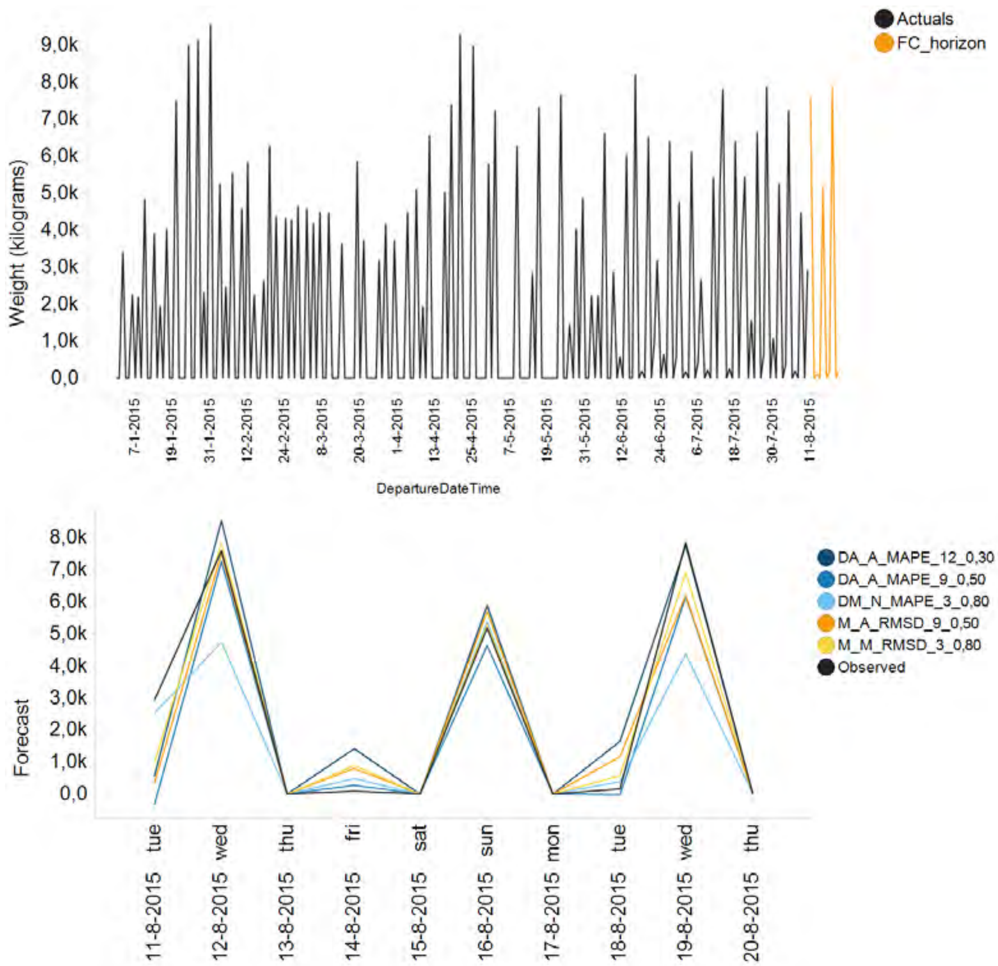

Figure 7.6: Flight KL0574 Dar es Salaam (Tanzania) - Amsterdam (Netherlands), M_M - WAPE of 18.85%

Case 3: more variable behaving flow - KL0438 Abu Dhabi (United Arab Emirates) - Amsterdam (Netherlands)

For the last case, flight KL0438 has been considered, a flight that is operated throughout the whole year (see Figure 7.7). Looking at the time series of the transported cargo weight, this time series has a highly variable behavior. No real seasonal component can be detected visually. Therefore, we assume that a forecasting model without seasonal component would perform best. Inspection of the figure also suggests that there is a rapidly fluctuating, non constant trend. So probably the best performing forecast model does not contain a trend component. From the computed WAPE values per forecasting method, we can obtain that an ARMA model fitted by using an Akaike Information Criterion (AIC) (number of AR and MA lags of respectively 4 and 2) and with a *non constant* trend component performs best (WAPE value of 9.46%). The error range of the top 20 best performing methods runs from 9.46% - 11.27%. Looking at the amount of history that has been taken into account by the model, a history of 6 months performs best.

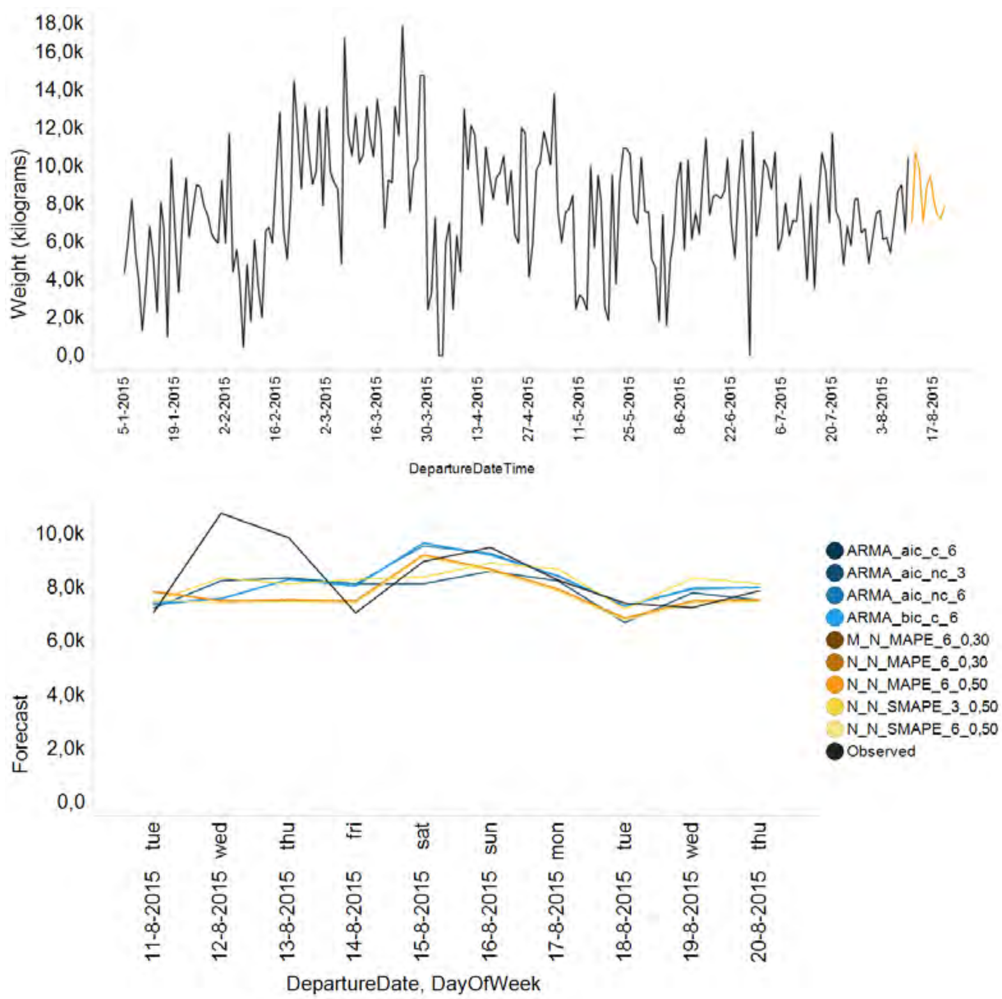

Figure 7.7: Flight KL0438 Abu Dhabi (United Arab Emirates) - Amsterdam (Netherlands), ARMA aic nc - WAPE of 9.46%

Case 4: counter example flow - KL0438 Abu Dhabi (United Arab Emirates) - Amsterdam (Netherlands)

From the previous cases, we can not conclude that it is always possible to determine upfront what kind of model would perform best for the different kinds of time series. It could always be that there are hidden dynamics in the time series that are better explained by using a completely different kind of forecasting method as expected. Also, it can happen that the when forecasting 'new' data another conclusion about the best performing forecasting technique can be drawn. To illustrate that it is not always the case that upfront we can determine what forecast technique will perform best, a counter example has been showed in Figure 7.8. In this example, a highly seasonal flight (from Mexico City to Amsterdam) has been best forecasted with an ARMA(4,2) model with constant trend (fitted by using an AIC criterion). Furthermore, 12 months of history were taken into account by this model, while at this period this flight was only operating during the last 2 months of 2014 and since the beginning of July 2015.

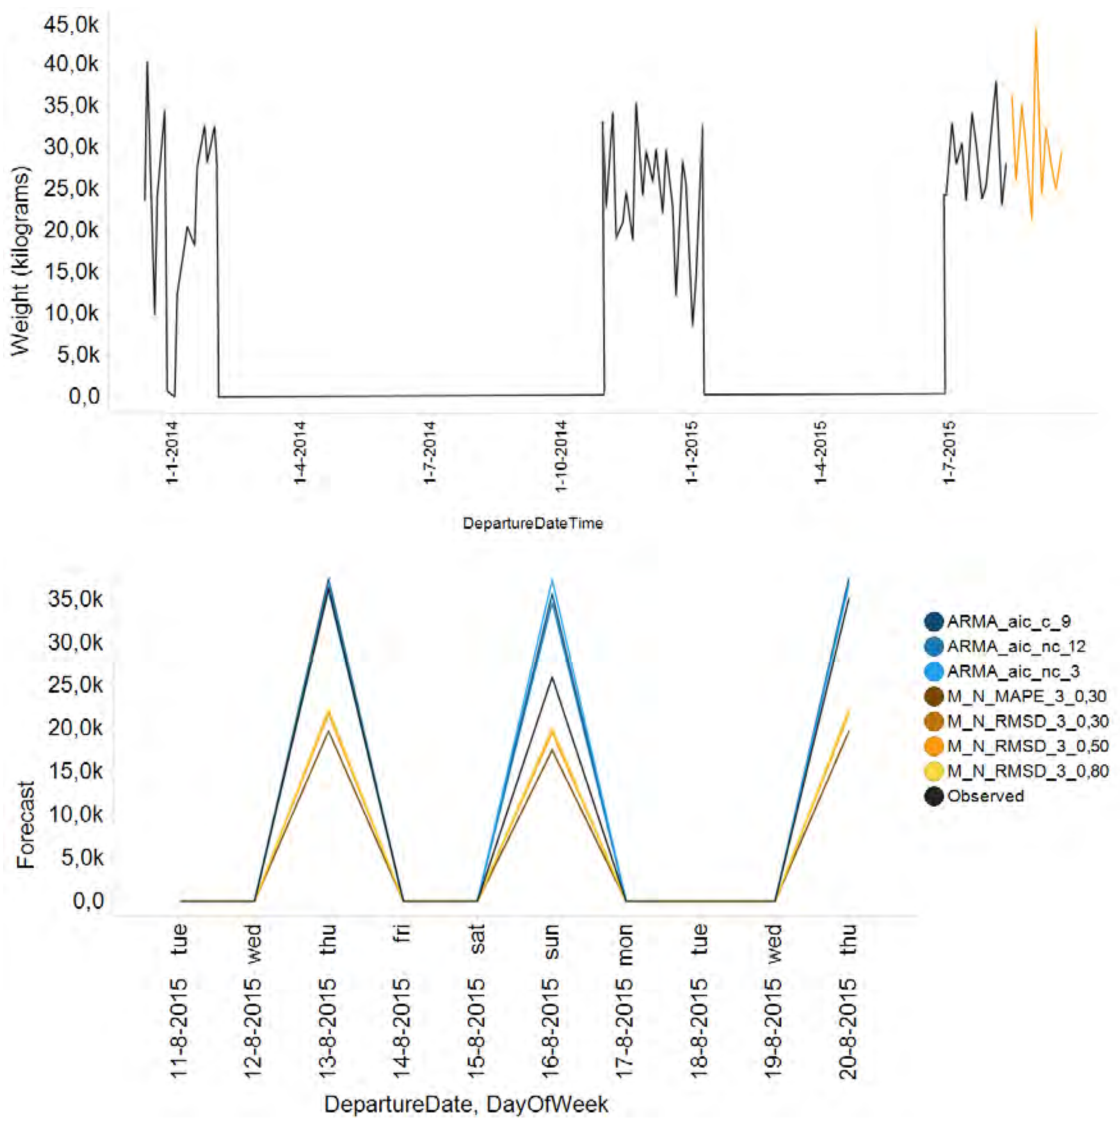

Figure 7.8: Flight KL0687 Amsterdam (Netherlands) - Mexico City (Mexico), ARMA aic c - WAPE of 10.68%

7.1.3 Aggregated flight forecast

It is obvious that some single flight forecasts are more accurate than others. Aggregating the single flight forecasts might result in an improvement of the overall forecast accuracy. Considering the full forecasting scope, looking at all flights that have flown between 11th and 20th of August, first a frequency table can be made of the overall best performing forecasting methods, see Table 7.3.

Model	Specifications	Frequency
ARMA	ARMA_aic_nc	7
	ARMA_aic_c	1
	ARMA_bic_c	1
Exponential Smoothing	M_M	53
	DA_M	19
	N_M	14
	DA_N	12
	A_M	11
	M_A	10
	DA_A	9
	N_N	8
	A_A	6
	DM_N	5
	N_A	4
	A_N	3
	DM_M	3
	M_N	3
	DM_A	1

Table 7.3: Frequency table of the best performing methods

Considering those results, we can easily conclude that, by far, most of the time an exponential smoothing model has been used (in 161 out of the 170 cases). Looking at the resulting forecast when aggregating all single flight forecasts, the total forecast can be obtained (see Figure 7.9) with a WAPE score of only 2.53%.

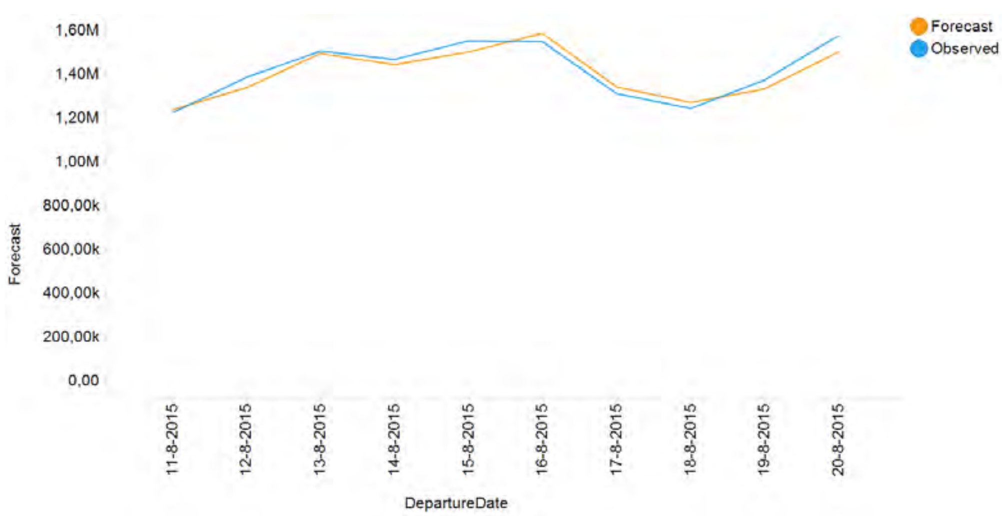

Figure 7.9: Total day level forecast, WAPE of 2.53%

The research question of this thesis is to create a forecast per freight building. Thus, the total day forecast has to be split by dividing the flows in *Arrival AMS* and *Departure AMS* flows. The resulting forecasts for both freight buildings are shown in Figure 7.10, where the *Arrival AMS* flights are assigned to freight building 2 and the *Departure AMS* to freight building 3 (see also Section 2.3.1, Figure 2.5). The corresponding WAPE values for both freight building forecasts can be found in Table 7.4. As we can observe from the table, the WAPE for the forecast on day level per freight building are both low, 4% and 2% for freight buildings 2 and 3, respectively. Looking back at Figures 6.2 and 6.3, a possible explanation for the difference in forecast accuracy between the freight buildings 2 and 3 can be observed. Looking at the outgoing cargo (freight building 3) a stronger relationship between transported tonnes of cargo and the total capacity used (the number of aircrafts) is shown, where incoming cargo (freight building 2) shows a more or less stable pattern regardless of the capacity used. It could be that the forecasting models that are used are better able to forecast the amount of outgoing cargo. We conclude that the forecast on flight level could support the Cargo Operations department estimating the expected workload.

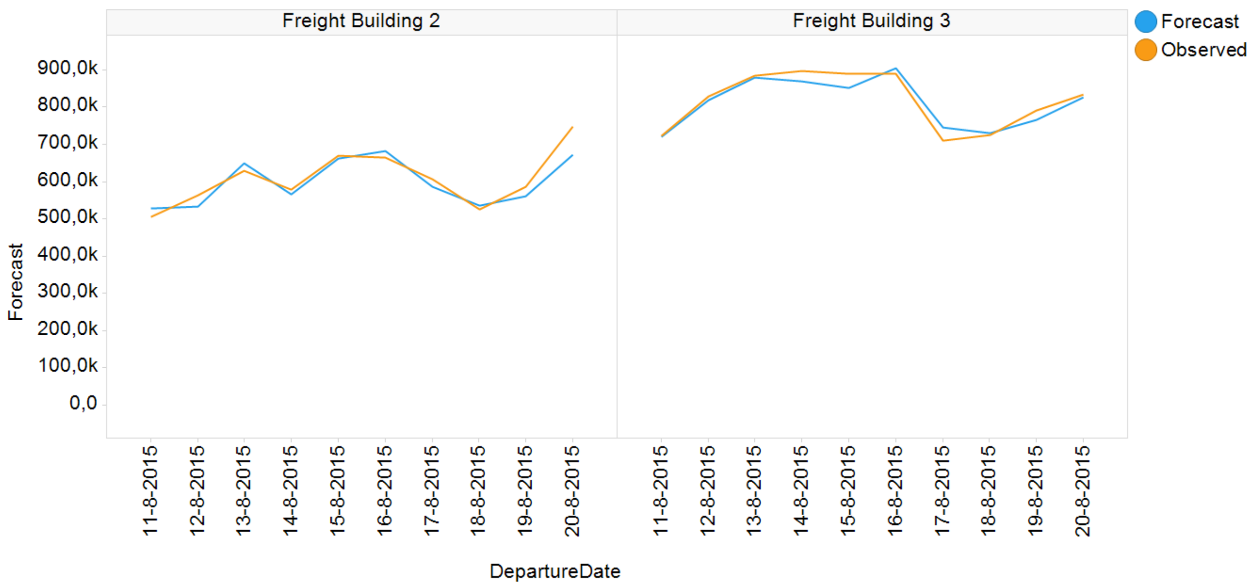

Figure 7.10: Total day level forecast per freight building

Freight Building	WAPE
Freight Building 2 (‘Europort’)	4.08%
Freight Building 3 (‘World Port’)	2.09%

Table 7.4: Forecast accuracy per freight building forecast

7.2 Booking forecast

For the forecast on booking level, first a linear regression model has been estimated for the Bookings at hand versus the materialization of those bookings. Because from the data analysis (see Figure 6.5) we observed that there is a significant difference between arriving and departing amount of cargo, this analysis has been performed for *board point* (outbound flights containing cargo boarded in Amsterdam) and *off point* (inbound flights containing cargo that is offloaded in Amsterdam) to be able to distinguish the incoming and outgoing cargo flows.

7.2.1 Bookings at hand versus materialization of the bookings

In Figure 7.11 the results of this linear regression for different DBDs (DBD 0 to 3) are shown. An analysis on all DBDs within the booking window can be found in Appendix C, see Figures C.4 and C.5.

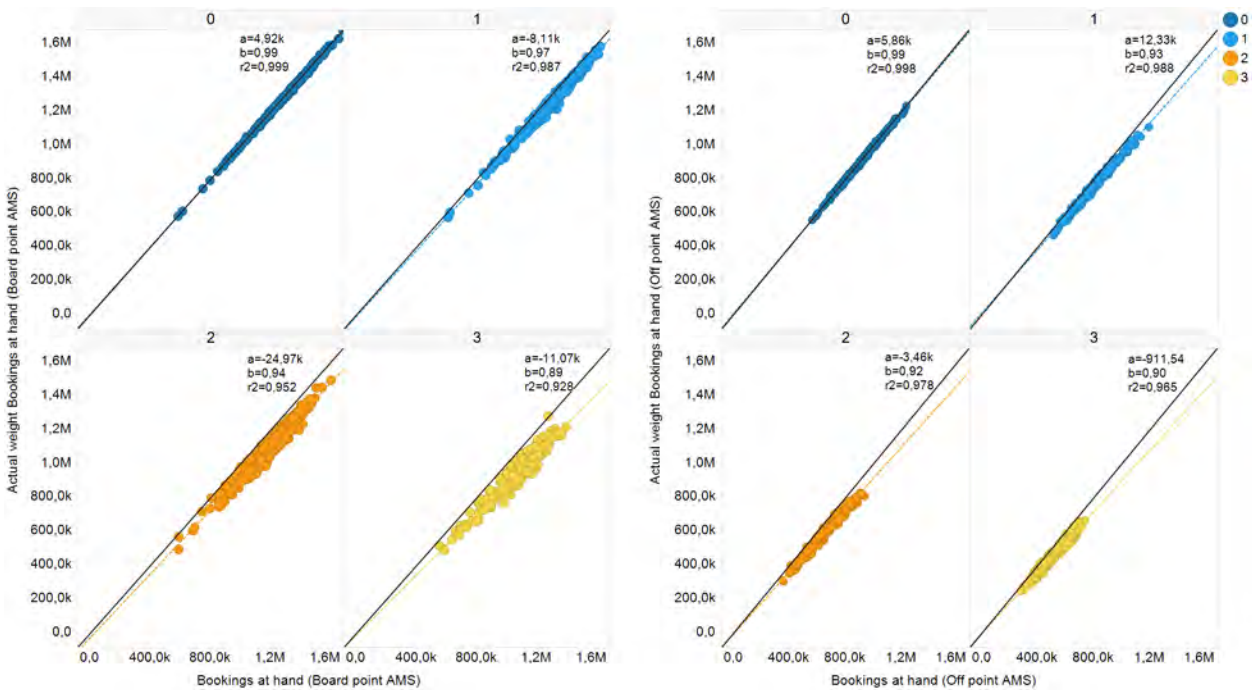

Figure 7.11: Estimated regression line bookings at hand versus materialization, split by board and off point AMS

Looking at the figure, which contains all booking from the start of 2015, several things can be observed. First of all, overall there is more outgoing cargo (cargo boarded in Amsterdam) than incoming cargo. Looking at the slopes (noted as value of b in the figure) as the day of operations approaches, it seems that the slope of the regression line gets closer to 1 for both board point as off point. This implies that the slope moves to the slope of the (black) diagonal y(x) = x. On the other hand, the intercept (noted as value of a in the figure) shows some intriguing behavior: it turns from negative long in advance to positive. The interpretation for this is that long in advance the model will overestimate the materialization of the bookings at hand. When the DBD 0 approaches, this overestimation of the model turns into an underestimation. Regarding the reliability of the estimated regression models, the statistical measure R^2 gives an indication of how well the regression line explains the data. The R^2 is defined as:

$$R^2 = \frac{\text{Explained variation}}{\text{Total variation}}$$

In other words, it gives an indication of how close the data are to the estimated line. The higher the R^2, the better the regression line fits the data. Looking at the R^2 values for the different DBDs and for the different flows (board point and off point AMS), those values seems to be quite high (all above the 90%). Therefore, we can conclude that the corresponding estimated regression lines fit the data very well. This implies that a couple of days before the day of operations, the materialization of the bookings can be forecasted quite accurately.

Booking data versus Opportunity Assessment To see whether Opportunity Assessment (OA) could improve the model, a comparison between the booking error and the OA error has been made. Here, the *error* has been defined as the actual materialization of the booking minus the booked weight or the forecasted weight by OA. In Figure 7.12 the aggregated OA error are plotted against the aggregated booking error for different DBD. Also, the (black) diagonal and the (red) estimated regression lines are plotted in the figure. On the horizontal axis we can see the error on booked weight as forecast, and on the vertical axis the error when using OA as forecast. A first thing that can be observed from the figure, is that a lot of negative errors implying that overall the sum of the booked weights is higher than the sum of the actual materialization (so in general people are overbooking). When looking at the Figure we should keep in mind that only data for the months August and September are presented, due to limited availability of OA data.

Looking at the (red) regression lines, we can interpret these as follows: given a certain error you would make using the booked weight as forecast, you can compute the error you would make using the OA forecast as forecast. Since the slope of all red lines are always decreasing, moving to the horizontal axis where the OA error is zero, it means that in general you are improving. This can also be concluded from the slope of the estimated regression lines, which is lower than the slope of the diagonal. Interesting however is to look at the intercept (see first numbers in the formulas of the estimated regression models in the figure): the interpretation for it is the error you will make using the OA forecast, when the error on booked weight as forecast is zero. The intercept is increasing from being negative long time before departure, to becoming positive as the day of departure approaches. Overall we can conclude, based on this analysis, that the OA model probably could improve the model. However, because of the limited amount of OA data, for now the decision has been made to use the booking data as input for the booking forecast model.

7.2.2 Bookings at hand versus Materialization of all Bookings

When having the forecast of the materialization of the bookings at hand for a certain DBD x, there still is a missing part of bookings that will be made after DBD x. That's why also an estimation is required of the weight of the remaining bookings. To create a forecast of the final weight of all bookings that will depart at the day of operations (hereafter called *'the total forecast'*), two different approaches have been examined: a *direct* and an *indirect* forecasting approach.

The *direct* forecasting approach directly creates the total forecast, using a linear regression model based on the booked weight of the bookings at hand and the materialization of <u>all</u> bookings (so of the bookings at hand plus the bookings which will be made after DBD x). Next to the direct approach, where we estimate the materialization of the bookings at hand as well as the materialization of the remaining bookings at once, also an approach has been examined that first estimates the materialization of the bookings at hand and then uses that prediction to estimate the materialization remaining bookings. This, so called, *indirect* forecasting approach consists of two steps. During the first step an estimation of the materialization of the bookings at hand is made as described in Section 7.2.1. In the second step a linear regression model based on the materialization of the bookings at hand and the materialization of all bookings will be used to create the total forecast.

Direct forecasting approach The direct forecasting approach creates the total forecast directly from the booked weight of the bookings at hand at a particular DBD. This has been done by creating a linear regression model based on the booked weight of the bookings at hand versus the actual weight of all bookings at the day of operations. The resulting models per DBD can be found in Figure 7.13. Furthermore, a split

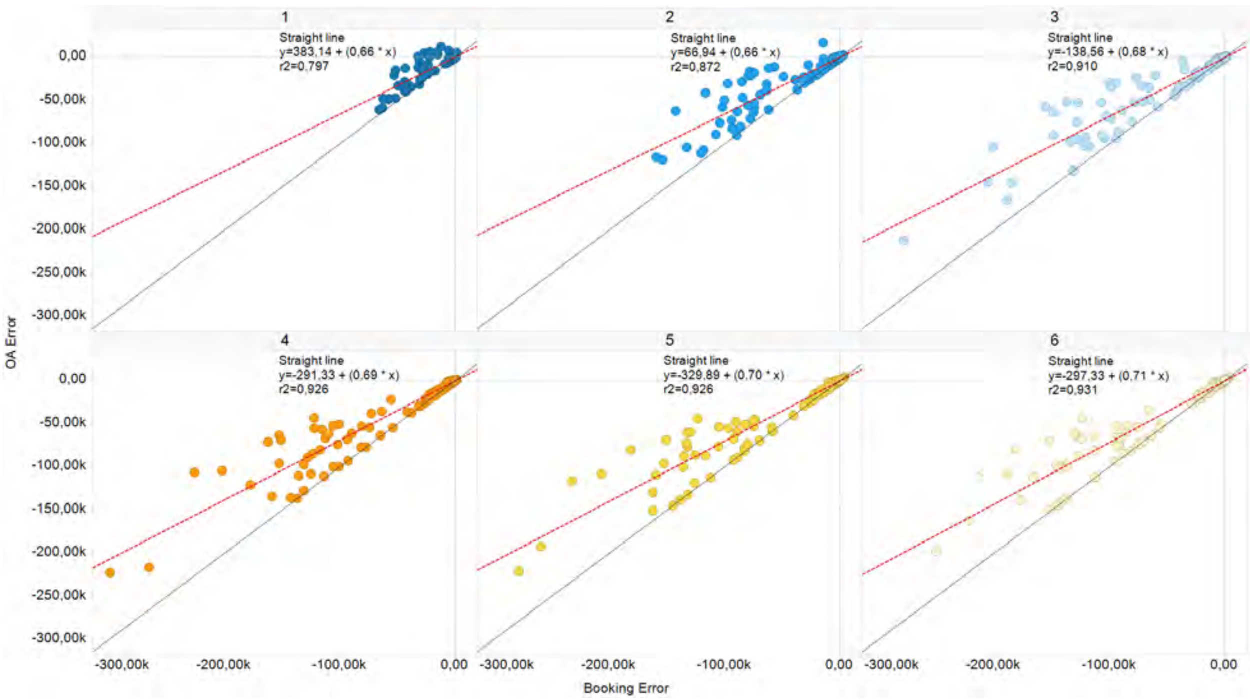

Figure 7.12: Booking error versus Opportunity Assessment (OA) error

has been made between board point AMS and off point AMS, the left and right figure respectively. The figures for all DBDs can be found in Appendix C (Figures C.6 and C.7).

First of all, it seems that, for each DBD the data points can be separated in two 'clouds'. For example when looking at DBD 3 in the off point AMS figure, there could be observed one cloud of points with an actual total weight of less than 250,0k and one with actual total weight values higher than 250,0k. Some further analysis showed that those 'clouds' belong to different product categories. An example of the product differentiation within the data can be found in the Appendix (Figure C.8). However, because the product differentiation is out of the scope of this research, we will not take into account the differences between regression lines for different products.

Secondly, looking at the development of the estimated regression lines as the day of departure approaches, the slope (noted as b in the figure) of the estimated regression models turns to the slope of the diagonal (which is of course equal to 1). This implies that as the day of departure approaches the bookings at hand gets closer to the total bookings. So, long in advance the portion of remaining bookings will be greater than shortly before the day of operations. This will result in a steeper regression line for DBDs long in advance, because an estimation has to be made of this large portion of remaining bookings. Looking at the intercept (noted as a in the figure), the same behavior can observed as before when looking at the bookings at hand versus the materialization of those bookings at hand: it goes from negative long in advance to positive as the departure day gets closer.

Finally, considering the R^2 of the linear regression models, the R^2 values increases for both board point AMS and off point AMS as the day of operations approaches. We can also observe that the R^2 values for board point AMS are higher than for off point AMS. Therefore, we can conclude that the models for board point AMS fit the data slightly better.

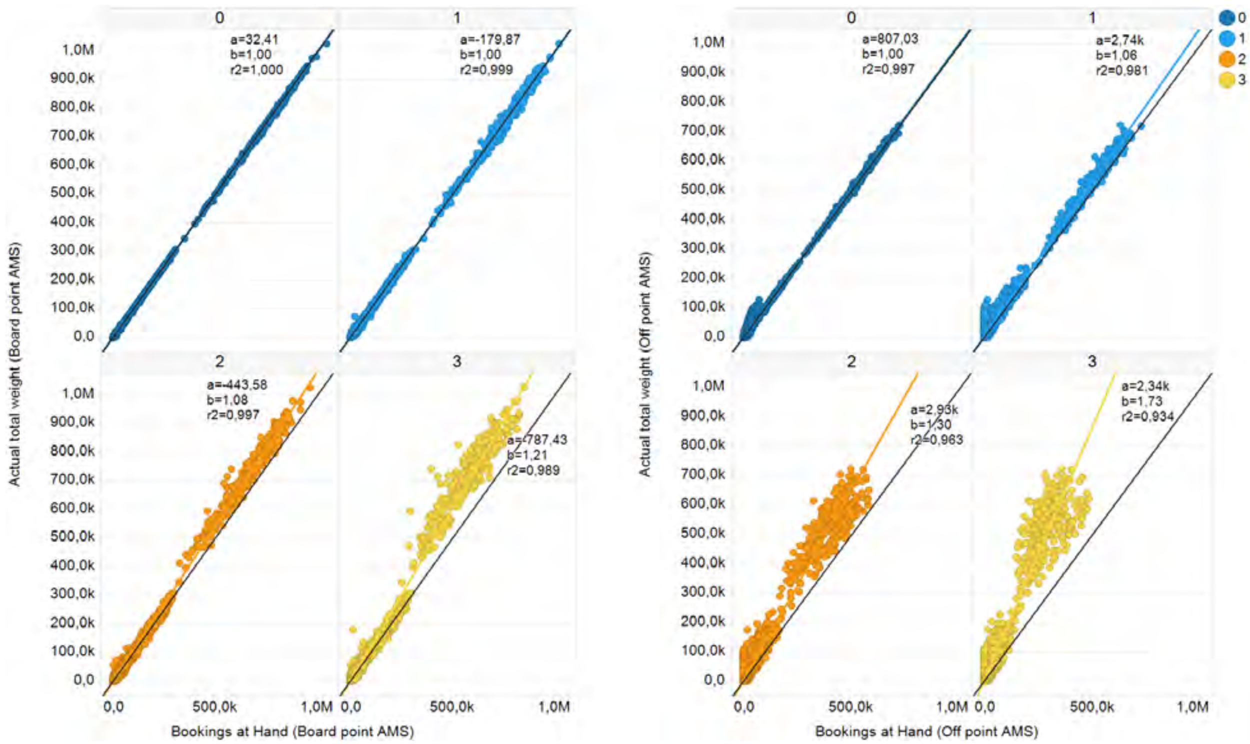

Figure 7.13: Estimated regression line bookings at hand versus total materialization (bookings made plus bookings to come), split by board and off point AMS

Now we have the estimated linear regression models for the different DBDs and different flows (arrival and departure), the total forecast can be created. Therefore, the sum of the booked weights of the bookings at hand will be substituted into the equation of the estimated regression line. The results of the total forecast can be found in Figure 7.14 and Table 7.5. From the resulting figures, we can observe that the forecasts created for board point AMS are more accurate in comparison with the off point AMS forecasts. This is to be expected, because of the higher R^2 values for the board point regression models. Although, the off point forecasts are getting worse when considering DBDs longer in advance, the overall performance (combine the both forecasts) seem to perform well. Therefore, depending on the degree of uncertainty that we would accept, the model could be used for forecasting the expected workload.

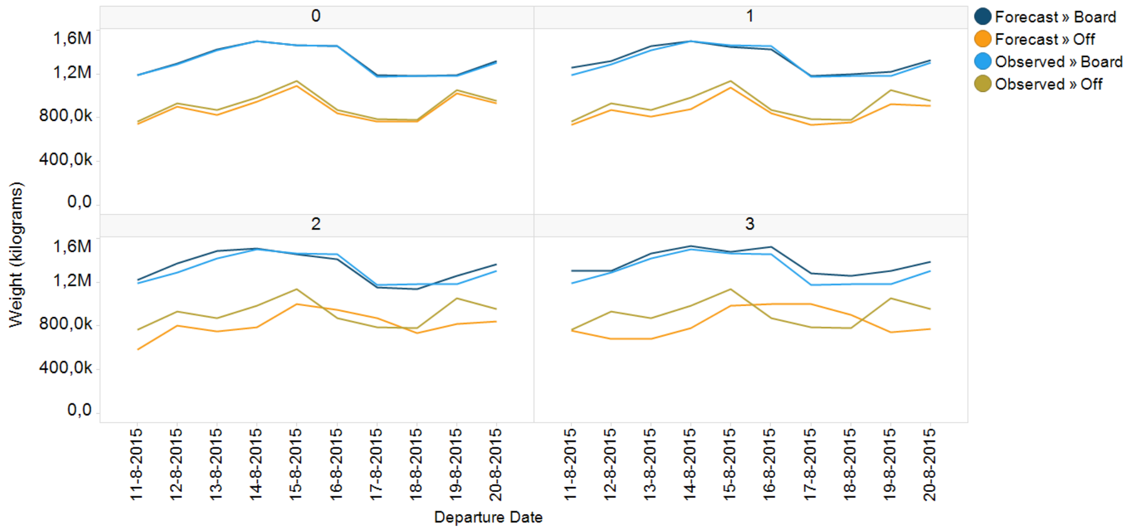

Figure 7.14: Direct forecasting approach

Days Before Departure (DBD)	Board point AMS WAPE	Off point AMS WAPE	Overall WAPE
0	0.43%	3.28%	1.09%
1	2.14%	6.45%	2.21%
2	3.42%	14.35%	4.36%
3	5.14%	19.18%	8.11%

Table 7.5: Performance direct forecasting approach

Indirect forecasting approach For the indirect forecasting approach we first predict how the bookings at hand will materialize at the day of operations (according to the method as described in 7.2.1). When we have a prediction for the materialization of the bookings at hand, we have to estimate what the weight of the remaining bookings to come will be. For this, linear regression models are created for the different DBDs and flows (board point AMS and off point AMS). The resulting models are shown in Figure 7.15 (all models can be found in Figure C.9 and C.10 in Appendix C). From those figures, we can conclude several things. The slope of the estimated models decreases as the day of operations approaches, this implies that during the booking period more bookings come in and we can predict more accurately the final amount of cargo that will arrive (hence the R^2 values are high (above the 90%) and also increases as the day of operations approaches). Also, as seen before in the direct forecasting approach, the goodness-of-fit of the models for board point AMS is better compared to the off point models. Therefore, we would expect that the board point AMS forecasts will be more accurate.

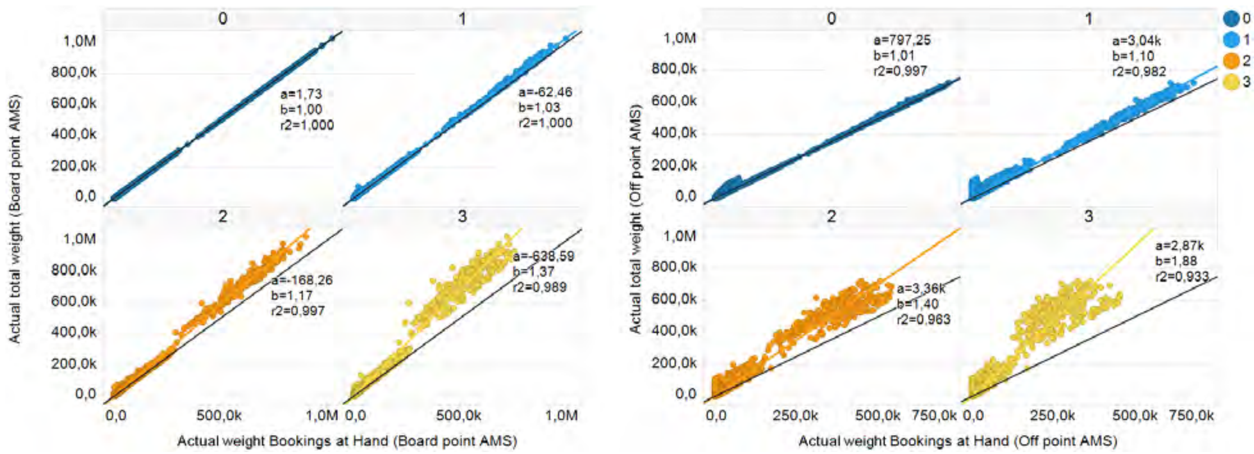

Figure 7.15: Estimated regression line materialization bookings at hand versus materialization of all bookings, split by board and off point AMS

Now we have the models to predict the materialization of all bookings based on the materialization of the bookings at hand, we can create the total forecast. Again, we showed the total forecasts for the board point AMS and the off point AMS in one figure (see Figure 7.16 and Table 7.6). From those results, we can conclude similar things as for the direct model. Comparing both approaches shows that for board point AMS the indirect forecasting approach performs better, while for the off point AMS the direct forecasting approach shows better results. When looking at the performance of both approaches of the overall forecast, the direct model performs best considering DBD 1, 2 and 3. Only for DBD 0 the indirect approach seems to be more accurate. Considering the aim of this research, a more accurate forecast shortly before the day of operations is more preferred (at DBD 0, no changes can be made to the resource planning).

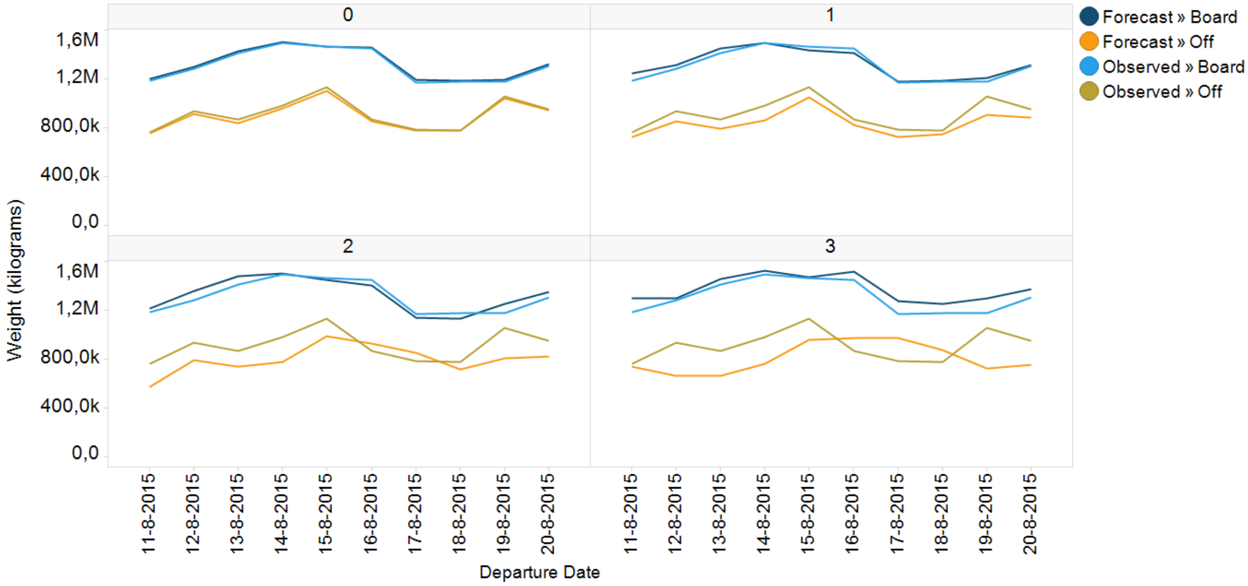

Figure 7.16: Indirect forecasting approach

Days Before Departure (DBD)	Board point AMS WAPE	Off point AMS WAPE	Overall WAPE
0	0.83%	1.77%	0.46%
1	1.75%	8.19%	3.15%
2	3.41%	15.1%	4.76%
3	4.8%	19.83%	8.24%

Table 7.6: Performance indirect forecasting approach

Chapter 8

Conclusions

This chapter provides the conclusions that can be drawn from this research.

The final results of this research provide an answer to the research question of this thesis:

Is it possible to create a reliable forecast of the expected cargo workload (within an error of 10%), split by freight building?

The answer to this question is: Yes, it is possible to create a forecast with an error less than 10%. For this, two approaches have been followed: one based on historical post-departure *flight* data and one on pre-departure *booking* data. Both approaches resulted in reliable forecasts that are consistent with each other with respect to the different freight buildings. To be able to measure the quality of the forecast results, a specific time period in the past has been forecasted, i.e. August 11 until August 20, 2015. The forecasted weights are compared with the actual values to measure the accuracy.

For the forecasting part based on historical flight data, we first created single flight forecasts. Next we split the single flight forecasts in incoming (freight building 2) and outgoing (freight building 3) flights. Finally we aggregated these forecasts and aggregated them. The resulting forecasts are shown in Figure 8.1. Those results show that it is indeed possible to create a reliable forecast per freight building based on historical flight data. Looking at the final forecasts for freight buildings 2 and 3, the accuracy of these forecasts were measured in terms of WAPE. This WAPE key performance indicator, weighs the forecast errors with respect to the level of actuals. Considering the WAPE values of those forecasts, the accuracy for freight buildings 1 and 2 was 4.08% and 2.09% respectively. For freight building 1, it was not possible to create a forecast, because of the missing European flight data in the data source that was used.

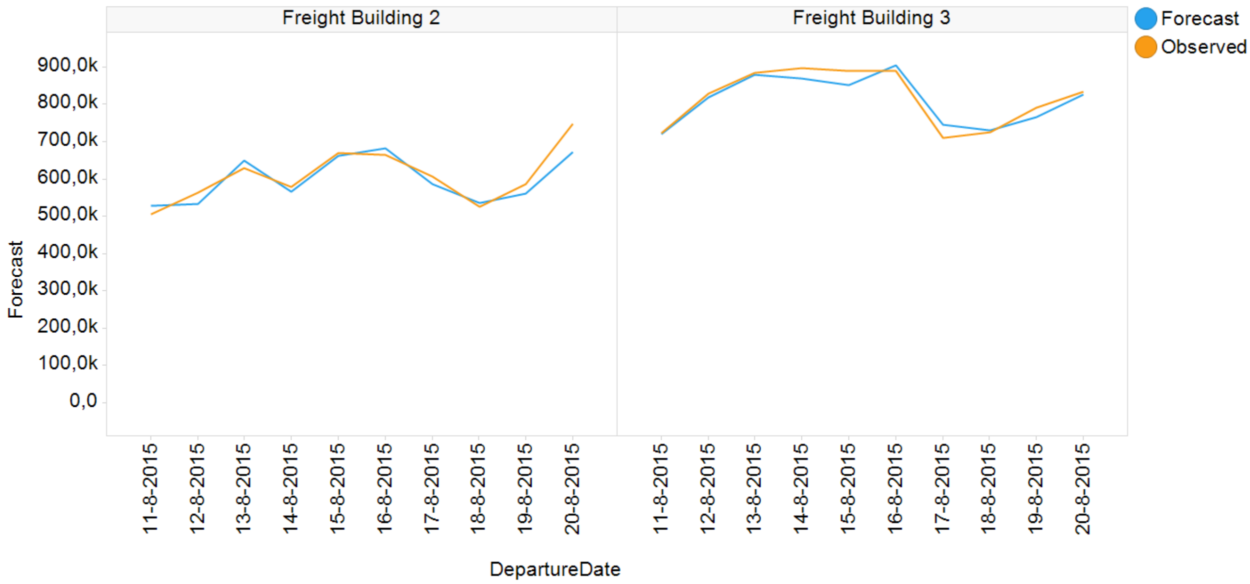

Figure 8.1: Total day level forecast per freight building

Looking at the forecasts based on pre departure booking data, two different approaches were examined: the *direct* and *indirect* forecasting approach. The direct approach uses the booked weight of the Bookings at hand to estimate the materialization of these bookings at hand as well as the materialization of the remaining bookings at once. On the other hand, the indirect approach first estimates the materialization of the bookings at hand and then uses that prediction to estimate the materialization remaining bookings. The final results were, in consultation with the business, split by incoming (off point AMS) and outgoing (boardpoint AMS) flights, because there were no clear business rules for splitting the total forecasts into freight building forecasts. Comparing the both approaches (direct and indirect) shows that for board point AMS the indirect forecasting approach performs better, while for the off point AMS the direct forecasting approach shows better results. Considering the overall performance of the forecasts created with the different approaches, the direct approach seems to perform slightly better in terms of accuracy (WAPE) for the 1,2 and 3 days before departure forecasts. Only for the DBD 0 forecast the indirect approach performs better. Because, the aim of the research is to be able to adapt the resource planning before the day of operations, we should use the forecasts created by using the direct approach to estimate the expected cargo workload. The resulting forecasts and corresponding forecast accuracy values can be found in Table 8.1 and Figure 8.2.

Days Before Departure (DBD)	**Board point AMS WAPE**	**Off point AMS WAPE**	**Overall WAPE**
0	0.43%	3.28%	1.09%
1	2.14%	6.45%	2.21%
2	3.42%	14.35%	4.36%
3	5.14%	19.18%	8.11%

Table 8.1: Performance direct forecasting approach

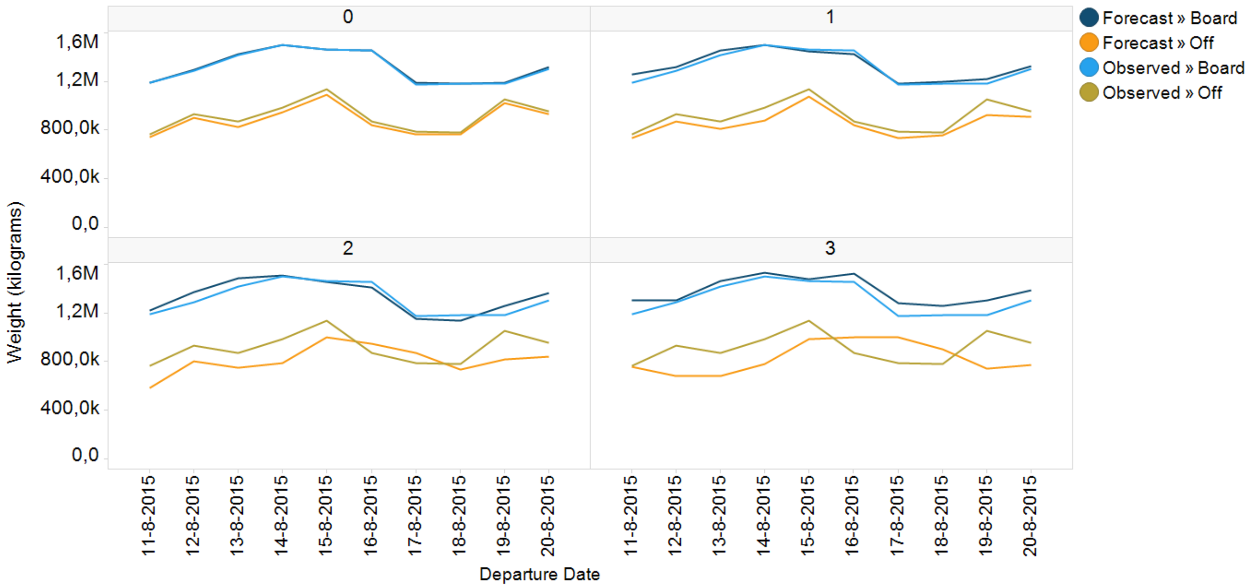

Figure 8.2: Direct forecasting approach

Comparing both forecasts (flight forecast versus booking forecast), it seems that accuracy of both forecasts is more or less equal shortly before the day of operations. Longer in advance, the forecast based on historical flight data is more accurate, but this is obviously due to the low number of booking that is made long in advance. Looking at the shape of both forecasts, the same behavior is observed.

Chapter 9

Discussion and Recommendation

This chapter provides discussion about the work that has been performed and some recommendations for further research.

9.1 Business recommendations

Measuring over and undercapacity To be able to measure the impact of implementing a solution (like using the results from this research), the feeling KLM Cargo has of having over/undercapacity should be quantified. Therefore, we would suggest the KLM Cargo division to measure the productivity of the employees, for example, by counting the number of processed ULDs or processed amount of cargo weight per time unit. With this information a 'standard' workpackage could be created, for example by creating a kind of Pareto curve where on the x-axis the amount of processed cargo (e.g. in terms of ULD) per resource per hour is stated and on the y-axis the service level of the cargo that has been processed on time. From such a figure, in combination with the final forecasted expected workload, then an indication of the number of resources needed can be determined.

Determine clear business rules for freight building split At the time of this research, we were not able to make a forecast per freight building, because there were no clear business rules for making this split. We recommend KLM Cargo to determine/create clear sets of business rules for this. This could be done by looking, for example, to the possibility of assigning specific product codes, flight numbers, flights from/to particular origin/destinations to the different freight buildings.

9.2 Data recommendations

Better registration of booking data Currently, the Commercial Product Code column in the ODS booking data files often contains the value 'Dummy Product Code'. People from the RM Cargo department mentioned that this is not a real existing product code, but a kind of default value when there is no product information available about the booking. Having more accurate product information of the booking upfront, might result in an improvement of the forecast accuracy.

Completion of the historical flight data Looking at the historical flight data that has been used, we found out that the European flights were missing. Because the majority of those flights contain cargo that has to be processed by freight building 1, the forecast for this freight building can not be made.

Flight versus booking data At the time of this research, we encountered that there was a difference between the flight and booking data source. First of all, it turns out that the flight data source does not contain European flight (EUR) flights. However, according to the business, this does not explain the significant difference in total tonnes of cargo that has been transported. From a second analysis, it turns out that the sum of all actuals on week level was much higher (around 50%) in comparison to the flight data and the weekly Cargo reports of the KLM Cargo division. Therefore, the origin, possible pre-process steps and precise content of the both data sources that have been provided should be made clear.

9.3 Model recommendations

Add Opportunity Assessment (OA) data to the booking forecast model From the comparison of the booking data (ODS) and the OA data, we concluded that the performance of the booking forecast model could be improved by using the OA forecast. However, because of the small amount of OA data (only 2 months) at the time of this research, we have chosen to restrict to the ODS booking data for the moment. When a larger amount of OA data is available, this data should be added to the model.

Extend the booking forecast model When estimating the linear regression models for the booking forecast, we only created separate models for incoming (off point AMS) and outgoing (board point AMS) flights. As we concluded from the Figures 7.13 and C.8, it might be interesting to see whether creating separate models for different feature splits (for example by product category, high/low season, a combination of different features, etcetera) could improve the accuracy of the model. Therefore, a linear regression model or decision tree model incorporating all features that might be interesting could be examined, to see what features are significant for the model.

Forecast on cargo load factor From the preliminary forecasting studies that have been reviewed, forecasting load factor instead of demand has several advantages, such as (partially) taking out the effect of aircraft capacity and schedule changes. During the forecasting process, a time series has been constructed based on the carrier, flight number, origin and destination information. However, it is known that some flights are performed with multiple aircraft types, and therefore different capacities, through the years. Due to this aircraft capacity changes, the constructed time series might show a very variable behavior which can be difficult to forecast. Therefore, it has to be (carefully) examined whether flights with different aircraft capacities can be better forecasted when considering the load factor.

Hierarchical forecasting From the literature, we found a forecasting methodology known as 'Hierarchical forecasting'. This way of forecasting combines forecasts of different levels (top-down and bottom-up approaches) in an optimal way. Hyndman, the founder of this way of forecasting, shows that a combination of forecasts can improve the accuracy of the forecast, because at different forecasting levels, the different dynamics of the time series that is forecasted are taken into account.

Flow characteristics From the flight forecasts, not only best forecast per flight is obtained, but also the best performing method with corresponding scenario settings. When many forecast runs have been performed, an analysis could be done to keep track of the characteristics of a particular flight. In other words, looking at the best performing method of each forecast run per single flight, could give some information of the characteristics of this flight. If it turns out that flights with the same kind of characteristics are better forecasted with specific kind of models (with a specific trend/seasonality component), then upfront one could choose to use this method (or set of methods) to forecast those flights.

Forecast stability The forecasts created for this internship, were only made for one period (11 - 20 August, 2015). To see whether for other moments in time the forecast has the same forecast accuracy, more forecasts have to be made. When there are more forecasts available, also an analysis can be performed to see how, for a particular departure date, the forecast accuracy changes over time. For example, for departure date 1 December, an analysis on the accuracy of the 10, 9,..., 1 days ahead forecasts for this departure date can be performed.

Abbreviations

AF	Air France.
AFKL	Air France KLM.
AFLS	Accenture Freight & Logistics Software.
AIC	Akaike Information Criterion.
ARIMA	Autoregressive Integrated Moving Average.
ARMA	Autoregressive Moving Average.
AWB	Airway Bill.
BB ULD	'Bijbouwer' Unit Load Device.
BIC	Bayesian Information Criterion.
CDG	Charles de Gaulle.
DBD	Days Before Departure.
DL	Delta Air Lines.
DoW	Day of Week.
EDF	Earliest Deadline First.
ES	Exponential Smoothing.
EUR	European flight.
FAP	Flight as Planned.
FAVAR	Factor-Augmented Vector Autoregressive Model.
FTK	Freight Tonne-Kilometers.
G1XL	Gare n1 Extra Large.
HOF	head operations flow.
IC	Information Criterion.
ICA	Intercontinental flight.
KL	KLM.
KLM	Koninklijke Luchtvaart Maatschappij.
KPI	Key Performance Indicator.
M-ULD	Mixed Unit Load Device.
MAPE	Mean Absolute Percentage Error.

OA	Opportunity Assessment.
OAG	Official Airline Guide.
OR	Operations Research.
PAX	Passenger.
PCHS	Pallet Container Handling System.
PRI	Planning, Roostering en Indeling.
RM	Revenue Management.
RMSD	Root Mean Square Deviation.
RTK	Revenue Tonne-Kilometers.
SLSQP	Sequential Least SQuares Programming.
SMAPE	Symmetric Mean Absolute Percentage Error.
SMPE	Symmetric Mean Percentage Error.
T-ULD	Through Unit Load Device.
ULD	Unit Load Device.
VAR	Vector Autoregressive Model.
WAPE	Weighted Absolute Percentage Error.

List of Terms

Airway Bill	A receipt of the contract of carriage, issued by an airline.
Altea	Product of Amadeus (company which develops software for the travel sector) to support the planning and loading of Passenger Flights.
Bijbouwer unit load device	A T-ULD that does have space left for some more shipments with the same destination.
Booking	Request for reservation of space on a flight for a specific date, which is (or has to be) confirmed by the airline.
Booking behavior	Describes how the bookings coming in over time during the entire booking window.
Bookings at hand	All bookings made up to a certain day. For example, the bookings at hand for DBD 3 are the bookings that have been made on DBD 3, 4, 5,....
Capacity	The total amount of volume/weight that an aircraft can carry.
Cargo load factor	Percentage of the available cargo space in an aircraft, that is used.
Code share	An marketing arrangement between airlines, which enables the option for an airline to sell tickets for a flight which is operated by another airline.
Colli	Package.
Coordinated Universal Time	The basis for civil time today. This time standard is based on the combination of atomic clocks with the Earth's rotation. Therefore it is highly precise.
Europort	Freight Building 2 at Amsterdam Schiphol Cargo Hub. Takes care of the handling of all cargo destined for trucks.
Freight tonne-kilometers	A measure of how much freight business an airline gets. One FTK is one metric tonne of revenue load carried one kilometer.
Hub	Airline hubs are airports that an airline uses as a transfer point to get passengers and freight to their intended destination.
Inbound	Incoming. Inbound flights on AMS are flights which arriving at Amsterdam Schiphol.
Manifest	Transport document which includes detailed information about the shipment.

Mixed unit load device	A ULD that does have to be broken down or build-up.
Outbound	Outgoing. Outbound flights from AMS are flights which are leaving from Amsterdam Schiphol.
Pallet and container handling system	A handling system for the fully automated handling of ready for carriage pallets and containers, from delivery to loading. PCHS offers 1200 storage positions.
Passage	Department of KLM which manages everything that has to do with passenger flights.
Show up rate	The percentage of the booked amount of cargo that will show up at the day of departure.
Through unit load device	A ULD that does not have to be broken down or build-up.
Traffic restriction code	An marketing arrangement between airlines, which enables the option for an airline to sell tickets for a flight which is operated by another airline.
Unit load device	A container or pallet for the transportation of cargo.
Workload	The number of employees needed to handle the arriving and departing cargo.
World Port	Freight Building 3 at Amsterdam Schiphol Cargo Hub. Takes care of the handling of all cargo destined for ICA flights.

Appendix A

Data

In this appendix a description is given of the different data sources that have been used.

A.1 Flight data

Column Name	**Description**
AircraftGroupCode	Aircraft Group Code
Arrival	Arrival city code
AWBVolAllo	AWB volume allotment bookings
AWBVolFreeSale	AWB volume free sales bookings
AWBWgtAllo	AWB weight allotment bookings
AWBWgtFreeSale	AWB weight free sales bookings
Departure	Departure city code
DepartureDate	Date of departure
FlightNumber	Flight Number
FlightNumberSuffix	Additional character that is added to the flight number when the flight deviates from its original flight plan
FlightPrefix	Carrier code
InvActualMain	Inventory cargo weight on main deck
InvActualVolume	Inventory weight of cargo
InvActualWeight	Inventory aircraft type
InvAcType	Inventory aircraft type
InvMainCapacity	Inventory weight capacity on main deck
InvVolumeCapacity	Inventory volume capacity
InvWeightCapacity	Inventory weight capacity
OpsActualCTR	Number of containers
OpsActualLDPositions	Number of filled ULD positions on the lower deck
OpsActualMain	Weight of all the main deck cargo
OpsActualMDPositions	Number of filled ULD positions on the main deck
OpsActualULDPositions	Total number of filled ULD positions
OpsActualVolume	Cargo volume in m3
OpsActualWeight	Cargo weight in kg
OpsAcType	Operating aircraft type
OpsCTRConfig	Total number of container positions
OpsEICPalletsWeight	Express cargo pallet weight in kg

OpsLDConfig	Total number of ULD positions on the lower deck
OpsMailWeight	Mail weight in kg
OpsMainCapacity	Weight capacity of the main deck
OpsMDConfig	Total number of ULD positions on the main deck
OpsULDConfig	Total number of ULD positions
OpsVolumeCapacity	Volume capacity of the aircraft
OpsWeightCapacity	Total weight capacity of the aircraft
Route	Flight route
SchAcType	Scheduled aircraft type
SchVolumeCapacity	Scheduled volume capacity
SchWeightCapacity	Scheduled weight capacity

Table A.1: Description of the flight data

A.2 OAG data

Column Name	Description
*Car	Carrier code
Arr	Arrival city code
Arrtm	Arrival UTC time
ArrtmLT	Arrival local time
Cnty Arr	Arrival country code
Cnty Dep	Departure country code
Dep	Departure city code
Deptm	Departure Coordinated Universal Time time
DeptmLT	Departure local time
Flt	Flight number
Freq	Flight frequency (day numbers)
Legnr	Leg number
OpCar	Operating carrier
OpFlt	Operating flight number
Seat	Number of seats
Traf	Traffic restriction code
Type	Aircraft type

Table A.2: Description of the OAG data

A.3 Booking data

Column Name	Description
AddlSCb	Additional shipment contribution, part of the revenue that can be contributed to the leg
AgentURN	Unique agent id
AircraftSubType	Aircraft sub type
Alloted	Indicator if the booking was allotted
AllotmentReference	Allotment booking reference
AllotmentReleased	Indicator if the allotment booking already released
AllotmentType	Allotment type
Arrival	Leg arrival city code
ArrivalCountry	Leg arrival country code
AWBNumber	AWB number
AWBPrefix	AWB prefix (carrier oriented)
BoardPoint	Board point city code
BoardPointCountry	Board point country code
BookingChannel	Booking channel
BookingKey	Booking key
BookingVersion	Booking version
BucketID	Bucket id
CancelIndicator	Indicator whether the booking is canceled
CargoReceiptFlag	Indicator if the cargo is already in the warehouse
CommercialProductCode	Commercial product code
CommodityCode	Commodity code
CommodityDescription	Commodity description
Departure	Leg departure city code
DepartureCountry	Leg departure country code
DepartureGMTOffset	Time difference between arrival and departure station
Destination	Destination city code
DestinationCountry	Destination country code
Dmin	Days before departure
FlightLegSequenceNumber	Leg number
FlightNumber	Flight number
FlightPrefix	Carrier code
LastUpdateDateTime	Last update date time booking
LegDepartureDate	Leg departure date
LegDepartureTime	Leg departure local time
LegPieces	Number of pieces of which the booking consists
LegStatus	Leg Status
LegVolume	Cargo volume in m3
LegWeight	Cargo weight in kg
OffPoint	Off point city code
OffPointCountry	Off point country code
Origin	Origin city code
OriginCountry	Origin country code
PartShipmentSuffix	Additional character that is added to the booking when the flight plan of the booking deviates from its original flight plan
ProductCode	Product code

SCb	Shipment contribution
SHC	Special handling codes
ShipperURN	Unique shipper id

Table A.3: Description of the booking data

Appendix B

Models

Below a table can be found where all standard exponential smoothing methods and their corresponding formulas are stated. Furthermore a list of definitions is given.

Symbol	Definition
α	Smoothing parameter for the level of the series
γ	Smoothing parameter for the trend
δ	Smoothing parameter for seasonal indices
ϕ	Autoregressive or damping parameter
β	Discount factor, $0 \leq \beta \leq 1$
S_t	Smoothed level of the series, computed after X_t is observed. Also the expected value of the data at the end of period t in some models
T_t	Smoothed additive trend at the end of period t
R_t	Smoothed multiplicative trend at the end of period t
I_t	Smoothed seasonal index at the end of period t. Can be additive or multiplicative
X_t	Observed value of the time series in period t
m	Number of periods in the forecast lead-time
p	Number of periods in the seasonal cycle
$\hat{X}_t(m)$	Forecast for m periods ahead from origin t
e_t	One-step-ahead forecast error, $e_t = X_t - \hat{X}_{t-1}(1)$. Note that $e_t(m)$ should be used for other forecast origins
C_t	Cumulative renormalization factor for seasonal indices. Can be additive or multiplicative
V_t	Transition variable in smooth transition exponential smoothing
D_t	Observed value of nonzero demand in the Croston method
Q_t	Observed inter-arrival time of transactions in the Croston method
Z_t	Smoothed nonzero demand in the Croston method
P_t	Smoothed inter-arrival time in the Croston method
Y_t	Estimated demand per unit time in the Croston method (Z_t/P_t)

Figure B.1: Exponential Smoothing Methods: list of definitions (Gardner, 1985)

Trend	Seasonality: N None	A Additive	M Multiplicative
N None	$S_t = \alpha X_t + (1-\alpha)S_{t-1}$ $\hat{X}_t(m) = S_t$	$S_t = \alpha(X_t - I_{t-p}) + (1-\alpha)S_{t-1}$ $I_t = \delta(X_t - S_t) + (1-\delta)I_{t-p}$ $\hat{X}_t(m) = S_t + I_{t-p+m}$	$S_t = \alpha(X_t / I_{t-p}) + (1-\alpha)S_{t-1}$ $I_t = \delta(X_t / S_t) + (1-\delta)I_{t-p}$ $\hat{X}_t(m) = S_t I_{t-p+m}$
	$S_t = S_{t-1} + \alpha e_t$ $\hat{X}_t(m) = S_t$	$S_t = S_{t-1} + \alpha e_t$ $I_t = I_{t-p} + \delta(1-\alpha)e_t$ $\hat{X}_t(m) = S_t + I_{t-p+m}$	$S_t = S_{t-1} + \alpha e_t / I_{t-p}$ $I_t = I_{t-p} + \delta(1-\alpha)e_t / S_t$ $\hat{X}_t(m) = S_t I_{t-p+m}$
A Additive	$S_t = \alpha X_t + (1-\alpha)(S_{t-1} + T_{t-1})$ $T_t = \gamma(S_t - S_{t-1}) + (1-\gamma)T_{t-1}$ $\hat{X}_t(m) = S_t + mT_t$	$S_t = \alpha(X_t - I_{t-p}) + (1-\alpha)(S_{t-1} + T_{t-1})$ $T_t = \gamma(S_t - S_{t-1}) + (1-\gamma)T_{t-1}$ $I_t = \delta(X_t - S_t) + (1-\delta)I_{t-p}$ $\hat{X}_t(m) = S_t + mT_t + I_{t-p+m}$	$S_t = \alpha(X_t / I_{t-p}) + (1-\alpha)(S_{t-1} + T_{t-1})$ $T_t = \gamma(S_t - S_{t-1}) + (1-\gamma)T_{t-1}$ $I_t = \delta(X_t / S_t) + (1-\delta)I_{t-p}$ $\hat{X}_t(m) = (S_t + mT_t)I_{t-p+m}$
	$S_t = S_{t-1} + T_{t-1} + \alpha e_t$ $T_t = T_{t-1} + \alpha\gamma e_t$ $\hat{X}_t(m) = S_t + mT_t$	$S_t = S_{t-1} + T_{t-1} + \alpha e_t$ $T_t = T_{t-1} + \alpha\gamma e_t$ $I_t = I_{t-p} + \delta(1-\alpha)e_t$ $\hat{X}_t(m) = S_t + mT_t + I_{t-p+m}$	$S_t = S_{t-1} + T_{t-1} + \alpha e_t / I_{t-p}$ $T_t = T_{t-1} + \alpha\gamma e_t / I_{t-p}$ $I_t = I_{t-p} + \delta(1-\alpha)e_t / S_t$ $\hat{X}_t(m) = (S_t + mT_t)I_{t-p+m}$
DA Damped Additive	$S_t = \alpha X_t + (1-\alpha)(S_{t-1} + \phi T_{t-1})$ $T_t = \gamma(S_t - S_{t-1}) + (1-\gamma)\phi T_{t-1}$ $\hat{X}_t(m) = S_t + \sum_{i=1}^{m} \phi^i T_t$	$S_t = \alpha(X_t - I_{t-p}) + (1-\alpha)(S_{t-1} + \phi T_{t-1})$ $T_t = \gamma(S_t - S_{t-1}) + (1-\gamma)\phi T_{t-1}$ $I_t = \delta(X_t - S_t) + (1-\delta)I_{t-p}$ $\hat{X}_t(m) = S_t + \sum_{i=1}^{m} \phi^i T_t + I_{t-p+m}$	$S_t = \alpha(X_t / I_{t-p}) + (1-\alpha)(S_{t-1} + \phi T_{t-1})$ $T_t = \gamma(S_t - S_{t-1}) + (1-\gamma)\phi T_{t-1}$ $I_t = \delta(X_t / S_t) + (1-\delta)I_{t-p}$ $\hat{X}_t(m) = (S_t + \sum_{i=1}^{m} \phi^i T_t)I_{t-p+m}$
	$S_t = S_{t-1} + \phi T_{t-1} + \alpha e_t$ $T_t = \phi T_{t-1} + \alpha\gamma e_t$ $\hat{X}_t(m) = S_t + \sum_{i=1}^{m} \phi^i T_t$	$S_t = S_{t-1} + \phi T_{t-1} + \alpha e_t$ $T_t = \phi T_{t-1} + \alpha\gamma e_t$ $I_t = I_{t-p} + \delta(1-\alpha)e_t$ $\hat{X}_t(m) = S_t + \sum_{i=1}^{m} \phi^i T_t + I_{t-p+m}$	$S_t = S_{t-1} + \phi T_{t-1} + \alpha e_t / I_{t-p}$ $T_t = \phi T_{t-1} + \alpha\gamma e_t / I_{t-p}$ $I_t = I_{t-p} + \delta(1-\alpha)e_t / S_t$ $\hat{X}_t(m) = (S_t + \sum_{i=1}^{m} \phi^i T_t)I_{t-p+m}$
M Multiplicative	$S_t = \alpha X_t + (1-\alpha)(S_{t-1}R_{t-1})$ $R_t = \gamma(S_t / S_{t-1}) + (1-\gamma)R_{t-1}$ $\hat{X}_t(m) = S_t R_t^m$	$S_t = \alpha(X_t - I_{t-p}) + (1-\alpha)S_{t-1}R_{t-1}$ $R_t = \gamma(S_t / S_{t-1}) + (1-\gamma)R_{t-1}$ $I_t = \delta(X_t - S_t) + (1-\delta)I_{t-p}$ $\hat{X}_t(m) = S_t R_t^m + I_{t-p+m}$	$S_t = \alpha(X_t / I_{t-p}) + (1-\alpha)S_{t-1}R_{t-1}$ $R_t = \gamma(S_t / S_{t-1}) + (1-\gamma)R_{t-1}$ $I_t = \delta(X_t / S_t) + (1-\delta)I_{t-p}$ $\hat{X}_t(m) = (S_t R_t^m)I_{t-p+m}$
	$S_t = S_{t-1}R_{t-1} + \alpha e_t$ $R_t = R_{t-1} + \alpha\gamma e_t / S_{t-1}$ $\hat{X}_t(m) = S_t R_t^m$	$S_t = S_{t-1}R_{t-1} + \alpha e_t$ $R_t = R_{t+1} + \alpha\gamma e_t / S_{t-1}$ $I_t = I_{t-p} + \delta(1-\alpha)e_t$ $\hat{X}_t(m) = S_t R_t^m + I_{t-p+m}$	$S_t = S_{t-1}R_{t-1} + \alpha e_t / I_{t-p}$ $R_t = R_{t-1} + (\alpha\gamma e_t / S_{t-1}) / I_{t-p}$ $I_t = I_{t-p} + \delta(1-\alpha)e_t / S_t$ $\hat{X}_t(m) = (S_t R_t^m)I_{t-p+m}$
DM Damped Multiplicative	$S_t = \alpha X_t + (1-\alpha)(S_{t-1}R_{t-1}^{\phi})$ $R_t = \gamma(S_t / S_{t-1}) + (1-\gamma)R_{t-1}^{\phi}$ $\hat{X}_t(m) = S_t R_t^{\Sigma_{i=1}^{m} \phi^i}$	$S_t = \alpha(X_t - I_{t-p}) + (1-\alpha)S_{t-1}R_{t-1}^{\phi}$ $R_t = \gamma(S_t / S_{t-1}) + (1-\gamma)R_{t-1}^{\phi}$ $I_t = \delta(X_t - S_t) + (1-\delta)I_{t-p}$ $\hat{X}_t(m) = S_t R_t^{\Sigma_{i=1}^{m} \phi^j} + I_{t-p+m}$	$S_t = \alpha(X_t / I_{t-p}) + (1-\alpha)(S_{t-1}R_{t-1}^{\phi})$ $R_t = \gamma(S_t / S_{t-1}) + (1-\gamma)R_{t-1}^{\phi}$ $I_t = \delta(X_t / S_t) + (1-\delta)I_{t-1}$ $\hat{X}_t(m) = (S_t R_t^{\Sigma_{i=1}^{m} \phi^j})I_{t-p+m}$
	$S_t = S_{t-1}R_{t-1}^{\phi} + \alpha e_t$ $R_t = R_{t-1}^{\phi} + \alpha\gamma e_t / S_{t-1}$ $\hat{X}_t(m) = S_t R_t^{\Sigma_{i=1}^{m} \phi^i}$	$S_t = S_{t-1}R_{t-1}^{\phi} + \alpha e_t$ $R_t = R_{t-1}^{\phi} + \alpha\gamma e_t / S_{t-1}$ $I_t = I_{t-p} + \delta(1-\alpha)e_t$ $\hat{X}_t(m) = S_t R_t^{\Sigma_{i=1}^{m} \phi^j} + I_{t-p+m}$	$S_t = S_{t-1}R_{t-1}^{\phi} + \alpha e_t / I_{t-p}$ $R_t = R_{t-1}^{\phi} + (\alpha\gamma e_t / S_{t-1}) / I_{t-p}$ $I_t = I_{t-p} + \delta(1-\alpha)e_t / S_t$ $\hat{X}_t(m) = (S_t R_t^{\Sigma_{i=1}^{m} \phi^j})I_{t-p+m}$

Figure B.2: Exponential Smoothing Methods (Gardner, 1985)

Appendix C

Results

In this appendix additional figures can be found.

C.1 Raw versus Decompose

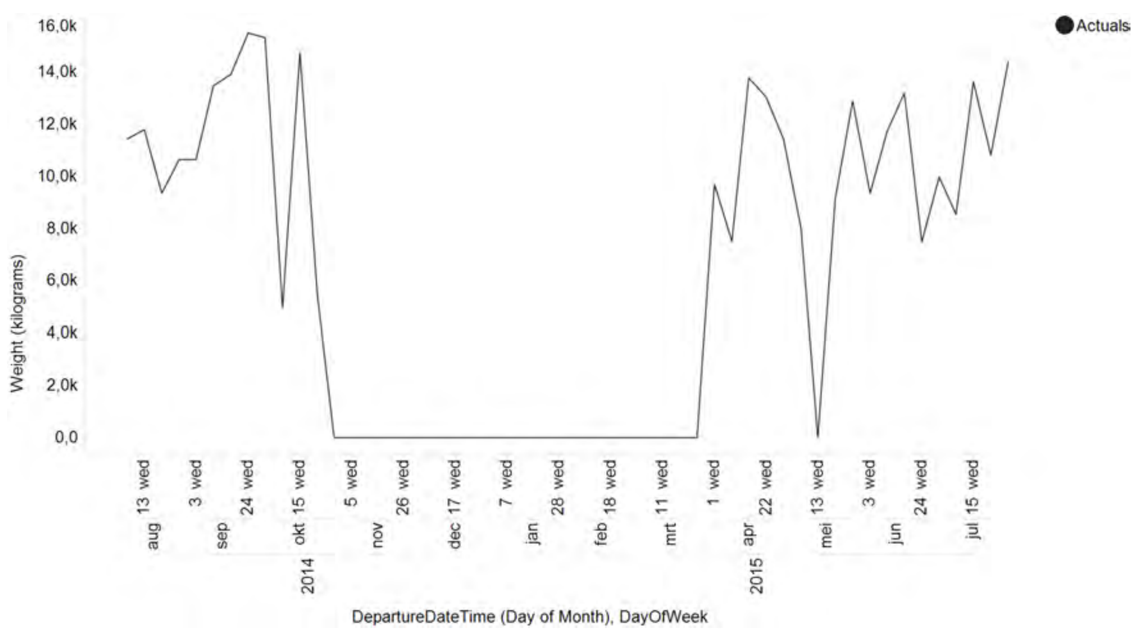

Figure C.1: Flight KL0884 Xiamen (China) - Amsterdam (Netherlands): Frequency change Wednesdays

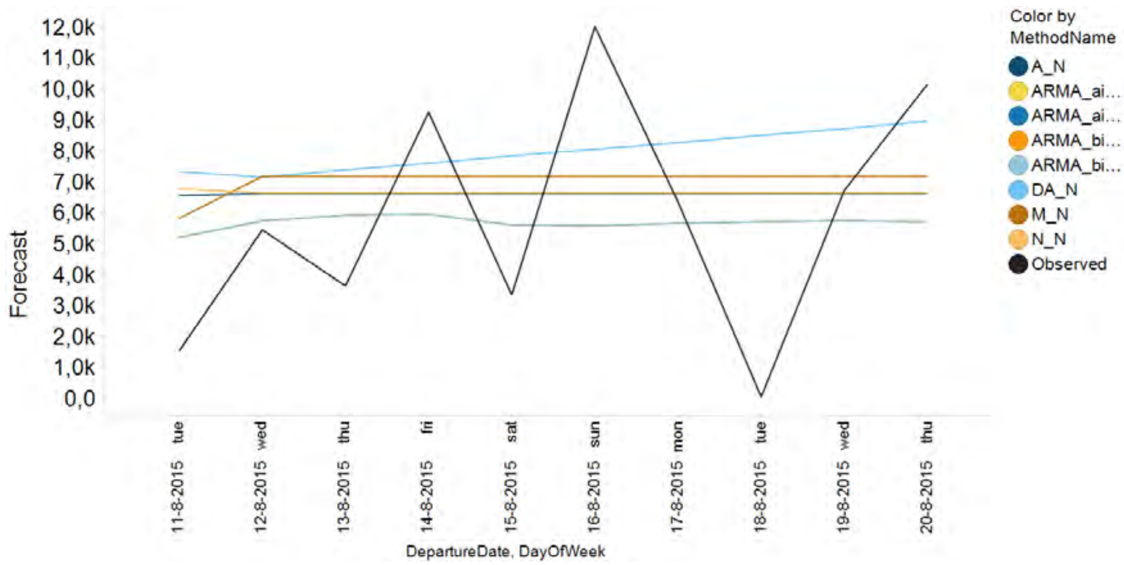

Figure C.2: Raw: Flight DL0230 Boston (United States) - Amsterdam (Netherlands), WAPE of 51.09%

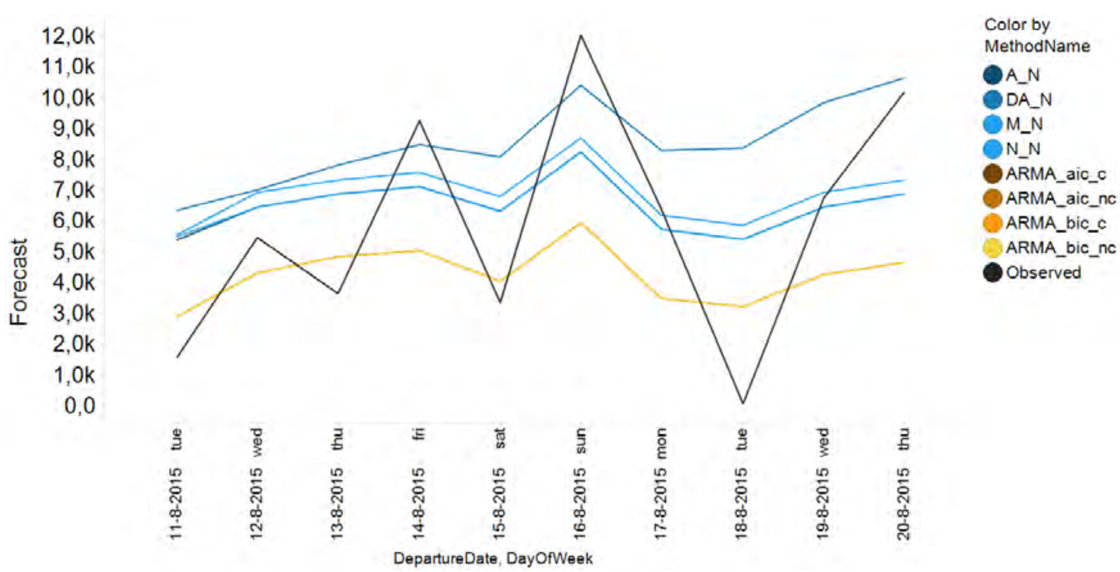

Figure C.3: Decompose: Flight DL0230 Boston (United States) - Amsterdam (Netherlands), WAPE of 45.06%

C.2 Bookings at hand versus materialization

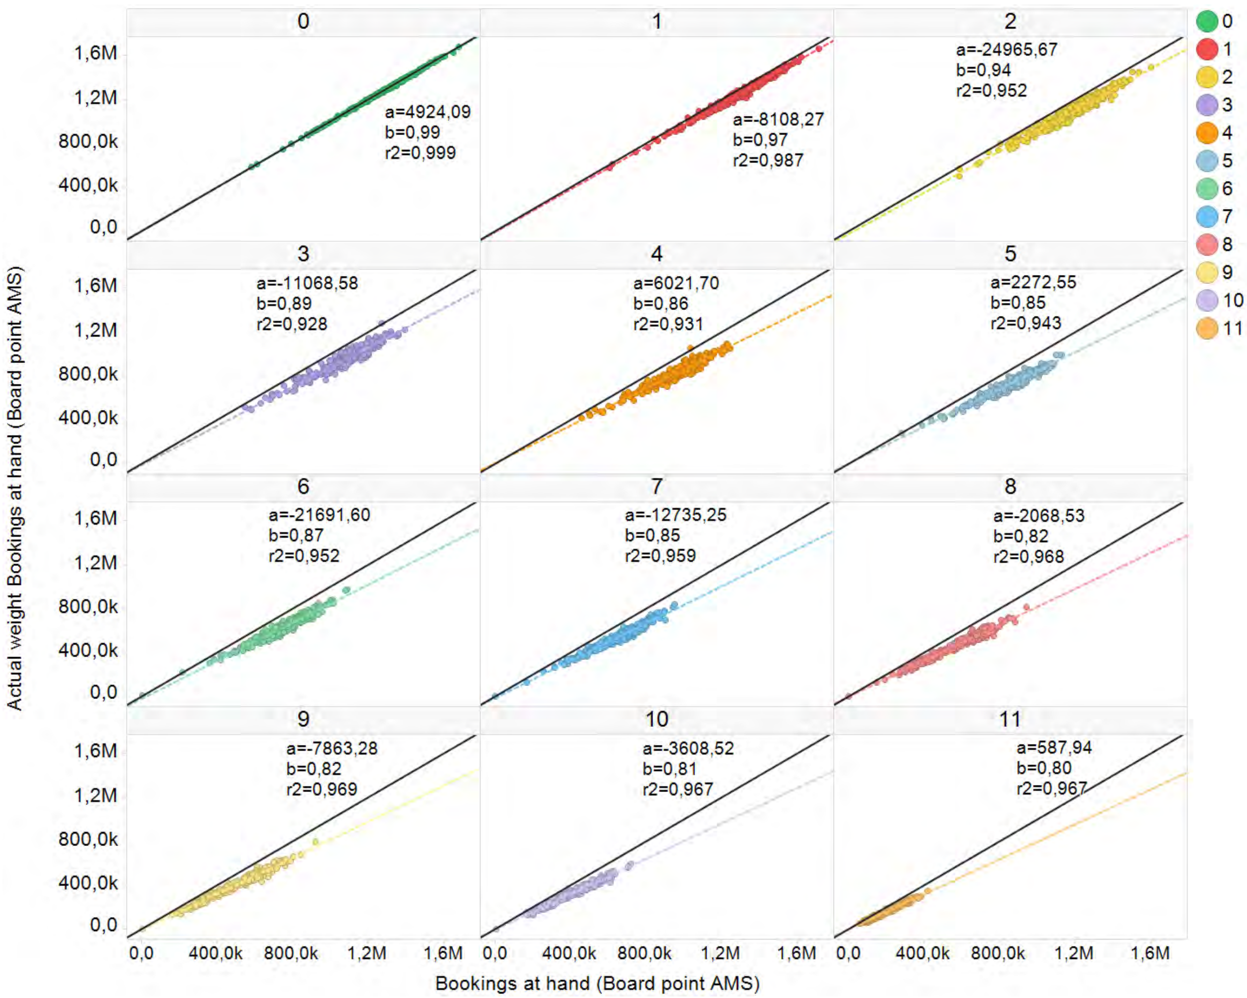

Figure C.4: Bookings at hand (Board point AMS) versus materialization for all DBDs in booking window

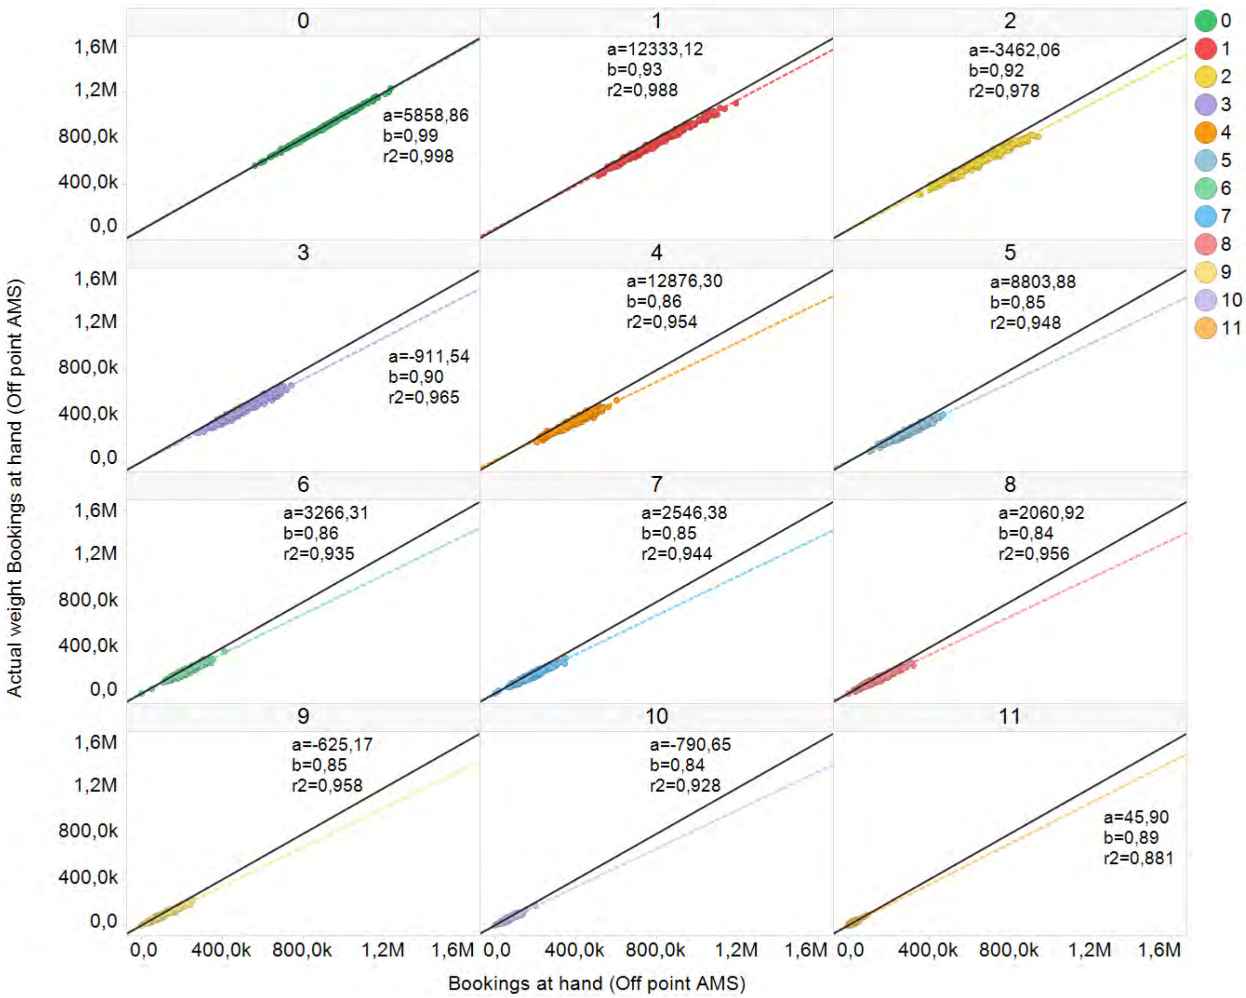

Figure C.5: Bookings at hand (Off point AMS) versus materialization for all DBDs in booking window

C.3 Bookings at hand versus total materialization

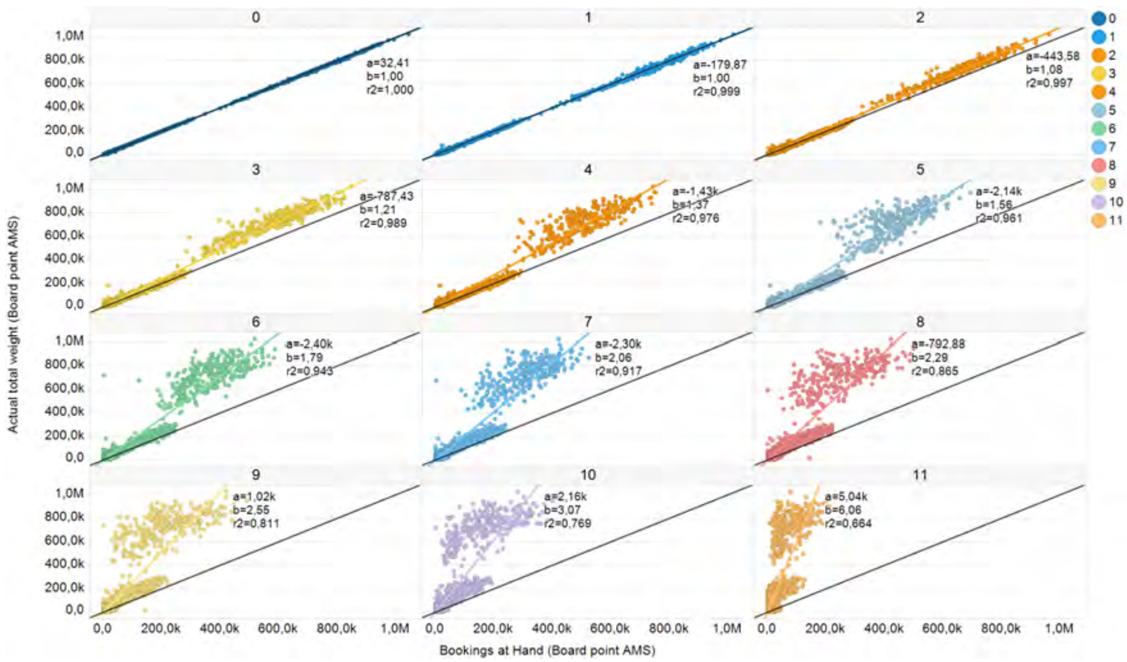

Figure C.6: Bookings at hand (Board point AMS) versus total materialization (bookings at hand plus bookings to come) for all DBDs in booking window

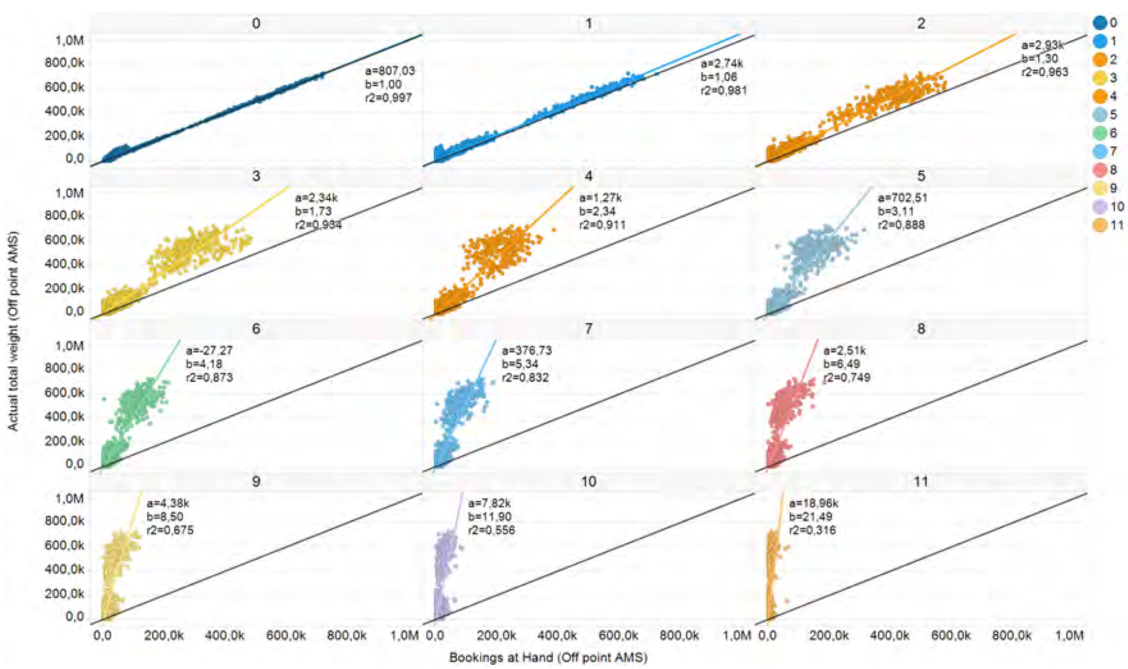

Figure C.7: Bookings at hand (Off point AMS) versus total materialization (bookings at hand plus bookings to come) for all DBDs in booking window

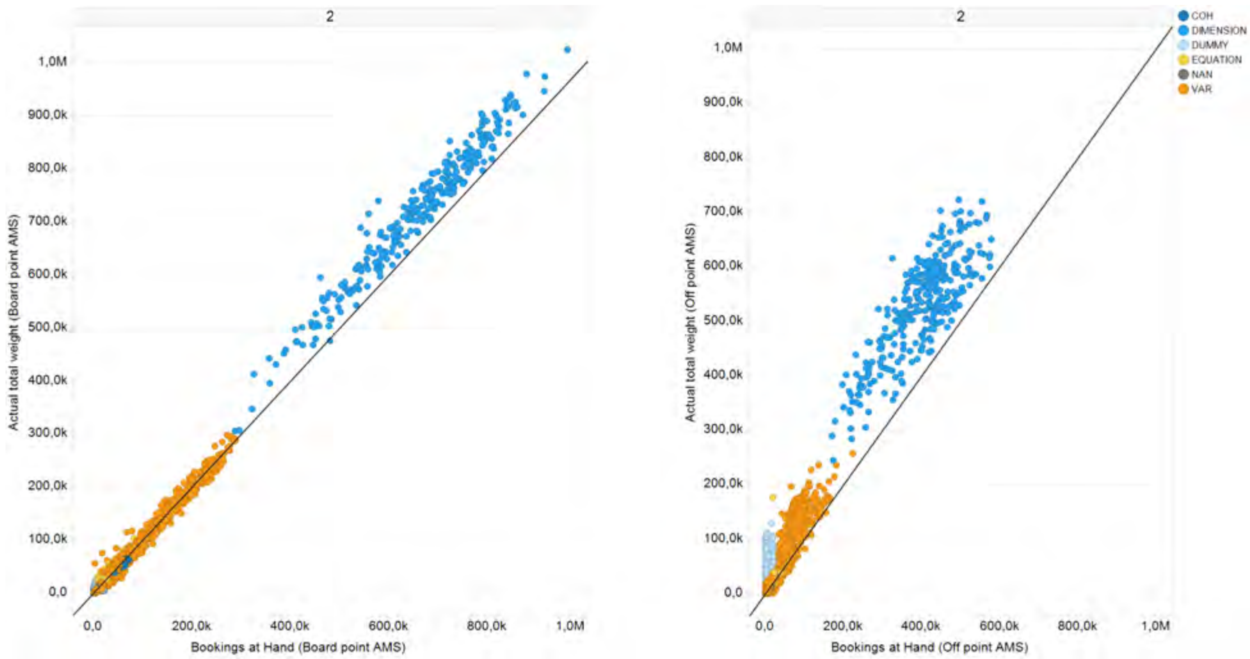

Figure C.8: Example of the product differentiation within bookings at hand versus total materialization (bookings made plus bookings to come) for DBD 2

C.4 Materialization of bookings at hand versus total materialization

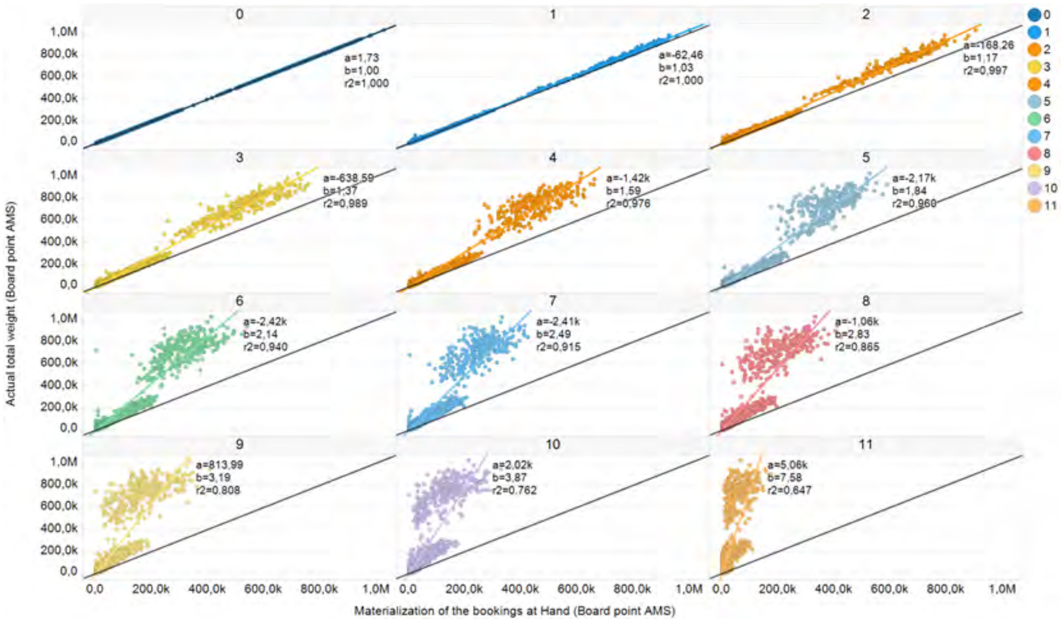

Figure C.9: Materialization of the bookings at hand (Board point AMS) versus total materialization (bookings at hand plus bookings to come) for all DBDs in booking window

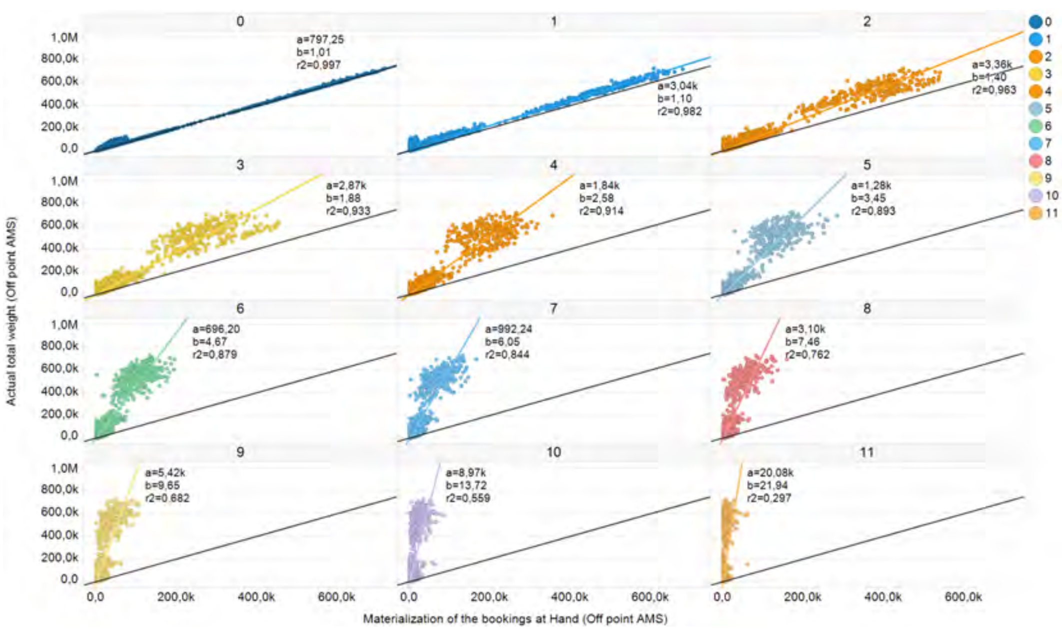

Figure C.10: Materialization of the bookings at hand (Off point AMS) versus total materialization (bookings at hand plus bookings to come) for all DBDs in booking window

Bibliography

AFKL (2015a). About air france klm. http://www.klm.com/corporate/en/about-klm/air-france-klm/index.html. 6

AFKL (2015b). About air france klm cargo. https://www.afklcargo.com/WW/en/local/about_us/about_us.jsp. 6, 7, 10

Benfatma, W. and Garcia, F. (2014). *Opportunity Assessment*. Confidential Functional Design Document (FDD). 15, 18

Boeing (2014). Air cargo market: Rtk over time. http://www.boeing.com/commercial/market/long-term-market/business-and-market-environment/. 3

Bokern, B. and Mulder, J. (2015). Long term demand forecast. Internal KLM presentation. 25

Coghlan, A. (2010). Using r for time series. https://media.readthedocs.org/pdf/a-little-book-of-r-for-time-series/latest/a-little-book-of-r-for-time-series.pdf. 20, 24

de Jong, J. (2014). *Procesmodel PRI*. Confidential Powerpoint. 10

de Jong, J. (2015). *Personal interview with Joke de Jong*. Projectmanager PRI. 8, 13

d'Engelbronner, R. (2012). *Air Cargo Handling Demand Forecasting*. KLM Cargo. 8, 15

Garcia, F. (2014). *Cargo Demand Forecast Short Term*. Confidential Document. 15, 16, 19

Gardner, E. S. (1985). Exponential smoothing: The state of the art. *Journal of forecasting*, 4(1):1–28. 20, 74, 75

Huang, K. and Hsu, W. (2005). Revenue management for air cargo space with supply uncertainty. In *Proceedings of the Eastern Asia Society for Transportation Studies*, volume 5, pages 570–580. 15

Hyndman, R. J. and Athanasopoulos, G. (2014). *Forecasting: principles and practice*. OTexts. 20, 21, 24

IATA, I. A. T. A. (2015). Air cargo market: Ftk and volume over time. http://www.iata.org/publications/economics/Pages/index.aspx. 4

Kalekar, P. S. (2004). Time series forecasting using holt-winters exponential smoothing. *Kanwal Rekhi School of Information Technology*, 4329008:1–13. 20, 21

Maddala, L. (1992). *Introduction to econometrics*, volume 2. Macmillan New York. 24

Meko, D. (2009). Geos 585a, applied time series analysis. *University of Arizona. http://www. ltrr. arizona. edu/˜ dmeko/geos585a. html (accessed 15 February 2012)*. 23

Popescu, A., Keskinocak, P., Johnson, E., LaDue, M., and Kasilingam, R. (2006). Estimating air-cargo overbooking based on a discrete show-up-rate distribution. *Interfaces*, 36(3):248–258. 5

SBRANA, G. (2012). Damped trend exponential smoothing: Prediction and control. *Journal of Quantitative Economics*, 10(2). 21

Spotters, P. (2015). Fleet details and history. `https://www.planespotters.net/airline/`. 7

Stigler, M. (2014). Stationary models. `http://matthieustigler.github.io/Lectures/Lect2ARMA.pdf`. 24

Veldman, A. (2015). *Personal interview with Aalko Veldman*. Resource Planner Freight Building 3. 12

Versnel, L., Bokern, B., Mulder, J., and Fontaine, M. (2015). Long term cargo demand forecast tool. Internal KLM application. 4, 16

Wever, G. (2011). *Improving resource planning at KLM Cargo*. KLM Cargo. 6

Zahedi, H. (2008). *Data mining to examine the treatment of osteomyelitis*. ProQuest. 22