NTNU
Norwegian University of
Science and Technology

Developing a Forecast-Based Optimization Model for Fantasy Premier League

William Eilertsen
Akash Gupta
Bjørn Kåre Kristiansen

Industrial Economics and Technology Management
Submission date: June 2018
Supervisor: Magnus Stålhane, IØT
Co-supervisor: Xin Wang, IØT
Giovanni Pantuso, University of Copenhagen

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

Preface

This Master's Thesis is written as a concluding part for achieving a Master of Science in Industrial Economics and Technology Management at the Norwegian University of Science and Technology. It was produced in the courses TIØ4905 - Managerial Economics and Operations Research, Master's Thesis and TIØ4900 - Financial Engineering, Master's Thesis. All the authors have technical background in Mechanical Engineering. Two of the authors have specialized in Empirical and Quantitative Methods in Finance, while one have specialized in Applied Economics and Optimization. The study has been conducted over the spring of 2018.

The scope of this thesis has been determined by the authors of the thesis along with their supervisors Magnus Stålhane, Xin Wang and Giovanni Pantuso. The objective is to create an optimization model for the Fantasy Sports, Fantasy Premier League, a rather unfamiliar field of research.

June 10, 2018
Trondheim

William Eilertsen Akash Gupta Bjørn Kåre Kristiansen

Summary

Fantasy Premier League is an online game where participants assemble imaginary teams consisting of real-life football players competing in the top tier of English football, the Premier League. Participants gain points week-by-week based on the real-life performances of the players they have selected. In this thesis, a mathematical model describing Fantasy Premier League is developed. Moreover, a forecast-based optimization model is developed to make decisions in Fantasy Premier League.

The mathematical model has been solved by using a rolling horizon heuristic with forecasts of player points as input. Three forecasting methods are developed in order to generate forecasts. The first approach is centered around the most recent average points obtained by each individual player. The second approach is based on regression on multiple explanatory variables. The third and last approach utilizes bookmakers' odds to predict points, combining team odds and individual player odds. In order to obtain data for the last approach, we have cooperated with Sportradar, a Norwegian company providing odds for international bookmakers. The mathematical model developed is also solved with realized points to obtain the optimal solution.

An additional feature of Fantasy Premier League is a concept called gamechips. Gamechips can be used to improve the performances in particular gameweeks. These gamechips are modeled in the mathematical model. Strategies for the implementation of gamechips are developed and their impact is tested. Furthermore, the effect of adding risk handling constraints has been analyzed. It is tested whether there exists a trade-off between risk and reward, similar to that found in portfolio optimization.

The model has been run for the first 35 gameweeks of the 2017/2018 Fantasy Premier League season, and the results are compared to the performance of Fantasy Premier League participants. The last part of the thesis is dedicated to a discussion of further research possibilities.

Sammendrag

Fantasy Premier League er et online spill hvor deltakere samler imaginære lag bestående av ekte spillere i engelsk fotballs toppdivisjon, Premier League. Deltakerne blir hver uke tildelt poeng basert på spillernes prestasjoner i ekte Premier League-kamper. I denne oppgaven utvikles en matematisk modell som beskriver Fantasy Premier League. Videre er en prognosebasert optimeringsmodell utviklet for å gjøre beslutninger i Fantasy Premier League.

Den matematiske modellen har blitt løst ved hjelp av en rullende horisont heuristikk med prognoser for spillerpoeng som input. Tre metoder er utviklet for å generere prognoser. Den første tilnærmingen er sentrert rundt gjennomsnittet av poeng oppnådd tidligere for hver enkelt spiller. Den andre tilnærmingen er basert på regresjon på flere forklarende variabler. Den tredje og siste tilnærmingen utnytter odds satt av spillselskaper for å forutsi poeng. Både odds relatert til lag og individuelle spillere er benyttet. For å få tilgang til data for den siste metoden har vi samarbeidet med Sportradar, et norsk selskap som tilbyr odds til internasjonale bookmakere. Den matematiske modellen som er utviklet løses også med realiserte poeng for å finne den optimale løsningen.

Fantasy Premier League inneholder et konsept kalt gamechips. Gamechips kan brukes for å forbedre prestasjonen i spesielle uker. Disse gamechipsene er modellert i den matematiske modellen. Strategier for implementering av gamechiper er utviklet og gamechipenes innflytelse har blitt testet. Videre er effekten av å legge til begrensninger relatert til risikohåndtering analysert. Det er testet om det eksisterer en avveining mellom risiko og belønning tilsvarende det som finnes i porteføljeoptimering.

Modellen har blitt kjørt for de første 35 ukene av 2017/2018 Fantasy Premier League sesongen, og resultatene har blitt sammenlignet med prestasjonene til Fantasy Premier League-deltakere. Den siste delen av oppgaven er dedikert til en diskusjon av fremtidige forskningsmuligheter.

Acknowledgements

First, we want to thank our supervisors, Magnus Stålhane, Xin Wang and Giovanni Pantuso. Their commitment and interest have been essential for the progress of this thesis and an important source of motivation. Their guidance and quick and thorough feedback have been a crucial contribution in making this thesis a reality.

Additionally, we want to thank Sportradar, and in particular Espen Rødsand, for providing us with sophisticated data used in this thesis and related technical assistance. With their help, we have been able to apply our knowledge to an area we share a combined passion for, namely football. This has been an immense motivational factor throughout the process of writing this thesis.

Table of Contents

List of Tables

List of Figures

Chapter 1

Introduction

For the past decade, a phenomenon called Fantasy Sports has experienced enormous growth, both in terms of active participants and media attention. According to Fantasy Sports Trade Association (2017), it is estimated that 59.3 million people in the U.S. and Canada are competing in some kind of Fantasy Sports. Furthermore, the Fantasy Sports industry is expected to grow annually by 41% and generate $14.4 billion by the year 2020 (Heitner, 2017). Fantasy Sports is an online game where participants assemble imaginary teams consisting of real players within a professional sport. The teams compete in both private- and public leagues, where they gain points based on players' week-by-week performances in real-life matches. Fantasy Premier League (FPL) is the Fantasy Sports of England's top division football league, Premier League. It is the largest Fantasy Sports in Europe, with more than 5.9 million active participants.[1]

Typically, Fantasy Sports participants select a squad consisting of a given amount of players in each position on the field. Each real-life player is given a buying price based on their skill level and ability to perform well in the Fantasy Sports. The performance of each player in the actual games is measured in terms of points, where an action like scoring a goal will provide the player with a specified amount of points. To keep participants from only selecting top-rated real-life players, a restricted budget is usually given.

Fantasy Sports has its origin from the 1950s when a limited partner of the Oakland Raiders, Wilfred Winkenbach, invented a Fantasy Sports for golf. Each participant selected a team of professional golfers, and the one who suggested the team with the best score would win the tournament. In 1962, Winkenbach suggested the first fantasy league for American Football, initiating the birth of what has become a billion-dollar business (McCracken, 2012). Nowadays, Fantasy Sports are available for several sports, including American football, basketball, ice hockey, baseball, football and cricket.

[1] The figure is taken from FPL's homepage: https://fantasy.premierleague.com/

Over the last years, the interest for Fantasy Sports has proliferated among commercial actors. For instance, a phenomenon called Daily Fantasy Sports (DFS) has arisen. It is a version of Fantasy Sports, allowing participants to compete for money over short time-periods, typically for one day or a weekend. In the U.S., FanDuel and DraftKings are the leading providers of DFS, with FanDuel being valued to more than $1 billion (Bertoni, 2015). Both FanDuel and DraftKings were granted licenses from the United Kingdom Gambling Commission in 2015 (Purdum, 2017). Furthermore, in 2017, the Scandinavian bookmaker Unibet launched Fantasy Sports Beta. Fantasy Sports Beta allows customers to compete for money within Fantasy Premier League. In addition, the Norwegian bookmaker GUTS included bets on Fantasy Premier League in 2016. They offer bets on whether one player will perform better than another for the upcoming week.

In the U.S., an estimated 30% of Fantasy Sports participants use additional websites to research athletes and other factors when they construct their fantasy teams. In total, these participants spend $656 million annually to purchase additional information and decision-making tools (Burke et al., 2016). There exist several optimization models for Fantasy Sports, for instance for ice hockey (Summers et al., 2005) and American football (Fry et al., 2007). However, compared to the models suggested for the mentioned Fantasy Sports, there is a difference when modeling Fantasy Premier League. For instance, classical American fantasy leagues consist of selecting a team for an entire season and do not allow for transfers. On the contrary, Fantasy Premier League allows for making player transfers during the season. Thus, FPL participants have to consider making new decisions each gameweek in a season. To the author's knowledge, an optimization model making optimal decisions in Fantasy Premier League has not been developed before. Given the number of participants and commercial actors, and their willingness to invest both money and time in Fantasy Sports, a potential for developing a model optimizing decisions in the Fantasy Premier League is seen. Sportradar, a Norwegian company providing data for international bookmakers, has assisted in writing this thesis by providing relevant data.

An optimization model for the FPL could serve two purposes. First, many actors may be interested in detailed analysis of the FPL run ex-post, i.e., after points are realized. Such a model could provide answers to questions like, "What is the optimal team so far?" and, "What are the optimal transfers for each particular gameweek?". Furthermore, the model can be used to determine what would have been the optimal transfers in a particular gameweek for a given participant's team. Moreover, the optimal solution can be compared to the score of the participant who finished on top in the overall rankings, and the gap can be assessed. Secondly, the model can be run ex-ante, i.e., before points are realized. Thereby, it can be used to determine how well the model performs when competing by the same rules as FPL participants and subsequently for decision support. However, in this case, forecasts of points are required. Therefore, in this thesis, three methods of generating such forecasts are suggested. Thus, the scope of this thesis is two-folded. First, the aim is to formulate a mathematical model for the Fantasy Premier League. Secondly, the aim is to develop a model that can compete with human participants in terms of overall ranking.

The thesis is structured as follows. Chapter 2 provides the problem description. Here, the rules and point system of Fantasy Premier League are presented. Chapter 3 provides a literature review. First, studies of real-life football that can be related to Fantasy Sports are presented. Then, a presentation of Fantasy Sports optimization models, player performance forecast and a comparison to other well-known problems are given. Chapter 4 describes the mathematical formulation of the problem. Chapter 5 presents a solution framework, including a description of the solution method used for solving the mathematical model. In addition, three methods for generating forecasts are described. By combining all elements of the solution framework, a forecast-based optimization model is developed. Chapter 6 presents a computational study undertaken to set up the model presented in Chapter 5 before it is run on the 2017/2018 season of the FPL. In Chapter 7, the performance of the model in 2017/2018 season of the FPL are presented. In Chapter 8, the concluding remarks highlight the most important results and implications of the research. Finally, Chapter 9 addresses the potential for further research.

Chapter 2

Problem Description

This chapter presents the problem description of the thesis. First, the rules and point system of Fantasy Premier League are described in Section 2.1 and Section 2.2, respectively. The rules and point system are given in correspondence with the information found on FPL's homepage, www.fantasy.premierleague.com. Finally, in Section 2.3, the *Fantasy Premier League Decision Problem* is defined.

Before the rules of the FPL and the point system are presented, some terms are defined. First, it is important to make the distinction between a *manager* and a *player* clear. Henceforth, the term manager refers to a participant in the FPL, while player refers to a real-life football player competing in the English Premier League. Furthermore, by football, it is always referred to European Football, also known as soccer.

2.1 Rules of Fantasy Premier League

In order to understand the dynamics of Fantasy Premier League, a review of the game rules is necessary. The Premier League consists of 20 teams, and every team faces each other two times during a season. Thus, each season consists of 380 fixtures split into a set of 38 chronological *gameweeks*. In general, every gameweek consists of 10 fixtures featuring each team once. All the matches within a gameweek are usually played over a period of 3-4 days. However, due to matches in national and international cups, some of the Premier League matches are postponed or played ahead of its original dates. This implies that some of the gameweeks do not consist of exactly 10 fixtures. Gameweeks consisting of more than 10 fixtures are hereby referred to as *double* gameweeks, while gameweeks consisting of less than 10 fixtures are referred to as *blank* gameweeks.

Initially, a manager must select a team of 15 players consisting of exactly 2 goalkeepers, 5 defenders, 5 midfielders and 3 forwards. This is regarded as the *selected squad*. As each team has a registered squad of 25 players, the whole FPL database consists of approximately 500 players. This creates numerous ways of selecting the team. The total purchase price of the team may not exceed the initial budget limit of £100m. Notice that the budget limit may vary over time, as player prices increase and decrease during the season. Further, more than 3 players cannot be selected from the same club. After the FPL manager has picked the selected squad, a *starting line-up*, consisting of 11 players from the selected squad has to be chosen. A manager can compose the starting line-up in any formation as long as there is exactly 1 goalkeeper, at least 3 defenders, at least 3 midfielders and at least 1 forward. This is regarded as the *team formation criterion*. The players in the starting line-up are the only players that are awarded points. However, if any player from the starting line-up is not playing, he will be replaced by one of the four substitutes on the bench. As only 1 goalkeeper can feature in the starting line-up, the goalkeeper will only be substituted in if the goalkeeper in the starting line-up is not playing. For the remaining 3 substitutes, each manager must set a *substitution priority*, deciding the order of substitution. It is important to note that the formation criterion overrules the substitution priority. Consider a case with 3 defenders in the starting line-up and a midfielder with substitution priority 1. If a defender in the starting line-up does not play, the formation criterion ensures that the midfielder cannot replace the defender.

Each player in the FPL database is assigned with a price, depending on how well the player is expected to perform. Player prices change during the season depending on the popularity of the player. A player's price will increase when more managers select him and decrease when fewer managers select him. If the price of a player increases while a manager has the player in the selected squad, the manager can benefit from the price increase. However, if the player is sold, a *sell-on fee* of 50% is inferred. Hence, the manager only receives 50% of the profit (rounded down to the nearest £0.1m). For example, if a manager buys a player for £9.5m and the player's value increases to £10.0m in a later gameweek, the manager receives £9.7m if the player is sold, earning a profit of £0.2m. Thus, there is a distinction between a player's *value* and *sell price*. A player's value is the price a player is listed with in the FPL database. If a player is *not* in a manager's selected squad, this is the price the manager must pay for the player. A player's sell price is the price a manager can sell a player in the selected squad for. If the price of a player decreases, his sell-price simply equals his value.

For each gameweek, managers may use one *free transfer* if they want to replace one of their players in the selected squad with another player. It is possible to make more than one transfer for a gameweek, but each additional transfer will deduct 4 points from the total score. Hereby, transfers deducting points are referred to as *penalized transfers*. If a manager chooses not to make any transfers for a gameweek, the free transfer is saved for the next gameweek, and the manager can make two free transfers in the upcoming gameweek. However, a manager cannot save up more than two free transfers. That is, not making any transfers for two consecutive gameweeks does not give the opportunity to make three free transfers in the upcoming gameweek.

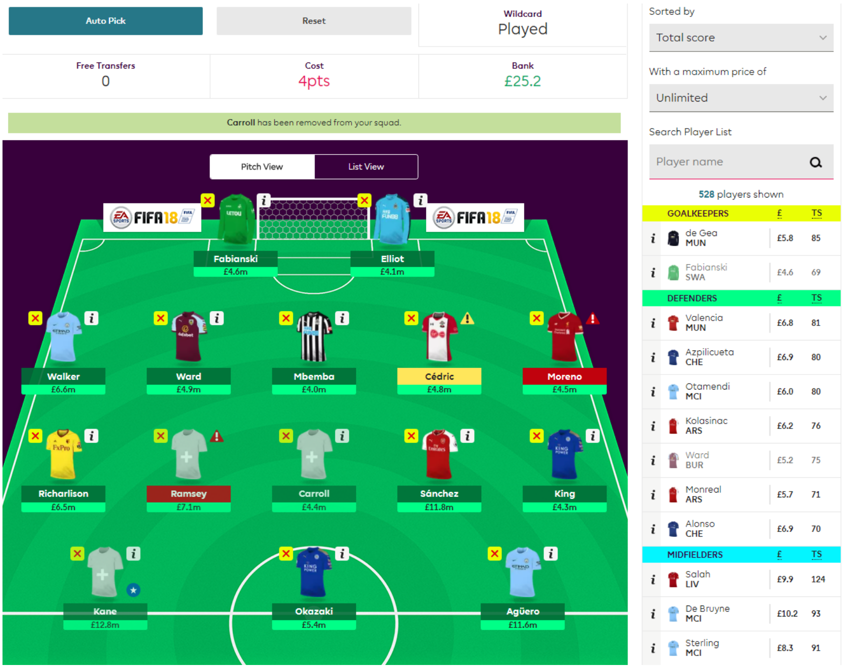

Figure 2.1: Illustration of the team selection page of Fantasy Premier League.
Source: https://fantasy.premierleague.com/a/squad/transfers (10.01.2018)

Figure 2.1 illustrates how a selected squad is picked in FPL. The grey shirts indicate that the manager is considering to remove Kane, Ramsey and Carroll from the selected squad. As one can see, this deducts 4 points. Hence, the manager had two free transfers for this gameweek. Further, the manager has £25.2m in the bank for new players if these transfers are carried out. Also, it can be observed that Ramsey and Moreno have a red triangle attached to them. This implies that they are either injured or suspended for the upcoming gameweek. Notice how the selected squad consists of exactly 2 goalkeepers, 5 defenders, 5 midfielders and 3 forwards as stated in the rules.

2.2 The Point System

The Fantasy Premier League managers are rewarded with points depending on their performance during a gameweek. The main point factors are related to goals scored, assists given and clean sheets. In the following, the point system is presented.

	Goalkeepers	Defenders	Midfielders	Forwards
For playing up to 60 minutes	1	1	1	1
For playing 60 minutes or more	2	2	2	2
For each goal scored	6	6	5	4
For each assist	3	3	3	3
For keeping a clean sheet	4	4	1	-
For every 3 saves made	1	-	-	-
For each penalty saved	5	-	-	-
For each penalty missed	-2	-2	-2	-2
For every two goals conceded	-1	-1	-	-
For each yellow card	-1	-1	-1	-1
For each red card	-3	-3	-3	-3
For each own goal	-2	-2	-2	-2

Table 2.1: Point system of FPL.

- In a match, the three best players are decided according to the FPL Bonus Points System and awarded a bonus of 1, 2 and 3 points. Bonus Points are calculated according to 32 match statistics, where goals scored, assists and clean sheets are the factors that are heaviest weighted. A complete overview of the Bonus Points System is given in Appendix A.
- For a goalkeeper or defender to receive points for a clean sheet, he has to play at least 60 minutes, excluding stoppage time.
- If a goal is scored on a direct free kick or a penalty, the player who got the free kick/penalty is awarded an assist.
- For each round, a manager chooses a captain and a vice-captain. If the captain is playing, he will be awarded double points for the entire round. However, if the captain is not playing, the vice-captain will be awarded double points. If neither the captain nor the vise-captain played, no players are given double points.

Notice that the point system does *not* consider the outcome of a match. Hence, players are not rewarded with points for playing on a team that wins or loses. In addition to the regular scoring system, each manager is awarded four different gamechips. Only one gamechip can be used in a gameweek. The gamechips are:

(i) **Wildcard**. The Wildcard allows the manager to replace his entire selected squad for free. As for the Wildcard squad selection, the same rules apply as in the regular fantasy team composition. Hence, one can only select a maximum of three players from each team, and the formation criterion must be upheld. When playing a Wildcard, a manager's budget is set to the sell price of his selected squad in that particular gameweek. Further, when playing a Wildcard, any saved transfers will be lost. The gamechip can be used twice a season, once in the first and once in the second half of the season.

(ii) **Bench Boost**. The Bench Boost allows a manager to receive points for all the 15 players in the selected squad. The gamechip can only be used once a season.

(iii) **Free Hit**. The The Free Hit allows a manager to replace the entire selected squad for one gameweek. However, for the next gameweek, the selected squad is reversed back to the squad from the previous gameweek. The Free Hit can only be used once a season. As for the Wildcard, the same transfer- and budget rules apply for the Free Hit. The gamechip can only be used once a season.

(iv) **Triple Captain**. The Triple Captain triples the points of the captain for a gameweek. If the captain does not play, the points of the vice-captain are tripled. If neither the captain nor the vise-captain play, no players are awarded triple points. The gamechip can only be used once a season.

2.3 The Fantasy Premier League Decision Problem

In Fantasy Premier League each manager attempts to maximize his or her total number of points over an entire season. As seen in this chapter, a number of decisions must be made each gameweek. The decisions include which players to add to the selected squad and which players to pick for the starting line-up. In addition, a captain- and vice-captain must be chosen. Further, one has to set a substitution priority for the players that are not selected for the starting line-up. As pointed out, a manager is allowed to perform transfers during the season. Therefore, a manager must also decide whether to make a transfer or not, and consequently which players to be transferred in and out for each gameweek. Hence, decisions are made in a multi-period manner. Hereby, the aforementioned decisions are referred to as the *Fantasy Premier League Decision Problem* (FPLDP).

Chapter 3

Literature Review

The literature review focuses on academic work related to the FPLDP. First, Section 3.1 compares Fantasy Football to real-life football. By Fantasy Football, all Fantasy Sports with basis in football are meant. Further, Section 3.2 focuses on the use of operations research in sports, Fantasy Sports and Fantasy Football. Next, Section 3.3 focuses on forecasting of player performance. This includes forecasting of individual player performance, rating the quality of football teams and works related to the existence of a home-field advantage. Section 3.4 draws parallels to other well-studied problems. Finally, the chapter is completed by emphasizing our contribution.

3.1 Football Team Composition

Multiple studies are devoted to the topic of team selection in football. For instance, Boon and Sierksma (2003) provided an optimization model for selecting the line-up of sports teams, accounting for questions like, "How to select an optimal line-up from the set of all candidates in different positions?". Further, Trninić et al. (2008) discussed the importance of individual player's roles and tasks when selecting the starting line-up, as they focused on finding the position and role most suitable for each player. They argued that consistency is an important consideration when selecting a line-up. Moreover, Ozceylan (2016) identified key performance criteria for each position and suggested an approach for selecting the players with the greatest contribution to the team. Similarly, Tavana et al. (2013) proposed a framework for player selection and team formation with regards to determining the collection of individual players that form an effective team. Pantuso (2017) proposed a stochastic programming approach for the composition of football teams. The approach takes the requirement of a mix of skills, competition regulations and budget limits into consideration.

When selecting a real-life football line-up, other considerations are faced than when selecting a Fantasy Football line-up. For instance, in real-life football it is important that the players hold a necessary set of skills that are compatible to the team's organization of play (Pantuso, 2017). Furthermore, a quality like leadership ability is considered essential when selecting the captain. In addition, different skills are required for different positions on the field (Boon and Sierksma, 2003). Moreover, it is vital that the selected players have a good chemistry and play each other better. Hence, in real-life football several factors are considered when selecting the team composition that are not important for the FPL. In FPL, it is only relevant to select players that perform well in terms of Fantasy Premier League points. Thus, the chemistry between players has limited impact, as the team mainly consists of players who do not play in the same club. In addition, while leadership ability is essential when selecting the captain of a real-life football line-up, it is an irrelevant consideration with regards to FPL line-ups. Moreover, player skills like stamina, interceptions and passing have a low impact for the points obtained by a FPL player. For instance, N'golo Kanté, the winner of the Professional Footballers' Association (PFA) Players' Player of the year award in 2017 (Skysports, 2017), is considered a hard-working midfield anchor. During the 2016-2017 season, Kanté earned a total of 83 points in FPL, scoring 1 goal, having 1 assist and receiving 9 yellow cards. By comparison, the best FPL midfielder in terms of FPL points, Alexis Sánchez, earned 264 points. However, Kanté's contribution was crucial for Chelsea who won the Premier League that season. Thus, although Fantasy Premier League is based upon the performances of real-life football, there are significant differences in terms of selecting the team. While FPL managers should mainly focus on players that score goals, have assists and keep clean sheets, actual football managers have to ensure that the team is composed with the mix of required and compatible skills (Pantuso, 2017). Furthermore, other distinctions such as transfer rules and the gamechips make the problem of picking a selected squad in FPL and a real-life football team different. Hence, uncritically applying models from real-life football to the FPL can be inappropriate.

3.2 Operations Research in Sports

In recent years, the development of literature in sport analytics related to decision making has been rapid. Well known journals in statistics, applied mathematics, operations research and economics have published a great number of articles on this field (Coleman, 2012). An area with great progress within operations research is called "Sports Scheduling". This field of research contains problems which involve deciding optimal league seasons' schedules for different sports. The arena of sports scheduling has existed for over 40 years, but only until recently the number of papers has increased significantly (Kendall et al., 2010). The use of operations research can be explained by the growing complexity of playoff structures, league divisions and opposing demands from constituent teams and other stakeholders. Its use has been investigated in sports including volleyball (Bonomo et al. (2012)), table-tennis (Schönberger et al. (2004); Knust (2010)), cricket (Wright (2005)), basketball (Wright (2006); van Voorhis (2002); Henz (2001)), baseball (Trick et al. (2012)), Canadian football (Kostuk and Willoughby (2012)) and softball (Saur et al. (2012)). In football, sports scheduling is also prominent. Bartsch et al. (2006) scheduled the professional soccer leagues of Austria and Germany. Besides constraints for inner-league requirements and preferences, the European top soccer leagues also have to take into account constraints related to European cup matches (Champions League, UEFA Cup). They developed models and applied greedy algorithms and a branch and bound method which yielded reasonable schedules quickly. The schedules were applied by the professional soccer league in both Austria and Germany. Further, Della and Oliveri (2006) scheduled the Italian football league where the schedules also were balanced with respect to additional cable television requirements. Further, Chile is perhaps one of the countries where sports scheduling in football has been most present. Since 2005, the Chilean professional soccer has adopted a scheduling system that is based on an integer linear programming model. Durán et al. (2007) generated schedules that gave teams benefits such as higher incomes, lower cost, higher fan attendance and fairer seasons. Their work was further expanded (Durán et al., 2012), when they considered challenges for Chile's Second Division soccer league. Similar scheduling problems are also studied in Belgium by Goossens and Spieksma (2009) and in Denmark by Rasmussen (2008). Finally, Kendall (2008) addressed the special case of English football fixtures over holiday periods. In England, during the Christmas and New Years period teams can play up to four matches in less than ten days. Kendall et al. (2010) considered the minimization of the travel distances by English football clubs during this period, and generated schedules which had 25% lower total travel distances than the fixtures actually used.

In the world of Fantasy Sports, the use of operations research is rather limited. Most of the research focuses on American football and the National Football League (NFL) Fantasy Draft. This could partly be explained by more attractive monetary prizes compared to for instance Fantasy Football leagues. For the NFL Fantasy Draft, optimization models and solution methods have been proposed by Fry et al. (2007), Gibson et al. (2007) and Becker and Sun (2016). Though Fantasy Premier League and NFL Fantasy Draft are both Fantasy Sports, they are essentially different problems. For instance, in NFL Fantasy Draft leagues only one manager within a league can own a given player - while in FPL this is not a restriction. Furthermore, for every gameweek, pairs of managers are matched head to head where the one with highest points wins - a feature that is not present in FPL. The most profound similarity that exists is that in each gameweek a fixed number of players in different roles have to be picked. In addition, decisions on starting line-ups are made.

Goossens et al. (2017) discussed the common characteristics in Fantasy Sports games, where game-rule characteristics are most relevant. They presented a mixed integer programming model for finding an optimal set of ex-post decisions in Fantasy Sports games under the assumption that all data are known in advance. The model was tested on a fantasy cycling game with all data known beforehand. Disregarding gamechips, captain and vice-captain, part of the model can with modification be used for the FPLDP. However, including the gamechips increases the complexity and calls for a more sophisticated mathematical model.

To the author's knowledge, there is no operations research literature which explicitly discusses the FPLDP. However, articles related to Fantasy Football exist. Matthews et al. (2012) were the first to develop a sequential team formation algorithm on Fantasy Premier League. They modeled decisions in FPL as a belief-state Markov Decision Problem and solved it using Bayesian Q-learning. Their model produced promising results as it outperformed most FPL managers. They only modeled the Wildcard gamechip, with the qualitative assessment of the opportunity to play it in gameweek 8 and 23. Their results motivate the development of analytic tools for the FPLDP.

Bonomo et al. (2014) presented two optimization models for the Argentinian Fantasy Football. The first model, called *a priori*, determines line-ups and transfers based on predictions of player points, while the second model, called *a posteriori*, determines the optimal line-ups with data known beforehand. The results obtained by the *a priori* model positioned within the top 0.2 % of all managers. They are the first to develop and solve an optimization model for Fantasy Football, and their results encourage the use of operations research in this field. The points in the Argentinian Fantasy Football are to a high degree allocated in the same manner as in the English Fantasy Premier League. They are set by objective statistics (goals, assists, yellow cards etc.) and subjective statistics determined by a newspaper that are somewhat analogous to the Bonus Points System in FPL (see Section 2.2 and Appendix A). Comparison of the optimization model proposed by Bonomo et al. (2014) with the one developed by Goossens et al. (2017), shows that many of the constraints can be related. Further, as with the points allocation, the rules in Argentinian Fantasy Football resemble the rules of FPL. Nevertheless, there are some significant exceptions:

- Gamechips are a feature which is only present in the FPL.
- Only in FPL, player values can increase or decrease based on how many managers select them.
- The Argentinian Fantasy League allows 4 free transfers in each gameweek, while FPL allocates 1 each gameweek. Further, penalized transfers are only featured in FPL.
- The decisions regarding captain and vice-captain choices are only featured in FPL.

As seen by the differences, a manager in FPL faces a more complex set of decisions than a manager in the Argentinian Fantasy League and thus an increased level of uncertainty.

3.3 Forecasting of Player Performance

The ability to predict future player performances is perhaps the most important skill required in Fantasy Premier League. As stated by Smith et al. (2006), fantasy managers tend to ignore historical statistics when selecting their line-up. Since one aim of this thesis is to compete against managers in FPL, forecasting of player performance is essential.

3.3.1 Individual Player Performance

Yang (2015) predicted the results of the 2016 NBA season by running a least-square linear regression model on player's individual statistics and the win-ratio of their team, thus predicting the results of the NBA matches based on player performances. The explanatory variables primarily used are well-known statistics such as points scored, assists, rebounds, steals and blocks. These variables were analyzed for the past 20 seasons in order to predict the outcome of the upcoming season. To evaluate the impact of the players, he assigned each player with a player efficiency rating, a number representing a player's effectiveness by a single number. A limitation of Yang's approach is that the data used mainly focused on offensive contributions. Hence, good defending players are not appropriately valued. Although this paper focuses on predicting a team's performance and not solely player performance, a similar approach may be used to predict the performance of individual Premier League players.

Sæbø and Hvattum (2015) proposed a top-down rating model for football players, using a regression model to capture how players perform relative to their teammates and the opposition. They provided a regression-based player rating using a plus-minus player statistics, measuring the number of goals scored minus the number of goals conceded when a player is used by the team. Further, they discounted earlier observations and placed greater emphasis on recent performances. As their model did not sufficiently differentiate between players from different divisions or different league systems, their model was further improved (Sæbø and Hvattum, 2017). The model was extended by adding a factor depending on a player's current league and division.

An approach used for rating the players in the Norwegian Tippeligaen, the top tier of Norwegian football, was suggested by Bjertnes et al. (2016). They applied a binary logistic regression in order to assign values to shots attempted in football, hence calculating the probability of a shot resulting in a goal scored. In addition, they developed two Markov game models in order to evaluate all player's actions, not only the shots. Their analyzes are made post-game. Thus, there is no forecasting of player performance. However, the regression variables used and the idea of valuing players using regression are of great interest.

As mentioned in Section 3.2, Bonomo et al. (2014) proposed a mathematical model for an Argentinian Fantasy League. In addition to the optimization model, they also forecast future player points for the individual players. The forecasts were calculated by averaging the points obtained in the three last gameweeks for each particular player. Further, a player's predicted points were multiplied by four deciding factors related to the upcoming match:

- 1.05 if the team was playing at home and 0.95 was the team is playing away.
- A value between 0.95 and 1.05, linearly dependent on the position in the league table. The players on the team leading the league were weighted by 1.05, while the players on the team in the bottom of the league were weighted by 0.95.
- 0.95 to 1.05, depending on a point streak factor.

- A starting line-up factor. This factor was set to 1 for players assumed to be starting in the next gameweek. The factor was based on coaches' announcements, press reports, or information posted on the game website.

As the problem of modeling the Argentinian Fantasy League resembles the FPLDP, a similar approach can be used for the English Fantasy Premier League. Therefore, in Section 5.2, the approach suggested by Bonomo et al. (2014) is modified for the FPL.

3.3.2 Rating Football Teams

The individual performance of a player can be assumed to be impacted by the performance of the player's team. In terms of FPL points, it is reasonable to assume that a player will perform better if he plays for a top-rated team than for a weaker team, as these teams score more goals and keep more clean sheets. Team strengths have been a hot topic in scoreline prediction, i.e., the prediction of final scorelines, for several years. Gamblers try to develop scoreline predictions in order to create realistic odds for upcoming matches, hence trying to beat the bookmakers. In addition to the development of profitable gambling, an appropriate scoreline prediction could be of great use when predicting the future performance of Premier League players.

Maher (1982) introduced attacking- and defensive strength parameters for each team in order to develop a model calculating the expected goals scored by both the home- and away team, where the number of goals was modeled as a Poisson distribution. This was further extended by Dixon and Coles (1997) who were able to calculate probabilities of different scorelines in a given match, by which they created a strategy for profitable gambling. Rue and Salvesen (2000) suggested that the attacking- and defensive strengths of a team were time-dependent, and updated the strength estimates using Bayesian methods. In order to do so, they used Markov Chain Monte Carlo iterative simulation techniques. Further, Crowder et al. (2002) developed a procedure for updating parameter estimates which was computationally less demanding.

Although the approach suggested by Maher (1982) has been improved over the years, there are other ways of determining the strength of a team. In 1978, the Elo rating was introduced by Elo (1978). The rating was initially developed for rating the strength of chess players but has been widely adopted to different sports, including golf, tennis and football. Hvattum and Arntzen (2009) used the Elo rating in order to rank teams when predicting match results in football. The ratings appeared to be useful in encoding the information of past results for measuring the strength of a team. However, when used in terms of match predictions, the predictions appeared to be considerably less accurate compared to market odds. Bjertnes et al. (2016) adopted the approach suggested by Hvattum and Arntzen (2009) for consideration of team strengths as they valued the individual player performance in Norwegian football. In addition, Leitner et al. (2010) assessed the Elo rating alongside with FIFA/Coca-Cola World Rankings (FIFA, 2012) in order to predict the outcome of tournament winners. However, both rating systems were found to be inferior to bookmakers' odds. Hence, it confirms the findings of Hvattum and Arntzen (2009).

Other rating systems investigated in the forecasting literature include team strengths estimated by adopting an ordered probit model with one parameter for each team (Graham and Stott, 2008), predictions based on *power scores* (ratings inspired by the Elo system) and an associated ranking for matches in the National Football League ((Boulier and Stekler, 2003); (Boulier et al., 2006)). In addition, Clarke and Dyte (2000) fitted differences in Association of Tennis Professionals (ATP) rating points to a logistic model and used resultant probabilities in a simulation to estimate each player's chance of victory.

3.3.3 Home-Field Advantage in Football

It is often presumed that football teams perform better when playing at home than when playing away. If this is the case, home-field advantage may play a role when forecasting the performance of a player for an upcoming match. In the following, academic works related to home-field advantage in football is presented.

Nevill and Holder (1999) identified likely causes of home-field advantage. They discussed four factors thought to be responsible for the home-field advantage. These factors were categorized under the general headings of crowd, learning, travel and rule factors. When considering the learning factor, they claimed that the players for the home team were used to play at that particular stadium. There was evidence to suggest that travel factors could be responsible for part of the home-field advantage, provided that the journey involved crossing a number of time zones. However, as most of the matches did not require long travels, these factors were not thought to be a primary cause of home-field advantage. Crowd factors appeared to be the most dominant cause of home-field advantage. Rule factors were found to play a minor role in contributing to home-field advantage.

Clarke and Norman (1995) used least squares to fit a model to the individual match results in English football and found home-field advantage effects on club level. Home-field advantages were calculated for all teams in the English Football League from 1981-82 to 1990-91. It was evident that teams with special facilities had significantly higher home-field advantage, and that London clubs generally had a lower home-field advantage. In addition, they found that the home-field advantage had more impact on the outcome than on goals margin.

Pollard and Pollard (2005) analyzed over 400,000 sports matches since the start of the main professional sports, including English top division football. They found that the home-field advantage in English football decreased from the 1980s to the 2000s. Further, in contrary to Nevill and Holder (1999), they claimed that travel and familiarity contributed to home-field advantage, and found limited support for the contribution of crowd effects.

3.4 Connections to Existing Literature

Although this thesis focuses exclusively on optimizing decisions related to Fantasy Premier League, parallels can be drawn to other scientific areas. The Fantasy Premier League

Decision Problem shares features with many well-studied problems.

Knapsack problems share several similarities with the FPLDP. A capacity (budget) has to be allocated to items characterized by a weight (price) and a reward (expected performance). Kirshner (2011) modeled the problem of selecting free agent players for an NBA squad by using a multiple-choice knapsack model. Here, the rewards were measured in terms of the ability of the players, weights as the salaries and the capacity as the NBA teams' salary cap. Furthermore, Gibson et al. (2007) used a stochastic knapsack problem in order to make team composition decisions regarding player drafts. Ahead of the American sports seasons, teams sign players according to a pick order. Therefore, players become stochastically unavailable over time, depending on the pick decisions of the teams. Thus, they modeled the drafting problem as a stochastic knapsack problem, where the future availability of the items was considered stochastic. There exist significant differences between the knapsack problem and the FPLDP. First, the FPLDP includes not only adding items to a knapsack, it also involves removing items from the knapsack. This is done by buying and selling players in a selected squad. In addition, since the price of a player changes over time, removal decisions do not only free space to the knapsack, it can also yield a reward in terms of increased budget.

Similarities can be found between the FPLDP and staffing problems, where one seeks to compose a set of personnel that satisfies the supply and demand of personnel from different categories (Komarudin et al., 2013; Bruecker et al., 2015). In general, the FPLDP makes decisions regarding buying and selling players, while staffing problems consider hiring and dismissal of personnel. Although the two problems resemble each other in terms of composing a team, there exist significant differences. First, the FPLDP considers players as assets that can increase the value (budget) of the team, and not only as items used to satisfy demand. Furthermore, in the FPLDP, players in the same role are heterogeneous, whereas in most staffing problems individuals in the same category are assumed homogeneous. For instance, Davis et al. (2007) aimed to determine the optimal composition of the pre-hospital medical response team and evaluate the importance of including a doctor to the team. No distinctions between workers in the same category for the staffing problems were made. In the FPLDP, on the other hand, although players in the same position are evaluated according to the same point system, individual factors such as ability of scoring a goal influence the expected performance.

In capacity renewal problems (Chand et al., 2000; Rajagopalan and Soteriou, 1994), machinery is replaced due to damages over time, improved equipment and changes of production (Hopp and Nair, 1994; Adkins and Paxson, 2013). Similarly, players are replaced in the FPLDP, but due to change in expected performance. This can be the result of injuries, suspensions or poor recent performance. However, a significant difference between the problems is that while players are considered assets for maximizing expected points, machinery is considered a necessary item in order to satisfy demand efficiently. Furthermore, in the machinery problems, individual machines are often indistinguishable from one another. Hence, different machines can satisfy the same requirements. For the FPLDP however, players are considered unique as their expected performances are dependent on factors such as team, opponent and recent individual performance.

In transportation and logistics, problems with similar features as the FPLDP can be found. For the long-term vehicle fleet composition (Jabali et al., 2012), different types of vehicles are needed to perform distribution activities, in which size and the mix of the fleet are important decision variables. As for the FPLDP, a mix of players are necessary in order to fulfill the formation criterion. Various vehicles may have the same characteristics and can hence be used for the same purposes. For the FPLDP however, although two players can be listed with the same purchase price, their performance in terms of expected points can be significantly different.

Finally, there is a large correspondence between the FPLDP and portfolio optimization problems (Markowitz (1952); Zenios (2005); Mansini et al. (2015)). The FPLDP can be interpreted as the problem of investing in a number of assets (football players) with stochastic returns in order to maximize total return (the total number of points). Furthermore, the budget limit in the FPLDP can be seen as a budget constraint in portfolio optimization. However, in the FPLDP the decisions are binary (a player is either bought or not) while in portfolio optimization fractions of wealth are allocated to different assets. Moreover, the fixed number of assets (players) in the FPLDP resembles cardinality-constrained portfolio optimization problems (Chang et al., 2000). Additionally, similar to how the formation criterion in the FPLDP gives a fixed number of players in each role, cardinality-constrained portfolio optimization problems may impose a minimum or maximum proportion of wealth to be allocated to a certain class of assets. Finally, variance in returns is a well-known measure of risk in portfolio optimization problems. For the FPLDP, variance can be regarded in a similar manner in order to model the uncertainty in future player performances. Newell and Easton (2017) suggested a stochastic integer program which maximizes the expected payout of a tiered Daily NFL Fantasy Sports by using a stochastic integer program. Here, variance in points obtained was used as a measure of risk.

3.5 Our Contribution

As discussed in this literature review, the FPLDP shares similarities with many well-studied optimization problems. Moreover, several optimization models have been suggested for football team composition as well as for Fantasy Sports (Fry et al. (2007); Gibson et al. (2007); Becker and Sun (2016)). However, to the best of the authors' knowledge, we have developed the first mathematical model for the Fantasy Premier League. Further, we have developed the first optimization driven framework that can compete on the same terms as human FPL managers.

Chapter 4

Model Formulation

In this chapter, the mathematical optimization model for the Fantasy Premier League Decision Problem (FPLDP) is presented. In Section 4.1, the sets, indices, parameters and variables used in the model are defined. In Section 4.2, the mathematical formulation is presented alongside explanations of the objective function and constraints.

The FPLDP is modeled as a stochastic mixed integer linear program. The points a player receives in each gameweek is stochastic and modeled as dependent on future random events during actual matches. Such random events include goals, assists, penalties, passes, saves etc. This is denoted in the model as $\rho_{pt}(\omega_t)$, where ω_t is the realization of these future random events in gameweek t. The objective is to maximize expected points over a number of gameweeks. From the output, the most interesting decisions are the selected squad, starting line-up, substitutes, captain, vice-captain as well as whether to use any gamechips in a gameweek.

The only simplification made from the original problem is the feature of a sell-on fee. Remember, as explained in Section 2.1, that if more managers select a player, the price the player is listed with in the FPL database will increase. Furthermore, a manager can benefit from a price increase if the player is in the selected squad. Thus, a distinction between a player's value and sell price is made. Hence, the prices are stochastic variables. In the model, however, this feature is simplified by assuming that a player's sell price and value are deterministic.

4.1 Definition of Sets, Indices, Parameters and Variables

Set		
$\mathcal{T}$	-	Set of gameweeks.
$\mathcal{P}$	-	Set of players.
$\mathcal{C}$	-	Set of teams.
$\mathcal{L}$	-	Substitution priorities, where 1 is first priority.

Table 4.1: Sets.

Subset			
$\mathcal{P}^D$	-	Subset of defenders,	$\mathcal{P}^D \subset \mathcal{P}$
$\mathcal{P}^M$	-	Subset of midfielders,	$\mathcal{P}^M \subset \mathcal{P}$
$\mathcal{P}^F$	-	Subset of forwards,	$\mathcal{P}^F \subset \mathcal{P}$
$\mathcal{P}^K$	-	Subset of goalkeepers,	$\mathcal{P}^K \subset \mathcal{P}$
$\mathcal{P}_c$	-	Subset of football players in a team c,	$\mathcal{P}_c \subset \mathcal{P}, \quad c \in \mathcal{C}$
$\mathcal{T}_{FH}$	-	Subset of the gameweeks in the first half of the season,	$\mathcal{T}_{FH} \subset \mathcal{T}$
$\mathcal{T}_{SH}$	-	Subset of the gameweeks in the second half of the season.	$\mathcal{T}_{SH} \subset \mathcal{T}$

Table 4.2: Subsets.

Index		
p	-	Players
t	-	Gameweeks
l	-	Priority

Table 4.3: Indices.

Parameters		
$\rho_{pt}(\omega_t)$	-	Points for a player p dependent on the realization of a random event, ω_t, in gameweek t.
ϵ	-	Constant such that $\epsilon << 1$.
κ_l	-	Constants such that $\kappa_l << \epsilon$ for all $l \in \mathcal{L}$, and $\kappa_1 > \kappa_2 > \ldots > \kappa_{\mathcal{L}}$.
C^S_{pt}	-	Sell price of player p in a gameweek t.
C^B_{pt}	-	Value of player p in a gameweek t.
R	-	Number of points deducted for each penalized transfer.
M^K	-	Number of goalkeepers required in the selected squad.
M^D	-	Number of defenders required in the selected squad.
M^M	-	Number of midfielders required in the selected squad.
M^F	-	Number of forwards required in the selected squad.
M^C	-	Maximum number of players allowed in the selected squad from the same team.
E	-	Number of players required in the starting line-up.
E^K	-	Number of goalkeepers required in the starting line-up.
E^D	-	Minimum number of defenders required in the starting line-up.
E^M	-	Minimum number of midfielders required in the starting line-up.
E^F	-	Minimum number of forwards required in the starting line-up.
B^S	-	Starting budget.
β	-	Sufficiently high constant.
$\bar{\alpha}$	-	Sufficiently high constant.
ϕ	-	Number of players which are substitutes.
ϕ^K	-	Number of goalkeepers which are substitutes.
$\overline{Q}$	-	Maximum number of free transfers possible to accumulate over gameweeks.
$\underline{Q}$	-	Number of free transfers given every gameweek.

Table 4.4: Parameters.

Variables		
x_{pt}	-	1 if a player p is picked for the selected squad in a gameweek t, 0 otherwise.
$x_{pt}^{freehit}$	-	1 if a player p is picked for the selected squad in a gameweek t where the Free Hit is used, 0 otherwise.
y_{pt}	-	1 if a player p is picked for starting line-up in a gameweek t, 0 otherwise.
f_{pt}	-	1 if a player p is picked as captain in a gameweek t, 0 otherwise.
h_{pt}	-	1 if a player p is picked as vice-captain in a gameweek t, 0 otherwise.
g_{ptl}	-	1 if a player p is picked as substitution priority l in a gameweek t, 0 otherwise.
u_{pt}	-	1 if a player p is transferred out of the selected squad in a gameweek t, 0 otherwise.
e_{pt}	-	1 if a player p is transferred in to the selected squad in a gameweek t, 0 otherwise.
w_t	-	1 if a Wildcard is used in a gameweek t , 0 otherwise.
c_{pt}	-	1 if the Triple Captain is used on a player p in a gameweek t, 0 otherwise.
b_t	-	1 if the Bench Boost is used in a gameweek t, 0 otherwise.
r_t	-	1 if the Free Hit is used in a gameweek t, 0 otherwise.
λ_{pt}	-	Binary auxiliary variable.
v_t	-	Remaining budget in a gameweek t.
q_t	-	Number of free transfers available in a gameweek t.
α_t	-	Number of penalized transfers in a gameweek t.

Table 4.5: Variables.

4.2 Mathematical Model

4.2.1 Objective Function

$$
\begin{aligned}
\max z = \mathbb{E}\Big[&\sum_{p\in\mathcal{P}}\sum_{t\in\mathcal{T}}\Big(\rho_{pt}(\omega_t)y_{pt} + \rho_{pt}(\omega_t)f_{pt} + \epsilon\rho_{pt}(\omega_t)h_{pt} + \sum_{l\in\mathcal{L}}\kappa_l\rho_{pt}(\omega_t)g_{ptl} \\
&+ 2\rho_{pt}(\omega_t)c_{pt}\Big)\Big] \\
&- R\sum_{t\in\mathcal{T}}\alpha_t
\end{aligned} \tag{4.1}
$$

The objective function maximizes expected points over a set of gameweeks $\mathcal{T}$ for the starting line-up. The term $\rho_{pt}(\omega_t)f_{pt}$ describes who is picked as captain, while $\epsilon\rho_{pt}(\omega_t)h_{pt}$ describes who is picked as vice-captain. Further, $\sum_{l\in\mathcal{L}}\kappa_l\rho_{pt}(\omega_t)g_{ptl}$ describes the substitution priority. Naturally, among the substitutions the player with the highest expected points is given the first priority, the player with the second highest expected points is given the second priority and so on. Hence, κ_l is defined as $\kappa_1 > \kappa_2 > \ldots > \kappa_l$. The Triple Captain awards triple points for one player in a gameweek and is described by the term $2\rho_{pt}(\omega_t)c_{pt}$.

If more transfers are made than the number of free transfers imposed by the rules, R points are deducted for each penalized transfer. This is described by the term $R\sum_{t\in\mathcal{T}}\alpha_t$.

Note that in principle all decision variables for a gameweek t are dependent on the realization of a random event ω_{t-1}, meaning that decisions are based only on information available before gameweek t starts. This ensures that the model is *non-anticipative*, meaning that decisions are based only on available information and do not anticipate the future. However, for the sake of legibility, the dependency is suppressed in the notation.

4.2.2 Constraints

Gamechips

$$\sum_{t\in\mathcal{T}_{FH}} w_t \leq 1 \tag{4.2}$$

$$\sum_{t\in\mathcal{T}_{SH}} w_t \leq 1 \tag{4.3}$$

$$\sum_{p\in\mathcal{P}}\sum_{t\in\mathcal{T}} c_{pt} \leq 1 \tag{4.4}$$

$$\sum_{t\in\mathcal{T}} b_t \leq 1 \tag{4.5}$$

$$\sum_{t\in\mathcal{T}} r_t \leq 1 \tag{4.6}$$

$$w_t + \sum_{p \in \mathcal{P}} c_{pt} + b_t + r_t \leq 1 \qquad t \in \mathcal{T} \tag{4.7}$$

Constraints (4.2) - (4.7) impose the gamechips rules. The first two constraints specify that the Wildcard can maximum be used one time in the first half of the season, and maximum one time in the other half of the season. Constraints (4.4) - (4.6) state that the gamechips Triple Captain, Bench Boost and Free Hit can each be used maximum one time during the season. The last constraints ensure that only one gamechip can be used in a gameweek.

Selected Squad Constraints

$$\sum_{p \in \mathcal{P}^K} x_{pt} = M^K \qquad t \in \mathcal{T} \tag{4.8}$$

$$\sum_{p \in \mathcal{P}^D} x_{pt} = M^D \qquad t \in \mathcal{T} \tag{4.9}$$

$$\sum_{p \in \mathcal{P}^M} x_{pt} = M^M \qquad t \in \mathcal{T} \tag{4.10}$$

$$\sum_{p \in \mathcal{P}^F} x_{pt} = M^F \qquad t \in \mathcal{T} \tag{4.11}$$

$$\sum_{p \in \mathcal{P}_c} x_{pt} \leq M^C \qquad t \in \mathcal{T}, \;\; c \in \mathcal{C} \tag{4.12}$$

Constraints (4.8) - (4.12) handle the number of players required in the selected squad. Explicitly, it is necessary to have M^K goalkeepers, M^D defenders, M^M midfielders and M^F forwards. Constraints (4.12) ensure that the maximum number of players allowed from same club is M^C.

$$\sum_{p \in \mathcal{P}^K} x_{pt}^{freehit} = M^K r_t \qquad t \in \mathcal{T} \tag{4.13}$$

$$\sum_{p \in \mathcal{P}^D} x_{pt}^{freehit} = M^D r_t \qquad t \in \mathcal{T} \tag{4.14}$$

$$\sum_{p \in \mathcal{P}^M} x_{pt}^{freehit} = M^M r_t \qquad t \in \mathcal{T} \tag{4.15}$$

$$\sum_{p \in \mathcal{P}^F} x_{pt}^{freehit} = M^F r_t \qquad t \in \mathcal{T} \tag{4.16}$$

$$\sum_{p \in \mathcal{P}_c} x_{pt}^{freehit} \leq M^C r_t \qquad t \in \mathcal{T}, \;\; c \in \mathcal{C} \tag{4.17}$$

Constraints (4.13) - (4.17) are the same as constraints (4.8) - (4.12), with the exception that they are only active in the gameweek the Free Hit is used. That is, they enforce $x_{pt}^{freehit}$ to only take value in that particular gameweek, while in the other gameweeks they are zero.

Starting Line-up Constraints

$$\sum_{p \in \mathcal{P}} y_{pt} = E + \phi b_t \qquad t \in \mathcal{T} \tag{4.18}$$

$$\sum_{p \in \mathcal{P}^K} y_{pt} = E^K + \phi^K b_t \qquad t \in \mathcal{T} \tag{4.19}$$

$$\sum_{p \in \mathcal{P}^D} y_{pt} \geq E^D \qquad t \in \mathcal{T} \tag{4.20}$$

$$\sum_{p \in \mathcal{P}^M} y_{pt} \geq E^M \qquad t \in \mathcal{T} \tag{4.21}$$

$$\sum_{p \in \mathcal{P}^F} y_{pt} \geq E^F \qquad t \in \mathcal{T} \tag{4.22}$$

Constraints (4.18) - (4.22) enforce that the formation criterion is upheld in each gameweek. Constraints (4.18) handle the number of players in the starting line-up, while constraints (4.19) - (4.22) ensure that the minimum number of players in each position is upheld. When the Bench Boost is used, all the substitutes are treated as a part of the starting line-up. This is ensured by the term ϕb_t and $\phi^K b_t$ in constraints (4.18) and (4.19), respectively.

$$y_{pt} \leq x_{pt}^{freehit} + x_{pt}(1 - r_t) \qquad p \in \mathcal{P}, \;\; t \in \mathcal{T} \tag{4.23}$$

Constraints (4.23) enforce that only players in the selected squad can be picked for the starting line-up. Remember, when the Free Hit is used, the whole team can be changed for one gameweek, but in the next gameweek it changes back. This means that the gameweek the Free Hit is used $r_t = 1$ and y_{pt} only takes value from $x_{pt}^{freehit}$, not x_{pt}. When the Free Hit is not used, i.e., $r_t = 0$, the constraints (4.13) - (4.17) ensure that $x_{pt}^{freehit}$ does not take value. Constraints (4.23) are non-linear. A linearization is given by constraints (4.24) - (4.26) where λ_{pt} is an auxiliary binary variable.

$$y_{pt} \leq x_{pt}^{freehit} + \lambda_{pt} \qquad p \in \mathcal{P}, \;\; t \in \mathcal{T} \tag{4.24}$$

$$\lambda_{pt} \leq x_{pt} \qquad p \in \mathcal{P}, \;\; t \in \mathcal{T} \tag{4.25}$$

$$\lambda_{pt} \leq (1 - r_t) \qquad p \in \mathcal{P}, \;\; t \in \mathcal{T} \tag{4.26}$$

When the Free Hit is used, $x_{pt}^{freehit}$ decides what players y_{pt} can select. In all other gameweeks, λ_{pt} enforces that x_{pt} decides what players y_{pt} can select. More specifically, when the Free Hit is used, constraints (4.26) force λ_{pt} to be zero for all players. Further, constraints (4.13) - (4.17) become active such that $x_{pt}^{freehit}$ equals 1 for the Free Hit players selected. Finally, constraints (4.24) impose that $x_{pt}^{freehit}$ decides what players y_{pt} can select. In all the other gameweeks, $x_{pt}^{freehit}$ is zero. Constraints (4.26) ensure that λ_{pt} can take non-zero values for all players in these gameweeks. Moreover, constraints (4.25) make sure λ_{pt} is only 1 for players in the selected squad decided by x_{pt}. Constraints (4.24) ensure x_{pt} decides what players y_{pt} can select.

Captain and Vice-captain Constraints

$$\sum_{p \in \mathcal{P}} f_{pt} + \sum_{p \in \mathcal{P}} c_{pt} = 1 \qquad t \in \mathcal{T} \tag{4.27}$$

$$\sum_{p \in \mathcal{P}} h_{pt} = 1 \qquad t \in \mathcal{T} \tag{4.28}$$

$$f_{pt} + c_{pt} + h_{pt} \leq y_{pt} \qquad p \in \mathcal{P},\ t \in \mathcal{T} \tag{4.29}$$

Constraints (4.27) - (4.29) set restrictions for captain and vice-captain in the starting line-up. Constraints (4.27) ensure that there can only be 1 captain in a gameweek. They also make sure that in the gameweek the Triple Captain is used, no captain is picked. Further, constraints (4.28) ensure that there can only be 1 vice-captain in a gameweek. The last constraints (4.29) enforce that a player can only be a captain, triple captain or vice-captain in a gameweek.

Substitution Constraints

$$y_{pt} + \sum_{l \in \mathcal{L}} g_{ptl} \leq x_{pt} + \beta r_t \qquad p \in \mathcal{P} \setminus \mathcal{P}^K,\ t \in \mathcal{T} \tag{4.30}$$

$$y_{pt} + \sum_{l \in \mathcal{L}} g_{ptl} \leq x_{pt}^{freehit} + \beta(1 - r_t) \qquad p \in \mathcal{P} \setminus \mathcal{P}^K,\ t \in \mathcal{T} \tag{4.31}$$

$$\sum_{p \in \mathcal{P} \setminus \mathcal{P}^K} g_{ptl} \leq 1 \qquad t \in \mathcal{T},\ l \in \mathcal{L} \tag{4.32}$$

Constraints (4.30) - (4.32) explain how the substitutions are handled. Both constraints (4.30) and (4.31) ensure that all the players who are not picked in the starting line-up are regarded as substitutes. A parameter β is a sufficiently high value, such that when the Free Hit is exercised, constraints (4.31) are restricting. In all other gameweeks, constraints (4.30) are restricting. Constraints (4.32) ensure that no substitutes have more than one priority.

Budget Constraints

$$B^S - \sum_{p \in \mathcal{P}} C_{p1}^B x_{p1} = v_1 \tag{4.33}$$

$$v_{t-1} + \sum_{p \in \mathcal{P}} C_{pt}^S u_{pt} - \sum_{p \in \mathcal{P}} C_{pt}^B e_{pt} = v_t \qquad t \in \mathcal{T} \setminus \{1\} \tag{4.34}$$

$$x_{p,(t-1)} + e_{pt} - u_{pt} = x_{pt} \qquad p \in \mathcal{P},\ t \in \mathcal{T} \setminus \{1\} \tag{4.35}$$

$$e_{pt} + u_{pt} \leq 1 \qquad p \in \mathcal{P},\ t \in \mathcal{T} \tag{4.36}$$

Constraints (4.33) - (4.36) handle the budget flow and how transfers affect the budget. In the first gameweek, constraints (4.33) assign the remaining budget of the starting budget

B^S to the variable v_1. Next, constraints (4.34) handle the flow in the variable v_t for the remaining gameweeks. More specifically, the remaining budget for the next gameweek is decided by the remaining budget from the previous gameweek plus the difference in budget generated by transferring players in and out. Further, constraints (4.35) ensure that a player in the selected squad in gameweek t was either in the selected squad in the previous gameweek or transferred in. Constraints (4.36) make sure that a player cannot be both transferred in and transferred out in the same gameweek.

$$\sum_{p \in \mathcal{P}} C^S_{pt} x_{p,(t-1)} + v_{t-1} \geq \sum_{p \in \mathcal{P}} C^B_{pt} x^{freehit}_{pt} \qquad t \in \mathcal{T} \setminus \{1\} \tag{4.37}$$

$$\sum_{p \in \mathcal{P}} u_{pt} \leq E(1 - r_t) \qquad t \in \mathcal{T} \setminus \{1\} \tag{4.38}$$

$$\sum_{p \in \mathcal{P}} e_{pt} \leq E(1 - r_t) \qquad t \in \mathcal{T} \setminus \{1\} \tag{4.39}$$

Constraints (4.37) ensure that the cost of the selected squad the gameweek the Free Hit is played does not exceed the value of the initially selected squad and remaining budget from the previous gameweek. Constraints (4.38) and (4.39) ensure that the selected squad in the gameweek after the Free Hit is played is the same as in the gameweek before.

Transfer Constraints

$$q_2 = \underline{Q} \tag{4.40}$$

$$Ew_t + q_t - \sum_{p \in \mathcal{P}} e_{pt} + \underline{Q} + \alpha_t \geq q_{t+1} \qquad t \in T \setminus \{1\} \tag{4.41}$$

$$\bar{\alpha}(\overline{Q} - q_{t+1}) \geq \alpha_t \qquad t \in \mathcal{T} \tag{4.42}$$

$$q_{t+1} \geq \underline{Q} \qquad t \in \mathcal{T} \tag{4.43}$$

$$q_{t+1} \leq \overline{Q} + (\underline{Q} - \overline{Q})w_t + (\underline{Q} - \overline{Q})r_t \qquad t \in \mathcal{T} \tag{4.44}$$

Constraints (4.40) - (4.44) handle the transfers. The first constraints (4.40) set the number of free transfers in the second gameweek to $\underline{Q}$. This is the first time the decision of transferring players is considered and consequently the first time the limitation of free transfers is treated. Constraints (4.41) handle free transfers available in the next gameweek. Each gameweek, $\underline{Q}$ free transfers are given. If they are not used, they are moved to the next gameweek. Further, the maximum number of free transfers is denoted by $\overline{Q}$. If the number of players transferred out exceeds the number of free transfers available, the variable α_t takes a value and each penalized transfer gets punished by R points in the objective function. The exception is when the Wildcard is played. In this case, the number of free transfers is unlimited. This is ensured by the term Ew_t, which makes sure that α_t does not take a value in the gameweek the Wildcard is used. Constraints (4.42) handle a special case in the modeling; it makes sure that α_t only can take the value of zero when q_t hits its upper limit, $\overline{Q}$. This will be further illustrated by an example where $\overline{Q} = 2$.

t	e_t	α_t	q_t	q_{t+1}
1	5	5	1	2
2	6	4	2	1

Solution 1.

t	e_t	α_t	q_t	q_{t+1}
1	5	4	1	1
2	6	5	1	1

Solution 2.

In the tables above, two solutions are presented with variables related to constraints (4.41). Without constraints (4.42), solution 1 and 2 are regarded as the same since the total number of penalized transfers over the gameweeks are equivalent. However, in practice only solution 2 is right. In both solutions, there is 1 free transfer available and 5 players are transferred out in the first gameweek. This means that α_1 is 4 and that free transfers next gameweek, q_2, is 1. Without constraints (4.42), it is possible that α_1 takes the value 5, q_2 takes its upper limit $\overline{Q} = 2$ and α_2 takes the value 4. However, constraints (4.42) ensure that α_t only take the value 0 when q_{t+1} hits its upper limit, $\overline{Q}$.

The last constraints (4.43) and (4.44) impose an upper and lower bound on the number of free transfers available in a gameweek. The terms $(\underline{Q} - \overline{Q})w_t$ and $(\underline{Q} - \overline{Q})r_t$ ensure that there are only $\underline{Q}$ free transfers the gameweek after the gamechips Wildcard and Free Hit are used.

Binary, Integer and Non-negativity Constraints

$$x_{pt} \in \{0,1\} \qquad p \in \mathcal{P},\ \ t \in \mathcal{T} \tag{4.45}$$

$$x_{pt}^{freehit} \in \{0,1\} \qquad p \in \mathcal{P},\ \ t \in \mathcal{T} \tag{4.46}$$

$$y_{pt} \in \{0,1\} \qquad p \in \mathcal{P},\ \ t \in \mathcal{T} \tag{4.47}$$

$$f_{pt} \in \{0,1\} \qquad p \in \mathcal{P},\ \ t \in \mathcal{T} \tag{4.48}$$

$$h_{pt} \in \{0,1\} \qquad p \in \mathcal{P},\ \ t \in \mathcal{T} \tag{4.49}$$

$$g_{ptl} \in \{0,1\} \qquad p \in \mathcal{P},\ \ t \in \mathcal{T},\ \ l \in \mathcal{L} \tag{4.50}$$

$$u_{pt} \in \{0,1\} \qquad p \in \mathcal{P},\ \ t \in \mathcal{T} \tag{4.51}$$

$$e_{pt} \in \{0,1\} \qquad p \in \mathcal{P},\ \ t \in \mathcal{T} \tag{4.52}$$

$$w_t \in \{0,1\} \qquad t \in \mathcal{T} \tag{4.53}$$

$$c_{pt} \in \{0,1\} \qquad p \in \mathcal{P},\ \ t \in \mathcal{T} \tag{4.54}$$

$$b_t \in \{0,1\} \qquad t \in \mathcal{T} \tag{4.55}$$

$$r_t \in \{0,1\} \qquad t \in \mathcal{T} \tag{4.56}$$

$$\lambda_{pt} \in \{0,1\} \qquad p \in \mathcal{P},\ \ t \in \mathcal{T} \tag{4.57}$$

$$v_t \geq 0 \qquad t \in \mathcal{T} \tag{4.58}$$

$$q_t \in \mathbb{Z}^+ \qquad t \in \mathcal{T} \tag{4.59}$$

$$\alpha_t \in \mathbb{Z}^+ \qquad t \in \mathcal{T} \tag{4.60}$$

Constraints (4.45) - (4.60) impose binary, integer and non-negativity constraints on the decision variables.

Chapter 5

Solution Framework

The FPLDP is a complex problem with a high degree of uncertainty. In each gameweek, a broad range of decisions must be optimized. Compared to other Fantasy Sports, the gamechips make it even more complicated. If the model presented in Chapter 4.2 is run ex-post, it yields the optimal solution as the realized points are known. However, in order to compete in FPL, decisions must be made ex-ante. Thus, estimations of expected points and strategies for when the gamechips should be played are required. The aim of this chapter is to propose a solution framework designed to maximize performance over the entire season. By a framework, a methodology for competing in the FPL on the same terms as managers are meant. Thus, the framework includes a solution method for the mathematical model, a method for generating forecasts of player points and strategies for the use of gamechips. In addition, an approach to handle variance in, and correlation between, player points is included in the framework. This procedure is hereby referred to as risk handling.

As seen throughout this chapter, the framework is intended to be solely data-driven. That is, a participant in the FPL applying the framework does not have to make any decision based on his or her own judgment. Thus, the inherent skill and knowledge of the FPL possessed by a participant applying the framework are essentially irrelevant for its performance.

In Section 5.1, a solution method suggested for solving the mathematical model is presented. In Section 5.2, three different methods for forecasting player performance are described. Further, in the Sections 5.3 and 5.4, respectively, methods for implementing gamechips and handling risk are proposed. Finally, in Section 5.5, the elements discussed throughout the chapter are combined to a solution framework called the *forecast-based optimization model*.

5.1 Solution Method

As stochastic programming models seek to leverage probability distribution of uncertain parameters, they are regarded as more computationally demanding than deterministic programming models (Shapiro and Philpott, 2007). A conventional method for solving a stochastic model is by the use of scenario generation. For the FPLDP, the number of scenarios explodes even for a modest number of scenarios in each gameweek. With only two scenarios in a gameweek, the number of scenarios already reaches over 1 million for 20 gameweeks ($2^{20} = 1\,048\,576$). Therefore, solving a stochastic model over a whole season is computationally infeasible. Hence, the FPLDP is solved with a deterministic approach. A rolling horizon framework is used where decisions are implemented for one gameweek at a time. When decisions are optimized for a given gameweek, a certain number of gameweeks ahead is considered and forecasts of player points are generated.

When forecasts are used, the objective function described in Section 4.2 is simplified from:

$$\begin{aligned}\max z = \mathbb{E}\Big[\sum_{p\in\mathcal{P}}\sum_{t\in\mathcal{T}}\Big(&\rho_{pt}(\omega_t)y_{pt} + \rho_{pt}(\omega_t)f_{pt} + \epsilon\rho_{pt}(\omega_t)h_{pt} + \sum_{l\in\mathcal{L}}\kappa_l\rho_{pt}(\omega_t)g_{ptl}\\ &+ 2\rho_{pt}(\omega_t)c_{pt}\Big)\Big]\\ &- R\sum_{t\in\mathcal{T}}\alpha_t\end{aligned}$$

to

$$\begin{aligned}\max z = \sum_{p\in\mathcal{P}}\sum_{t\in\mathcal{T}}&\Big(\hat{\rho}_{pt}y_{pt} + \hat{\rho}_{pt}f_{pt} + \epsilon\hat{\rho}_{pt}h_{pt} + \sum_{l\in\mathcal{L}}\kappa_l\hat{\rho}_{pt}g_{ptl}\\ &+ 2\hat{\rho}_{pt}c_{pt}\Big)\\ &- R\sum_{t\in\mathcal{T}}\alpha_t\end{aligned}$$

where $\hat{\rho}_{pt}$ denotes an estimation of expected points for each player p in a gameweek t.

5.1.1 Rolling Horizon

The idea behind a rolling horizon heuristic is to split the problem into several sub-problems along a time axis, and then solve them in a chronological order taking into account the solution of the previous sub-problem. The heuristic suits problems with a long planning horizon where there are limited dependencies between decisions made early and late in the planning horizon. Industrial planning and scheduling problems are common areas where the heuristic is used. The most important considerations are the number of sub-problems versus the length of each sub-problem, and also which decisions to fix in each sub-problem.

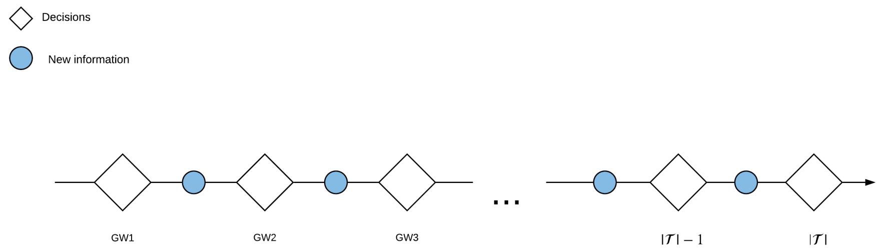

Figure 5.1: Information flow in the FPLDP.

Rolling horizon is an appropriate solution method for FPLDP as managers are provided new information and face new decisions every gameweek. That is, each gameweek is a sub-problem. As shown in Figure 5.1, decisions are made based on the information available. When a gameweek is played, updated information such as realized points, injuries and suspensions are made available. Hence, new decisions are made accordingly. Consequently, the problem is divided into $|\mathcal{T}|$ number of sub-problems along a time axis. That is, the decisions taken in a gameweek affect decisions in future gameweeks. However, decisions made early in the season have a limited impact on decisions made late in the season. It is fully possible to have a whole different selected squad in the last gameweeks compared to that of the first gameweeks. This is especially evident when considering the opportunity if playing the Wildcard.

In the rolling horizon heuristic, each sub-problem is solved over shorter sub-horizons. This corresponds well to the FPLDP, as FPL managers often consider several gameweeks when making decisions. Each sub-horizon is split in two time periods (TP). The two *TP*s are TP_k^C, which is the *central period* that will be implemented in sub-problem k, and the *forecasting period* TP_k^F. Before solving sub-problem $k+1$, variables x, v and q are frozen prior to $k+1$ and used as input. The gameweek in which decisions are already implemented is regarded as the *fixed period*. The sub-horizon is shifted so that the first part of TP_k^F becomes TP_{k+1}^C. If the central period and the forecasting period have equal length, the forecasting period becomes the new central period. To summarize, three periods are considered:

(i) **Fixed period (FP)**. A gameweek where decisions have been implemented and used as input for the next sub-problem. These decisions are the players that are selected in a gameweek, x, the remaining budget, v, and the number of free transfers available, q.

(ii) **Central period (CP)**. A gameweek where new decisions are implemented. The decisions are which players to include in the selected squad, x, determining the starting line-up, y, the choice of captain, f, vice-captain, h, whether to use a gamechip or not and which substitution priority, g, to assign.

(iii) **Forecasting period (PP)**. To avoid making too myopic decisions, a predefined number of gameweeks are selected and input data such as expected points are generated based on information available at that point in time. Decisions are made for these gameweeks, but not implemented. The decisions are subject to change when new and updated information becomes available.

The FP, CP and PP in a sub-problem are illustrated graphically in Figure 5.2. In the figure, the predicted part overlaps the central part in order to illustrate that the input data in the central part is also predicted. The sub-horizon is set to 4 gameweeks in the figure.

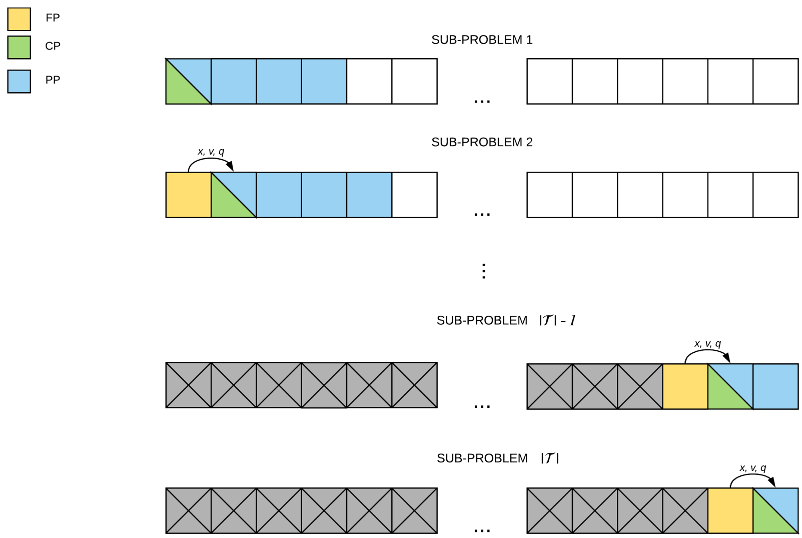

Figure 5.2: Fixed-, central- and forecasting period in a sub-problem.

An important decision is that of determining the length of the sub-horizon. Due to uncertainty in the forecasts, it is not necessarily ideal for the sub-horizon to be the remaining length of the season. Thus, several alternatives concerning the length of the sub-horizon should be tested in order to see what gives the best result. In Chapter 6, setting an appropriate length for the sub-horizon is further discussed. Note that in the framework, the aspect of value, sell price and sell-on fee with regards to a player's price is implemented. However, no attempt is made to forecast future prices. The prices are simply assumed to be what they are at the beginning in the entire sub-horizon. For example, if a player is bought for £9.0m in a sub-problem, his sell price is set to this for the sub-problem's entire sub-horizon. When solving for the next sub-problem, if his value increases to £9.2m, his sell price is updated to £9.1m and set to this for the entire sub-horizon.

5.2 Forecast of Player Points

To generate input for the forecasting period in the rolling horizon heuristic, forecasts of player points are required. In this thesis, three forecasting methods are considered. The methods are:

- **Modified Average**. The forecasting approach suggested by Bonomo et al. (2014) is adapted and modified for the FPL.
- **Regression**. From a pool of explanatory variables related to the FPL, variables are selected and a linear regression model is fitted.
- **Odds**. Bookmakers' odds are transformed into a probability and subsequently used to forecast player points. Odds related to factors such as scorelines, goals and assists are utilized in order to estimate expected points.

5.2.1 Modified Average

As discussed in the literature review, Bonomo et al. (2014) suggested a solution approach for the Argentinian Fantasy League and achieved promising results. Therefore, the first forecast method is based on their solution approach. However, the methodology is modified in order to make it applicable for the English Fantasy Premier League rather than the Argentinian Fantasy League. Moreover, some alterations are made as an effort to improve the method. Hence, the first forecasting method is called the Modified Average method.

In Section 3.3.1, it is explained that Bonomo et al. (2014) used the average points for the three last gameweeks to forecast player points in the upcoming gameweek. The forecast was further weighted by four factors: opponent strength factor, home-field advantage, whether a player is on a good performance streak and whether a player is assumed to be in the starting line-up. More specifically, Bonomo et al. (2014) estimated a player's expected points for the next gameweek according to Equation 5.1. In the following, it is explained how each factor of the equation was decided by Bonomo et al. (2014) and which modifications that are made.

$$\begin{aligned}\hat{\rho} = &\ \textit{Tracking average} \times \textit{Opponent strength factor} \times \textit{Field factor} \\ &\times \textit{Point streak factor} \times \textit{Starting line-up factor}\end{aligned} \tag{5.1}$$

Tracking Average and Opponent Strength Factor

Bonomo et al. (2014) calculated a *tracking average* by averaging the points obtained in the 3 previous gameweeks. The first modification suggested is to consider averages over more than 3 previous gameweeks. The number of previous gameweeks taken into consideration when forecasting is hereby referred to as the *track-length*. In addition, Bonomo et al. (2014) optimized for only 1 gameweek ahead when decisions were set in a current gameweek. Thus, Bonomo et al. (2014) used a sub-horizon of 1 gameweek and a track-length of 3 gameweeks when solving for the Argentinian Fantasy League.

Considering different track-lengths is not the only modification made with respect to the tracking average. Another adjustment is to weigh previous points based on the strength of the opposition faced in the previous matches. Bonomo et al. (2014) did not account for the strength of the opposition faced when averaging previous points. However, they used league table position for measuring the relative strength between adversaries when predicting future points. That is, when they determined *opponent strength factor*. Yet, league position suffers from numerous drawbacks which makes it unreliable for prediction. For example, the league table suffers from high variation in the early stage of the season and low variation in the late stage of the season. Further, competing teams may not have played an equivalent number of matches at each stage of the season due to postponements. Additionally, at a given point in time, the league table does not capture who has played against who during the season so far. Hence it is inherently biased until the final match of the season has been played (Constantinou and Fenton, 2012). Therefore, instead of weighting the teams based on their league position, measuring team strength with basis in Elo ratings are suggested.

The Elo system, briefly introduced in Section 3.3.2, is a rating system used for measuring the relative strength level of sports teams and individual athletes. The Elo ratings provide a sophisticated measurement of the relative strength between teams. Further, as the ratings are updated once a match has been played, the Elo system considers the aspect of form. In addition, the Elo ratings can incorporate all fixtures, not only league matches. In Appendix B, it is shown how Elo ratings are transformed into a factor of team strength.

Once Elo ratings are obtained, they can be used in order to weight the previous points of each Fantasy Premier League player. In Table 5.1, some Elo ratings for the first gameweeks of the English Premier League 2016/2017 season are listed. Note that the ratings are calculated ahead of each gameweek so that the ratings in the GW1 (gameweek 1) column are the Elo ratings of the teams before gameweek 1 had been played.

		Elo rating			
	Team	GW1	GW2	GW3	GW4
1	Chelsea	1793	1800	1798	1804
2	Tottenham	1804	1804	1800	1798
3	Man City	1848	1856	1858	1863
4	Man Utd	1789	1797	1799	1804
-					
-					
-					
17	Watford	1624	1631	1618	1612
18	Hull	1589	1603	1613	1608
19	Middlesbrough	1595	1597	1601	1604
20	Sunderland	1655	1654	1636	1641

Table 5.1: Elo ratings for Premier League 2016/2017.

Ratings from Table 5.1 can be used in order to weight the previous points of a player. For instance, receiving 9 points against Chelsea is more impressive than receiving 9 points against Sunderland based in their respective Elo ratings. Hence, a suggestion is to weight the points accordingly. Equation 5.2 shows how a player's tracking average is computed when the track-length is m and realized points in gameweek i is denoted ρ_i.

$$\textit{Tracking average} = \frac{1}{m}\sum_{i=1}^{m}\rho_i \cdot \frac{(\text{Opponent Elo rating})_i}{(\text{Team Elo rating})_i} \tag{5.2}$$

Furthermore, Elo ratings, rather than values based on the league table, are used to determine opponent strength factor, as seen in Equation 5.3. Thus, both previous points and future expected points account for faced and upcoming opposition. Notice that in Equation 5.2, Elo ratings for both opponent and team is denoted with i. This is done because the Elo ratings are updated each gameweek.

$$\textit{Opponent strength factor} = \frac{\text{Team Elo rating}}{\text{Opponent Elo rating}} \tag{5.3}$$

Field Factor

As pointed out in Section 3.3.3, numerous studies have concluded that a home-field advantage exists in football. That is, football teams have a tendency of performing better when playing at their home-field than when playing at an opponent's ground. In order to distinguish between home-field advantage and away-field (dis)advantage, the variable that takes such an advantage into account is called the *field factor*. Bonomo et al. (2014) set the home-field factor to 1.05 and away-field factor to 0.95 without further elaboration.

Studies of home-field advantages in a football, e.g., Pollard and Pollard (2005), usually concentrate on the outcome of the game, i.e., win, draw or loss. In Fantasy Premier

League, however, the outcome of a match has no direct impact on the point system. Further, an important aspect influencing points in FPL is goals scored. Thus, the field factor is calculated with basis in goals. Further, the field factor is calculated for all teams in the Premier League combined, not for each team specifically. The idea is that the field factor alongside the relative team strengths calculated by the Elo ratings will complement each other. The home-field factor for a particular season is calculated as:

$$\textit{Home-field factor} = \frac{\text{Total number of goals scored home}}{\text{Total number of goals scored}}$$

Similarly, the away-field factor is calculated as:

$$\textit{Away-field factor} = \frac{\text{Total number of goals scored away}}{\text{Total number of goals scored}}$$

As seen later in Section 6.2.1, several previous seasons are considered in order to find an appropriate empirical value for the field factor.

Point Streak Factor

Bonomo et al. (2014) adjusted their forecasts with a value between 0.95 and 1.05 to account for a player's point streak. They did not further describe how the *point streak factor* was determined. In this thesis, a player's point streak factor is determined in the following way: if a player receives more than X points for 2 gameweeks in a row, his point streak factor will be set to 1.01. Further, if the streak continues, the factor will increase by 0.01 for each match that his points exceed X number of points. When a player has received more than X points 6 gameweeks in a row his point streak factor is set to the maximum of 1.05. Hence, the value does not increase if his streak exceeds 6 matches. The same rule applies to a player scoring less than Y points 2 gameweeks in a row. In this case, his point streak factor will be set to 0.99. When obtaining less than Y points 6 matches in a row the point streak factor is set to a minimum of 0.95. Once a streak is broken, the player's point streak factor is set to 1. In Section 6.2.1, a method for determining X and Y is proposed.

Starting Line-up Factor

Bonomo et al. (2014) decided the *starting line-up factor* by considering coaches' announcements, press reports and information posted on the website of the Argentinian Fantasy League. Such considerations are of qualitative nature. As the solution framework is intended to be data-driven, an explicit attempt at implementing this factor for the FPL is not undertaken. Instead, as seen later in Section 6.1, all forecasts are adjusted by data concerning injuries and suspensions.

Summarizing Numerical Example

In order to summarize the Modified Average method, a numerical example for a track length of 2 for Manchester United's Paul Pogba is provided. In Tables 5.2 and 5.3, fictive values for realized points and Elo ratings are given.

	Realized points		
Opponent	GW1	GW2	GW3
Chelsea	6		
Sunderland		9	
Watford			?

Table 5.2: Realized points for Paul Pogba.

Team	GW1	GW2	GW3
Man Utd	1789	1797	1799
Chelsea	1793		
Sunderland		1654	
Watford			1618

Table 5.3: Elo ratings for three gameweeks.

Thus, Paul Pogba's tracking average can be calculated as:

$$\text{Tracking Average} = \frac{1}{2} \cdot (6 \cdot \frac{1793}{1789} + 9 \cdot \frac{1654}{1797}) \approx 7.15$$

and the opponent strength factor can be found as:

$$\frac{1618}{1799} \approx 0.90$$

Factor	Value
Tracking average	7.15
Opponent strength factor	0.90
Field factor	1.5
Points streak factor	1
Starting line-up factor	1

Table 5.4: Numerical values for Equation 5.1.

Finally, all factors required to calculate the expected points for Paul Pogba in the upcoming gameweek are summarized in Table 5.4. Thus, in compliance with Equation 5.1, Pogba's expected points for the upcoming gameweek are according to the Modified Average method:

$$7.15 \cdot 0.90 \cdot 1.5 \cdot 1 \cdot 1 \approx 9.7$$

5.2.2 Regression

As seen in the literature review, regression is commonly in research related to sports (Bjertnes et al. (2016); Yang (2015)). Therefore, as a second way of forecasting Fantasy Premier League points, a regression-based method is proposed. In order to create a regression model, the following steps are undertaken:

(i) separate players based on position,

(ii) perform a variable selection,

(iii) fit a linear regression model.

Variables

In order to use regression to estimate expected points, numerous numerical and categorical explanatory variables are considered. In the following, all variables considered are presented:

Dependent Variable

- **Realized points** - The dependent variable. The realized points are the actual points gained by players.

Explanatory Variables

Categorical Variables

- **Team** - Which teams the players play for. This variable can change during the international transfer window in August and January.
- **Position** - Describes whether a player is a goalkeeper, defender, midfielder or forward. The variable is constant throughout the season.
- **Home/Away** - Displays whether a match is played at home or away.

Numerical Variables

- **Previous points** - Points obtained in each of the previous gameweeks.
- **Price** - The purchase price (value) of a player.
- **Transfers in** - The amount of FPL managers that bought the player ahead of a given gameweek.
- **Transfers out** - The amount of FPL managers that sold the player ahead of a given gameweek.
- **Minutes Played** - Minutes played in the previous gameweek.
- **Yellow Cards** - Total amount of yellow cards received in the previous gameweeks.
- **Red Cards** - Total amount of red cards received in the previous gameweeks.
- **Goals** - Total amount of goals scored in the previous gameweeks.
- **Assists** - Total amount of assists in the previous gameweeks.
- **Penalties Missed** - Total amount of penalties missed in the previous gameweeks.
- **Penalties Saved** - Total number of penalties saved by a goalkeeper in the previous gameweeks.
- **Saves** - Total amount of saves made by a goalkeeper in the previous gameweeks.

- **Clean sheets** - Total number of matches a player was awarded a clean sheet in the previous gameweeks. As explained in Chapter 2, clean sheets are only awarded goalkeepers, defenders and midfielders who play a minimum of 60 minutes during the match.

Categorization by Position

As described in Chapter 2, points are awarded based on different criteria to players in different *positions*. For instance, goalkeepers, defenders and midfielders are rewarded points for keeping a clean sheet, but forwards are not. Based on this knowledge, it is decided to categorize the players based on their respective positions before continuing the regression procedure. If all available variables were regressed, such a distinction would perhaps not be necessary, as the coefficient related to the categorical variable position could compensate the differences. However, the aim is to limit the number of explanatory variables for each position by using variable selection. Therefore, the categorization is done.

In order to provide further justification for the categorization based on position, a box-plot and scatter-plots are assessed in Section 6.2.2. For example, in the box-plot, it is investigated whether the empirical distribution of points obtained differs between positions. Moreover, the scatter plots are visually inspected to see the relation between, for example, total points, goals and position.

Times Series Effect of Points

The presence of a scoring streak, or a ”hot hand”, has long been a topic of research (Bar-Eli et al. (2006); Tversky and Gilovich (1989)). If it is present in FPL, points gained in the most previous matches could be an interesting explanatory variable. If not, aggregated variables such as the total number of goals scored can be just as predictive. In order to assess if elements of a time series are correlated with each other, *autocorrelation* is a reasonable metric to asses. Just as correlation measures the extent of a linear relationship between two variables, autocorrelation measures the linear relationship between lagged values of a time series (Hyndman and Athanasopoulos, 2018). For example, r_1 measures the relationship between y_t and y_{t-1}. In general, there exist many autocorrelation coefficients, and the value of r_k can be written as in Equation 5.4.

$$r_k = \frac{\sum_{t=k+1}^{T}(y_t - \bar{y})(y_{t-k} - \bar{y})}{\sum_{t=1}^{T}(y_t - \bar{y})^2} \tag{5.4}$$

To test for autocorrelation, the Ljung-Box test and the Durbin-Watson test are conducted. The Ljung-Box test is a statistical test that often is conducted in order to test for autocorrelation in a time series for different lags (Hyndman and Athanasopoulos, 2018). The Durbin-Watson test is another test frequently conducted to test for the presence of autocorrelation for the first lag (Alexander, 2008). A detailed description of the tests can be found in Hyndman and Athanasopoulos (2018) and Alexander (2008), respectively.

Variable Selection

In order to select variables, *k-fold cross validation* is performed in combination with a *lasso regression*. Before these topics are treated, a brief discussion of training- and test sets and forecast error measurement is presented.

Training- and Test Sets

In general, in order to determine the accuracy of a forecast, out-of-sample forecasts should be evaluated (Hyndman and Athanasopoulos, 2018). To accomplish this, available data should be separated into two portions: training- and test data. A model is first fitted using the training data only. Then, the model's forecast accuracy is determined by comparing forecasts from the model with actual realizations contained in the test data. Furthermore, in general, if a model is fit on all available data, rather than separated into a training- and test set, the model would be subject to over-fitting (Hyndman and Athanasopoulos, 2018). That is, the model performs well on the sample data but does not necessarily forecast well out of sample.

Forecast Error Measurement

A forecast error is the difference between an observed value and its forecast (Hyndman and Athanasopoulos, 2018), as defined as in Equation 5.5. Here, $\{y_1, ..., y_T\}$ is the training set data, $\{y_{T+1}, , y_{T+2}, ...\}$ is the test set data and $\hat{y}_{t+h|T}$ is the forecast at time *T*+*h*, given *T*.

$$e_{T+h} = y_{T+h} - \hat{y}_{T+h|T} \tag{5.5}$$

A commonly used error measure is the Mean Squared Error (MSE, Equation 5.6) (Hyndman and Athanasopoulos, 2018).

$$MSE = mean(e_t^2) \tag{5.6}$$

k-fold Cross Validation

k-fold cross validation involves randomly dividing a set of observations into *k* groups, or *folds*, of equal size (Gareth et al., 2017). The first fold is treated as the test set, and the remaining $k-1$ folds are used to train a model. The procedure is repeated *k* times, and for each repetition, a new fold makes up the test set. The process results in *k* estimates of the test-error, MSE_1, MSE_2 ..., MSE_k. Finally, a *k*-fold average MSE, denoted $CV_{(k)}$ is computed by averaging the values, as shown in Equation 5.7 (Gareth et al., 2017).

$$CV_{(k)} = \frac{1}{k}\sum_{i=1}^{k} MSE_i \tag{5.7}$$

Lasso Regression

As seen previously, a number of explanatory variables are available. If all were to be used to build a regression model, the model could be subject to over-fitting (Gareth et al., 2017).

Therefore, in order to limit the number of explanatory variables, a *variable selection* is undertaken. That is, a subset of variables believed to be related to the dependent variable *realized points*, is selected. To perform the variable selection, lasso regression is used. In this thesis, only a brief introduction of the concept of lasso regression is given. For a more comprehensive review, the reader is referred to "An Introduction to Statistical Learning" by Gareth et al. (2017).

According to Gareth et al. (2017), lasso regression is a *shrinkage method* used in order to perform variable selection. By shrinkage method, the authors refer to methods were coefficients are sunken towards zero relative to least squares estimates. This shrinkage has the effect of reducing variance. Depending on what type of shrinkage is performed, some of the coefficients may be estimated to be exactly zero. Lasso regression is such a shrinkage method and can thus perform variable selection. The aim of lasso regression is to find the lasso coefficients $\hat{\beta}^L_\lambda$ by minimizing the expression found in Equation 5.8:

$$\sum_{i=1}^{n}\Big(y_i - \beta_0 - \sum_{j=1}^{p}\beta_j x_{ij}\Big)^2 + \lambda\sum_{j=1}^{p}|\beta_j| = RSS + \lambda\sum_{j=1}^{p}|\beta_j| \tag{5.8}$$

Here, RSS is the Residual Sum of Squares, which is minimized in multiple linear regression. However, in contrast to linear regression, Lasso regression does not only minimize RSS. It also penalizes the sum of the β-coefficients by absolute value. Or, more precisely, it uses an l_1 penalty. This has the desired effect of forcing some of the coefficient estimates to be exactly equal to zero when the tuning parameter λ is sufficiently large (Gareth et al., 2017). Hence, lasso regression performs a variable selection.

In order to decide the value of λ, and subsequently what variables to include in the model, *k*-fold cross validation is used. In Section 6.2.2, the average *k*-fold cross validation MSE, $CV_{(k)}$, associated with the coefficients $\hat{\beta}^L_\lambda$, is plotted against corresponding values of λ. As suggested by Gareth et al. (2017), two considerations are made when selecting a value of λ. First, the model with the smallest error is considered. However, if such a model includes a large number of explanatory variables, it can be subject to over-fitting. Therefore, the model with an MSE one standard deviation higher than that of the minimum is also considered. Afterwards, a variable selection is performed. That is, the variables associated with a non-zero $\hat{\beta}^L_\lambda$ are chosen.

Model Selection

It is important to note that lasso regression is only conducted in order to perform variable selection. The $\hat{\beta}^L_\lambda$-values from the regression are *not* carried forward. Instead, points are forecast by the use of linear regression on the variables selected by the lasso regression. Linear regression has an advantage over a lasso regression as it has more easily interpreted coefficients and well-defined variance (Gareth et al., 2017). This makes the calculation of *p*-values possible. Moreover, the use of linear regression makes the results more accessible to readers from different disciplines as it is a widely known procedure. Furthermore, a linear regression model has an advantage compared to for instance logistic regressions as points can be negative, which is a possibility in FPL.

The variable *previous points* is the only variable explicitly related to points earned. Other variables such as total points in the season so far and points from the Bonus Point System are not included. These variables would be highly correlated with variables such as goals. Thus, one rationale for excluding them is not to over-complicate the correlation between explanatory variables in the regression model. Furthermore, since the bonus points are to a substantial degree decided by the same factors as ”regular” points (e.g., goals, assists, clean sheets etc.), their effects are assumed included in the coefficients for these explanatory variables.

5.2.3 Odds

As seen in the literature review in Section 3.3.2, Hvattum and Arntzen (2009) found that Elo ratings appeared to be considerably less accurate than market odds when used for forecasts. Furthermore, in their project report, Kristiansen and Gupta (2017) created a scoreline prediction for the fixtures of the Premier League and used it in order to forecast future performance of the players in FPL. However, they found their scoreline predictions far less accurate than those of bookmakers' odds. Therefore, as a complementary method to the Elo based Modified Average method, a forecast method based on odds is suggested.

In order to obtain predictions of player points, odds must be transferred into expected points. Odds can be separated into *team odds* and *individual player odds*. Team odds include all odds related to a team's performance such as scoreline odds. Individual player odds include odds related to an individual player's performance such as goals scored, assists provided and yellow cards received by the player.

The team odds used is related to scorelines and serve two purposes. First, it is used to calculate the probability of a team keeping a clean sheet. By summing all the scoreline odds and dividing each distinct odds by that sum, the probability of each exact result can be obtained. In Table 5.5, an example is given for the odds and the corresponding probabilities. By summing all the probabilities for the scorelines where Leicester has a clean sheet, it can be found that the probability of Leicester keeping a clean sheet is 0.316. Thus, in order to calculate expected points from clean sheets, the probability is multiplied with points received for a clean sheet, i.e., 4 points for goalkeepers and defenders and 1 point for midfielders. Secondly, team odds is used to obtain a team's expected goals. This is done by multiplying goals scored in each scoreline with the probability of that scoreline occurring. In the example from Table 5.5, Leicester's expected goals is found to be 1.310.

Odds

1-0	7.00	0-0	7.50	0-1	8.00
2-0	11.00	1-1	6.00	0-2	13.00
2-1	9.50	2-2	15.00	1-2	10.50
3-0	23.00	3-3	67.00	0-3	29.00
3-1	21.00	4-4	301.00	1-3	23.00
3-2	31.00			2-3	35.00
4-0	61.00			0-4	81.00
4-1	56.00			1-4	67.00
4-2	81.00			2-4	91.00
4-3	181.00			3-4	201.00
5-0	201.00			0-5	276.00
5-1	181.00			1-5	226.00
5-2	276.00				

Probabilities

1-0	0.1043	0-0	0.0974	0-1	0.0913
2-0	0.0664	1-1	0.1218	0-2	0.0562
2-1	0.0769	2-2	0.0487	1-2	0.0696
3-0	0.0318	3-3	0.0109	0-3	0.0251
3-1	0.0348	4-4	0.0024	1-3	0.0318
3-2	0.0236			2-3	0.0209
4-0	0.0120			0-4	0.0090
4-1	0.0130			1-4	0.0109
4-2	0.0090			2-4	0.0080
4-3	0.0040			3-4	0.0036
5-0	0.0036			0-5	0.0026
5-1	0.0040			1-5	0.0032
5-2	0.0026				

Table 5.5: Result odds for Leicester-Burnley 14.04.18 and the corresponding probabilities.

In an analogous manner as with team odds, individual odds can be transformed into probabilities. However, now the probabilities are calculated as the inverse of the odds. Thus, if Leicester's Riyad Mahrez is assigned an odds of 4 for scoring one or more goals, his probability of scoring is $\frac{1}{4} = 25\%$. Note, however, that by using this method it is assumed that the odds are set in a "fair" way, i.e., the odds market is a perfect market. In reality, bookmakers are profit-seekers. Thus, if bookmakers are to profit, the odds they set must transfer into probabilities lower than the real ones, and the validity of assumption can be weakened. After odds are transferred into probabilities, the probabilities are converted to expected points. For points related to scoring a goal, for example, this is achieved by multiplying expected goals and points obtained for a goal. Points for assists are allocated in a similar manner. Note, that all goals are not assisted. However, it is assumed that bookmakers consider this when setting the odds for assists.

	Odds	Probability
Goal	4	0.25
Assist	2.5	0.40
Yellow Card	8	0.125

Table 5.6: Fictive individual odds for Riyad Mahrez.

Finally, the expected points of a player can be computed. In Table 5.6, a fictive set of individual odds and probabilities are given for the Leicester midfield player, Riyad Mahrez. The reader is reminded that a midfielder earns 5 points for a goal, 3 points for an assist, 1 point for a clean sheet and -1 point for a yellow card. In equation 5.9, it is shown that he is expected to obtain 3.40 points for the Leicester-Burnley match.

$$5 \cdot 1.310 \cdot 0.25 + 3 \cdot 1.310 \cdot 0.4 + 1 \cdot 0.316 - 1 \cdot 0.125 \approx 3.40 \tag{5.9}$$

Note that the probability of starting is not explicitly handled. Nor is points awarded for playing over 60 minutes. However, one can argue that the probability of starting is incorporated in individual odds, i.e., as the probability of scoring a goal or providing an assist.

5.3 Gamechips

Gamechip strategies are actions developed to decide when it is optimal to use the gamechips. There exists limited literature on this topic. Consequently, there is a leeway for new ideas. Many interesting approaches could be adopted to decide the optimal use of gamechips. The approach adopted in this thesis is mainly qualitative and aim to take advantage of future gameweeks which are *blank* or *double*. The reader is reminded a gameweek is blank when at least one team does not play, and a double if a team is featured more than once.

FPL managers must make decisions on when to use their gamechips, i.e., they have to decide whether it is best to play them in a given gameweek or if it is optimal to wait. As the model is solved using a rolling horizon heuristic, a strategy for when to use the gamechips is required. Also, there is nothing that prevents the model from not using the gamechips in the current sub-problem, unless the strategy addresses this issue. Many different approaches can be undertaken to develop such a strategy. One approach is to base the strategy on intuition and well-known reasoning among FPL managers.[1] As limited research is conducted on the use of gamechips, such strategies are considered an appropriate place to start.

5.3.1 First Wildcard

The first Wildcard, which has to be played within the first half of the season, is played at the end of the first quarter of the season. There is high uncertainty at the start of the sea-

[1] See websites https://www.premierleague.com/news/624013 and https://www.premierleague.com/news/404004

son related to starting line-ups and the form of players. In addition, the summer transfer window ends in gameweek 4, providing a possibility of players being sold, lend or bought. Thus, using the Wildcard after a few gameweeks allows managers to wait and consider such developments. However, if managers wait too long, they can miss the opportunity of adding players that over-perform compared to their initial expectations. Hence, playing the Wildcard in last part of the first quarter of the season takes both considerations into account. This is implemented by giving the mathematical model the opportunity to play a Wildcard in a specific gameweek in the first quarter of the season which is chosen beforehand.

5.3.2 Triple Captain

It is assumed sensible to play the Triple Captain in a double gameweek as players are expected to gain more points when they are featured twice. Thus, the model is given the opportunity to play the Triple Captain in the first double gameweek of the season.

5.3.3 Second Wildcard and Bench Boost

The second Wildcard should be played with caution, as it is sensible also to consider the Bench Boost when using it. As the Bench Boost allows each player in the selected squad to collect points, and not only the starting line-up, it is wise to play the Bench Boost in a double gameweek. Therefore, the model is allowed to play the second Wildcard in the gameweek ahead of a double gameweek.

5.3.4 Free Hit

The Free Hit allows a manager to perform an unlimited amount of free transfers for one gameweek. In the gameweek after, the selected squad is changed back. There are in general two cases where playing the Free Hit is reasonable. First, the Free Hit could be played in a blank gameweek to ensure that all the players in the starting line-up are featured for that particular gameweek. Secondly, the Free Hit could be played in a double gameweek to ensure that the majority of the starting line-up consists of players that are featured twice. As two other gamechips already are related to double gameweeks, the first method is implemented.

5.3.5 Gamechip Implementation in the Mathematical Model

The gamechips are implemented in the mathematical model by initially setting the right-hand side of constraints (4.2) - (4.6) equal to 0. The right-hand side of the constraints is only set to 1 in the respective constraints when a gameweek where a gamechip can be played is considered in the sub-horizon. Up to those gameweeks, the constraints are deactivated and have no effect on the objective function.

When a gameweek where a gamechip can be played appears in the sub-horizon, the related constraints are split in two. For example, if the sub-horizon is set to 4, and the first Wildcard is allowed to be played in gameweek 12, it is first considered when the mathematical model is run for sub-problem 9. In the model, the right-hand side of constraints (4.3) - (4.6) are set to 0, while constraints (4.2) is split in two:

$$\sum_{t \in \{9,10,11\}} w_t \leq 0$$

$$w_{12} \leq 1$$

In the next sub-problem, the range is moved so the constraints is now modelled as:

$$\sum_{t \in \{10,11,13\}} w_t \leq 0$$

$$w_{12} \leq 1$$

The other gamechips are implemented in a similar way as the Wildcard.

5.4 Risk Handling

As both goalkeepers and defenders in the same team earn 4 points for keeping a clean sheet, it is not hard to imagine that the points they obtain are positively correlated. Similarly, it is not unreasonable to assume that the points a goalkeeper and forward of opposing teams earn are negatively correlated. In this section, a method for dealing with this assumed correlation in points obtained, as well as the fluctuations of individual player's points, is presented. As described in Section 3.4, there is a large resemblance between the FPLDP and the well-known portfolio selection problem in financial optimization. This is taken advantage of by attempting to handle risk in an analogous manner in the FPLDP as is done in portfolio optimization. Thus, one might say that the objective of adding risk handling is *diversification*, a risk management technique with the purpose of yielding higher returns (expected points) without increasing risk.

5.4.1 Player Variance

Variance refers to the variability in each player's realized points. Some midfielders are for instance known to be consistently starting but seldom contribute in terms of assists or goals. It is reasonable to assume that such players have a low expected value of points, but also a low variability in points obtained. Other players might contribute a lot in terms of assists and goals when they are playing, but are perhaps often injured. Such players could have a high number of expected points, but also a high variance in points obtained.

5.4.2 Correlation

The variance of each player only considers each player's individual contribution to the team's total variance. As discussed, it is natural to assume that there exists a correlation between the points obtained by players. However, even if the points obtained by a goalkeeper and a defender from the same team are expected to be positively correlated, it is perhaps not so obvious whether or not there exists correlation between defenders of opposing teams. In order to test for correlation, Pearson's Product-Moment Correlation test is performed. By applying the test, a few assumptions are made. The variables are assumed approximately normally distributed and display *homoscedasticity* (constant variance). For a full description of Pearson's Product-Moment Correlation test and its underlying assumptions, the reader is referred to Lund and Lund (2017).

5.4.3 Team Variance

In financial optimization, it is common to examine the trade-off between expected return and variance (Markowitz, 1952). According to Zenios (2005), the variance of a portfolio containing *n* assets with weight *w* can be expressed as

$$\sigma^2(w) = \sum_{i=1}^{n}\sum_{i'=1}^{m} \sigma_{ii'} w_i w_{i'} = \sum_{i=1}^{n} \sigma_i^2 w_i^2 + \sum_{i=1}^{n}\sum_{\substack{i'=1;\\ i'\neq i}}^{n} \sigma_{ii'} w_i w_{i'} \tag{5.10}$$

where, w_i is the portfolio weight of asset i, σ_i^2 is the variance of asset i and $\sigma_{ii'}$ is the covariance between assets i and i'.

In the case of the FPLDP, the mathematical model optimizes over the starting line-up variable, y. In Equation 5.10, returns are expressed in percentage, as is common in finance. For the FPLDP, however, it is the variance of the sum of total points obtained each gameweek, rather than the variance in percentage return, that is of interest. Hence, Equation 5.10 can be rewritten by utilizing the fact that each "portfolio weight" is either 1 or 0. If the expression for covariance is also substituted with correlation, the expression simplifies to

$$\sigma^2(y) = \sum_{i=1}^{n} \sigma_i^2 y_i + \sum_{i=1}^{n}\sum_{\substack{j=1;\\ j\neq i}}^{n} \sigma_i^2 \sigma_j^2 y_i y_j \eta_{i,j} \tag{5.11}$$

where σ_i^2 is the variance of expected points obtained by player i, σ_j^2 is the variance of expected points obtained by player j and $\eta_{i,j}$ is the correlation between expected points obtained by player i and player j. In addition, y_i and y_j are binary variables indicating whether player i and player j are included in the starting line-up (portfolio) of players. Conventionally ρ is used to denote correlation. However, in this thesis, η denotes correlation as ρ is already reserved for denoting points.

In the mathematical model, an upper threshold is determined for the variance of the starting line-up at time t as described below. Note that in Equation 5.11, two binary variables are multiplied. Hence, the constraints are non-linear. In order to linearize the expression, an auxiliary variable, $\delta_{pp't}$, is introduced and the following set of constraints are inferred.

$$\sigma_t^2 = \sum_{p=1}^{P} \sigma_p^2 y_{pt} + \sum_{p=1}^{P}\sum_{\substack{p'=1;\\ p'\neq p}}^{P} \sigma_p^2 \sigma_{p'}^2 \eta_{p,p'} \delta_{pp't} \qquad t \in \mathcal{T} \tag{5.12}$$

$$\sigma_t^2 \leq \sigma_0^2 \qquad t \in \mathcal{T} \tag{5.13}$$

$$\delta_{pp't} \leq y_{pt} \qquad p \in \mathcal{P}, p' \in \mathcal{P}, p' \neq p, \; t \in \mathcal{T} \tag{5.14}$$

$$\delta_{pp't} \leq y_{p't} \qquad p \in \mathcal{P}, p' \in \mathcal{P}, p' \neq p, \; t \in \mathcal{T} \tag{5.15}$$

$$\delta_{pp't} \geq y_{pt} + y_{p't} - 1 \qquad p \in \mathcal{P}, p' \in \mathcal{P}, p' \neq p, \; t \in \mathcal{T} \tag{5.16}$$

Note that constraints (5.12) and (5.13) could be combined into one constraints. However, for the sake of legibility, they are both explicitly written. Observe also that term σ_0^2 in constraints (5.13) is the risk threshold, i.e., a parameter set according to the risk preference of an investor in the portfolio optimization setting. In Section 6.4, estimation of the parameters related to risk handling is treated. Furthermore, in Section 7.4 the effect of changing the risk threshold, σ_0^2, is analyzed.

5.5 The Forecast-Based Optimization Model

As seen throughout the chapter, methods for generating input data related to forecasts and risk handling, as well as strategies for the use of gamechips are developed. By taking all the aforementioned elements into consideration, the *forecast-based optimization model*, a solution framework for the Fantasy Premier League Decision Problem, is developed. Algorithm 1 gives a pseudo code for the forecast-based optimization model. Here, GC is the set of gameweeks where gamechips are implemented according to the gamechips strategies.

Algorithm 1 Forecast-based optimization model

```
k = 1
Set risk handling threshold, σ_0^2
while k ≤ |T| do
    Generate input data: expected points, variance and correlation for TP_k^C and TP_k^F
    if a gameweek in TP_k^F is in GC then
        Implement gamechip strategy
    end if
    Solve the mathematical model for the problem defined by TP_k^C and TP_k^F
    Fix variables x, v and q prior to solve TP_{k+1}^C and use as input
    k ← k + 1
end while
```

Chapter 6

Computational Study

In this chapter, a computational study for the forecast-based optimization model presented in Chapter 5 is conducted. The goal is to provide and justify numerical values for parameters used in the model before it is tested on the 2017/2018 season. For the purpose of readability, the results of the testing are presented in the next chapter, namely Chapter 7. From now on, the points obtained over a season is referred to as *performance*. Performance can either be measured in total points obtained over a season or in terms of mean points per gameweek.

First, in Section 6.1, the handling of the data used is discussed. Section 6.2 elaborates on how parameters and input associated with the three different forecasting methods are decided. Section 6.3 provides the decisions regarding which gameweeks to consider the use of gamechips. Finally, in Section 6.4, the numerical values related to the risk handling are presented.

6.1 Data

Data provided by the official website of Fantasy Premier League is used. A detailed record of such data for both the 2016/2017- and the 2017/2018 season is kept by Fantasy Overlord (see www.fantasyoverlord.com). As only data for two seasons are available, data for the first season is used for training and the latter for testing. Two seasons worth of data is considered a limited amount since it is reasonable to assume that if more data were available the model could be improved further and more robust statements could be made. Furthermore, Elo ratings for each Premier League team are obtained from ClubElo (see www.clubelo.com). In addition, Sportradar has provided odds for the fixtures of the 2017/2018 season.

A season in FPL consists of 38 gameweeks. However, because of restrictions on time, an early decision was made to solve the model only for the first 35 gameweeks. The transfer window closed after gameweek 26. Until then, 625 players had appeared in at least one game. Thus, the total number of players considered is 625.

6.1.1 Processing of Player's Data

Injuries and Suspensions

Once a player is listed as injured and thus unavailable for the upcoming gameweek, his expected points, $\hat{\rho}$, are automatically set equal to 0 in all the gameweeks in the sub-horizon. As the injury list is updated for each gameweek, a player's injury status can change when solving for the next sub-problem. Injury data is collected from Fantasy Premier League (see www.fantasypremierleague.com), ahead of each gameweek.

Suspensions are also accounted for by setting the expected points, $\hat{\rho}$, of suspended players equal 0. A player may be suspended for up to three games when receiving a red card, depending on whether it was a direct red card or not. Further, players are subject to a match ban when receiving their 5th, 10th or 15th yellow card of the season. In addition, a player can be suspended by the English Football Association if he is found guilty of unsportsmanlike conduct.

Promoted Teams

There does not exist Fantasy Premier League data for English football's second tier, the English Championship. Therefore, gathering data for newly promoted teams is difficult. In addition, comparing performances in the English Championship with performances in the English Premier League is unreliable. Thus, some simplifications are made. In general, a player is not taken into consideration until data on his previous performance in the FPL becomes available. However, if a player was transferred to a promoted team ahead of the season, and he played in the Premier League in the 2016/2017 season, the player would be included from the start of the season. For the Odds method, data exists for newly promoted teams as well. It is worth noticing that in the Modified Average method, newly promoted teams are considered from the start of the 2017/2018 season with regards to the Elo ratings.

Players Transferred Ahead of 2017/2018 Season

A question that arises with the international transfer window is how newly transferred players are going to perform in the English Premier League. Such players are handled in the same way as newly promoted players, i.e., by not considering them until they have featured in a sufficient amount of gameweeks. Other approaches include forecasting based on their performance in the foreign league or comparing them to players of the same price in FPL. In general, the decision of not including players when data is missing is supported by the fact that an aim is to test how the different approaches perform, and including such players induces more uncertainty.

Irregular Gameweeks

For blank gameweeks, expected points, $\hat{\rho}$, are set to 0 for players that are not featured. For double gameweeks, forecasts are made for both matches. It is assumed that information on whether a gameweek is blank or double is known at least as many gameweeks ahead as the length of the sub-horizon.

6.2 Forecast of Player Points

In this section, a computational study for the three different forecasting methods is presented. For all methods, numerical values for parameters such as the sub-horizon in the rolling horizon heuristic are set, and the decisions are justified. Further, the possibility of altering the penalty term, R, when optimizing is considered. For the Modified Average method, a detailed parameter study is undertaken to find the best combination of numerical values for parameters such as the sub-horizon and track-length. For the Regression method, a detailed variable selection procedure is performed. For the Odds method, the handling of data obtained by Sportradar is presented. Note that all computations conducted in this section are carried without taking the aspect of risk handling into consideration.

6.2.1 Modified Average

The aim of this study is to set numerical values to the factors presented in Equation 5.1 for the Modified Average method. As explained in Section 5.2.1, the relative team strength factors are calculated according to the Elo ratings. These ratings are updated for every gameweek, depending on the results in the previous games. An overview of the Elo ratings is attached in the Appendix B. Further, the starting line-up factor is assumed handled by data regarding suspensions and injuries. Thus, determining the field factor, streak factor variables *X* and *Y* and track-length remain.

Determining Field Factor and Streak Factor Variables *X* and *Y*

Table 6.1 gives the values for field factors for the past six Premier League seasons. As observed, the home-field factor range from 1.105 to 1.149 and the away-field factor range from 0.851 to 0.895. Furthermore, the average of factors over the 5 last seasons before the 2016/2017 and before the 2017/2018 season is presented in the table. The average before the 2016/2017 season (Pre 16/17 Avg) is based on season 2011/2012 to 2015/2016. Similarly, the average before the 2017/2018 season (Pre 17/18 Avg) is based on season 2012/2013 to 2016/2017. Thus, taking the average over the past 5 seasons yields almost identical factors before the 2016/2017 and before 2017/2018 season. Hence, the home-field factor is set to 1.13 and the away-field factor is set to 0.87.

Season	11-12	12-13	13-14	14-15	15-16	16-17	Pre 16-17 Avg	Pre 17-18 Avg
Home-field factor	1.133	1.114	1.137	1.149	1.105	1.141	**1.128**	**1.129**
Away-field factor	0.867	0.886	0.863	0.851	0.895	0.859	**0.872**	**0.871**

Table 6.1: Field factors for previous 6 seasons.

For the point streak factor, the variables X and Y regarding whether a player is on a point streak, have to be decided. Goalkeepers and defenders receive 4 points for keeping a clean sheet, all positions earn 3 points for an assist and midfielders and forwards earn 5 and 4 points for scoring a goal, respectively. Further, as a player earns 2 points for playing more than 60 minutes, it is appropriate to let the X variable take a value of 5. This means that if a player keeps a clean sheet, has an assist or scores a goal and additionally plays at least 60 minutes, he will receive enough points to be on a positive point streak given that he did not receive a yellow or red card. As for the Y variable, a player that does not contribute with either a goal, assist nor a clean sheet is awarded maximum 2 points. Thus, it is reasonable to set the Y equal to 2 points.

Determining Optimal Track-Length and Sub-Horizon

In the following, the aim is to determine an optimal combination of sub-horizon and track length based on data for the 2016/2017 season.

When testing combinations to determine the optimal pair of sub-horizon and track-length, the sub-horizon does not exceed the track-length. The rationale is that the track-length is not accurate when forecasting for more gameweeks than what it is composed of. As only 1 season of training data is available, the model's performance is expected to improve as the track-length increases. For instance, consider the case when the track-length is 7 gameweeks. As the method is trained on the 2016/2017 season, only forecasts for matches after gameweek 7 has been played can be made. Hence, the model gets an opportunity of initially selecting the players that have performed best over the 7 first gameweeks. Thus, the selection is somewhat unfair. Therefore, track-lengths greater than 6 gameweeks are not considered. By similar reasoning, it is more appropriate to measure performance in mean points rather than total points across different track-lengths. Notice that only the training data is affected. When the model is run for the first gameweeks of 2017/2018 season, the last gameweeks of the 2016/2017 season are used when computing tracking average.

Figure 6.1 displays the results when considering track-lengths from 3-6 gameweeks as well as sub-horizons from 1-6 gameweeks. Moreover, Table 6.2 provides the numerical values of the combinations in Figure 6.1 with the 10 highest means as well as average weekly penalized transfers made for these combinations.

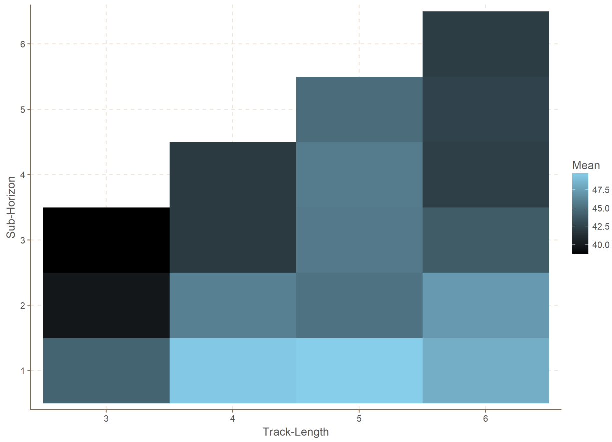

Figure 6.1: Performance with changing sub-horizon and track-length and penalty, R, set to 4.

Sub-horizon	Track-length	Mean	Mean penalized transfers
1	5	49.82	1.48
1	4	49.53	1.88
1	6	48.25	1.34
2	6	47.28	2.09
2	4	45.94	2.74
4	5	45.73	2.88
3	5	45.58	2.70
2	5	45.18	2.55
5	5	44.88	3.00
1	3	44.43	2.63

Table 6.2: The 10 best combinations of track-lengths and sub-horizons.

As observed in Figure 6.1 and Table 6.2, the model reaches its best performances with a sub-horizon of 1 gameweek combined with track-lengths of 4-6 gameweeks. However, as seen in Table 6.2, the performances generally increase as the number of penalized transfers decreases. Consequently, in the following, it is investigated how limiting the number of penalized transfers by adjusting the penalty parameter, R, affects performance. Note, as the penalty term is set by the game rules of FPL, altering the penalty term when optimizing is in reality only compensating for inaccurate forecasts. However, it is done as the end goal is to maximize the performance of the model.

Model Run by Adjusting Penalty Term, R

To set values for the parameters R, track-length and sub-horizon, different combinations are tested on the 2016/2017 season. Table 6.3 provides the numerical values from the 10 best combinations.

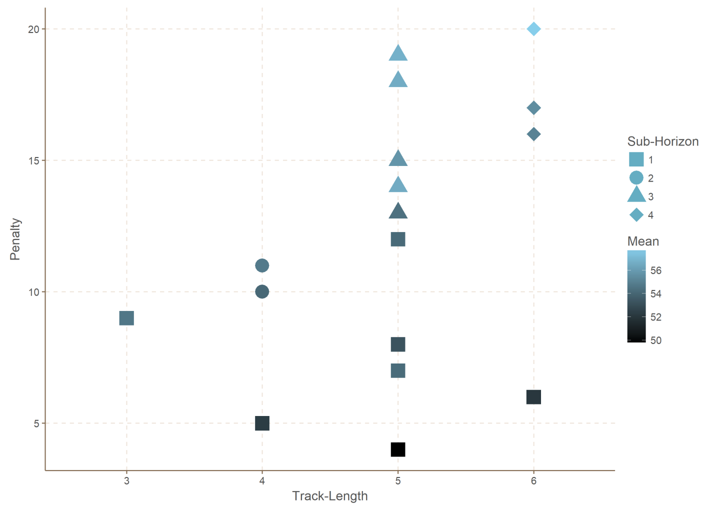

Figure 6.2: Performance with different penalty terms, sub-horizons and track-lengths.

R	Sub-horizon	Track-length	Mean	Mean penalized transfers
20	4	6	57.84	0.0313
20	4	5	57.55	0.0606
19	3	5	56.82	0.0606
18	3	5	56.64	0.0606
14	3	5	56.61	0.273
18	4	6	56.50	0.0938
19	4	6	56.44	0.0625
20	3	5	56.36	0.0303
20	5	5	56.09	0.212
18	6	6	56.03	0.313

Table 6.3: The 10 best combinations of parameters.

Figure 6.2 provides an overview of how the mean points each gameweek for the entire 2016/2017 season changes with different track-lengths, sub-horizons and penalty terms. In the figure, only the best run (maximum mean) for each different penalty value from 3-20 is plotted. According to Figure 6.2 and Table 6.3, sub-horizons of 3-4 gameweeks appear to be ideal, as they provide the combinations with the highest mean. Furthermore, penalties in the range of 14-20 yield the best results. As for the track-lengths, values of 5 and 6 gameweeks are clearly optimal. However, the latter is not unexpected due to the unfair player selection. From the table, it is clear that mean penalized transfers are significantly reduced for the best performing cases, compared to when R was 4.

Figure 6.3 and Table 6.4 display the best combinations when the track-length is held fixed for 3-6 gameweeks. Clearly, track-lengths of 5 and 6 gameweeks outperform those of 3 and 4 gameweeks based on mean points. Moreover, Table 6.3 shows that a track-length of 6 gameweeks combined with a sub-horizon of 4 gameweeks and a penalty term of 20 yields the highest mean for the 2016/2017 season. However, the next four highest means are achieved with a track-length of 5 gameweeks. Due to the relatively similar performances of a track-length of 5-6 gameweeks, but considerable differences in performance when the track-length is 4-5 gameweeks, a track-length of 5 gameweeks is selected. Additionally, according to Figure 6.3, a sub-horizon of 3 or 4 gameweeks is optimal. Combined with a track-length of 5 gameweeks, a sub-horizon of 3 gameweeks generally performs better than that of 4 gameweeks. Therefore, the sub-horizon is set to 3 gameweeks. Finally, Figure 6.3 illustrates that relatively high penalty terms yield the best means for the 2016/2017 season, with penalty terms ranging from 14-20. However, from Table 6.3, it is apparent that in combination with a track-length of 5 gameweeks and sub-horizon of 3 gameweeks, the best penalty term is in the range from 14-19. Also, from Table 6.4, it is clear that the penalty term is considerably lower for the best combinations with a track-length of 3 and 4 gameweeks and that mean penalized transfers are considerably higher for these track-lengths. By taking all this into consideration, a penalty term, R, of 16 is selected.

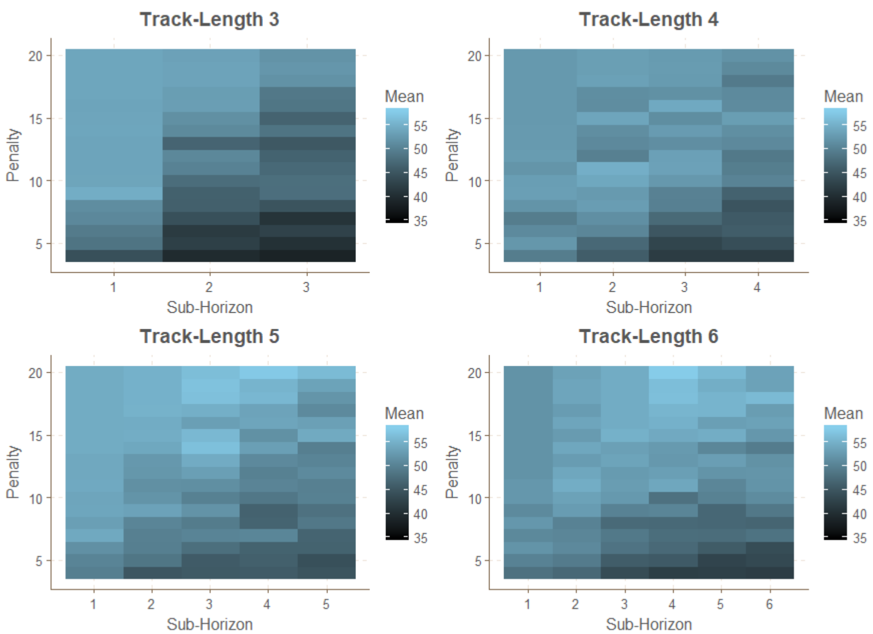

Figure 6.3: Performance with different penalty terms, sub-horizons, but fixed track-lengths.

R	Sub-horizon	Track-length	Mean	Mean penalized transfers
9	1	3	54.69	0.457
11	2	4	54.79	0.676
20	4	5	57.55	0.0606
20	4	6	57.84	0.0313

Table 6.4: The best combinations for each track-length.

In summary, the Modified Average approach on the 2017/2018 season is run using the parameters shown in Table 6.5.

Parameters	Value
Sub-horizon	3 gameweeks
Track-length	5 gameweeks
R	16

Table 6.5: Numerical values of parameters when Modified Average is run on 2017/2018 season.

6.2.2 Regression

In the following, a computational study for the Regression model based on the approaches outlined in Section 5.2.2 is conducted. When making decisions, both analytic and more rationally influenced arguments are considered. For instance, even if the points obtained by a forward might be correlated with the number of clean sheets his teams have collected, the variable could still be excluded. The intuition is that with training data limited for one season only, one cannot blindly trust the models.

Variables

In the first gameweek, where it is considered reasonable, data are set equal to the last gameweeks of the 2016/2017 season. This is for instance the case for *goals* and *assists*. For *transfers in* and *transfers out*, however, this is not a sensible approach, as managers do not base their transfer-decisions in the last gameweek of the season with the upcoming season in mind. Therefore, these values are set equal to 0 in the first gameweek. The variables *saves* and *penalty saves* are only defined for goalkeepers and are thus only considered for goalkeepers.

Categorization Based on Position

As outlined in Section 5.2.2, the players are separated based on their *positions*. To substantiate this decision, a visual assessment is undertaken. The analysis is done for the 2016/2017 season. First, Figure 6.4 displays a box-plot of the total points for each different position. The different positions appear to exhibit some differing characteristics. The range is larger for midfielders and forwards than what it is for defenders and goalkeepers.

Further, defenders and goalkeepers have the highest median scores and the most positive outliers are found among midfielders and forwards.

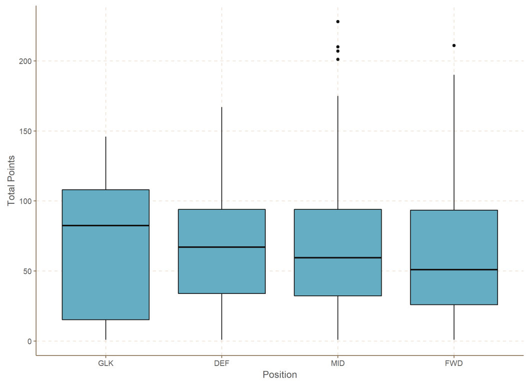

Figure 6.4: Box plot substantiating the decision of separating the players based on position.

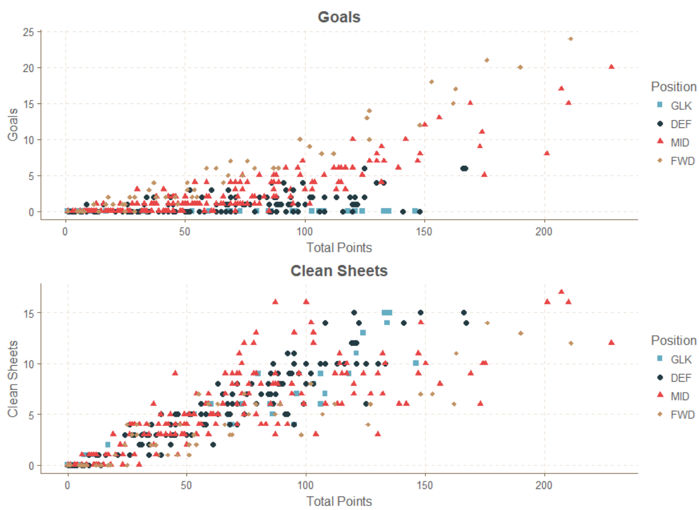

Figure 6.5: Scatter plots substantiating the decision of separating the players based on position.

Secondly, Figure 6.5 shows how total points are associated with the variables *clean sheets*, and goals. It is clear that goals are not contributing a lot in terms of total points for goalkeepers, but are important for midfielders and forwards. Figure 6.5 also manifest that clean sheets are important for goalkeepers and defenders in order to earn a high point score. Based on these observations, categorizing players by position seems reasonable.

Time Series of Points

To determine if the *previous points* a player has obtained should be treated as a time series (i.e., include the effect of form), the Durbin-Watson test and the Ljung-Box test with lags from 1-5 for autocorrelation are performed. Table 6.6 presents the results of the tests. Based on the results, 75-95% of the players exhibit insignificant autocorrelation in their time-series of points at a significance level of 10 %, and 80-99% at significance level 5%. The results indicate no autocorrelation effects. Therefore, the variable previous points is disregarded.

Test	Significance level, 10%	Significance level, 5%
Durbin-Watson	94 %	99 %
Ljung-Box lag 1	76 %	84 %
Ljung-Box lag 2	76 %	81 %
Ljung-Box lag 3	78 %	84 %
Ljung-Box lag 4	80 %	84 %
Ljung-Box lag 5	81 %	85 %

Table 6.6: Percentage of players with insignificant autocorrelation for different tests and significance levels.

Variable Selection

As described in Section 5.2.2, lasso regression with k-fold cross validation is performed in order to conduct a variable selection. Data from the entire 2016/2017 season is used and separated into $k = 10$ folds for the cross validation. In Figure 6.6, a plot of average k-fold MSE, $CV_{(k)}$, against values of log(λ) for each position is presented. The reader is reminded that λ is the tuning-parameter found in Equation 5.8. In each plot, there are two dotted vertical lines. The leftmost line marks the log(λ)-value with lowest MSE. The rightmost marks the log(λ)-value with one standard deviation higher MSE than the lowest MSE. Furthermore, the numbers in the top of the plot indicate the number of variables included for each value of log(λ). The reason why the number of variables may exceed the number of variables presented in Section 5.2.2 is that a dummy-variable is introduced for each level of a categorical variable. Thus, the variables *team* and *opponent* will be treated as 20 different variables each as there are 20 teams (levels) in the English Premier League. If only 1-3 of the dummy-variables are selected by the lasso-regression, the variable associated with them is excluded from consideration. Similarly, if more than 3 dummy-variables are selected, all levels are included. That is, some levels are included even if they are not selected by the lasso regression.

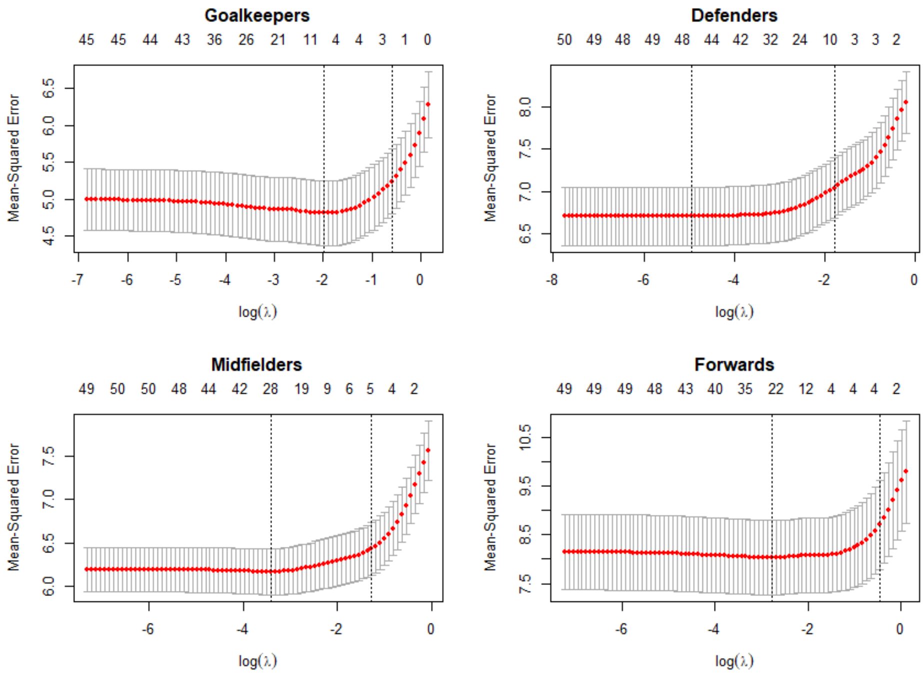

Figure 6.6: k-fold cross validated MSE for different values of log(λ) for the different positions.

Model Selection

For each position, a linear regression model is fitted based on the variables selected. In the following, a presentation of the selected variables and a summary of each model is presented. The β-coefficients presented are computed before the 2017/2018 season. Note that even though the variables are the same throughout the season, the β-coefficients are subject to change, as each model is refitted every gameweek when new data becomes available. Even if team or opponent are selected, they are not presented in the summary, as they would require 20 variables each.

Goalkeepers

No goalkeepers scored a goal in the 2016/2017 English Premier League season, and only one assist was provided by a goalkeeper. Therefore, it is decided not to consider these explanatory variables when performing the variable selection for goalkeepers. Figure 6.6 shows that for the lowest MSE, 6 variables are included in the model. With only 6 variables, the possibility of overfitting is considered limited. Therefore, all the variables associated with this log(λ)-value are considered. However, one of them were related to a team, and another to an opponent. As only 1 of 20 variables for team and opponent were included, both team and opponent are excluded from consideration. In Table 6.7, a summary of a regression fitted on the selected variables is presented. All variables have a sign consistent with their intuitive interpretation. The p-values indicate that all variables are significant.

Variable	Coefficient (β)	Std. Error	t-value	p-value
(Intercept)	-3.35E+00	7.1E-01	-4.7	2.4E-06
Price	8.33E-07	1.6E-07	5.3	1.4E-07
Transfers in	4.50E-05	7.0E-06	6.4	2.3E-10
Minutes	1.12E-02	2.0E-03	5.5	4.0E-08
Saves	1.76E-02	2.8E-03	6.4	2.4E-10

Table 6.7: Summary of the linear regression for goalkeepers.

Variable	Value
Price	$5.5 \cdot 10^6$
Transfers in	13 884
Minutes	90
Saves	20

Table 6.8: Numerical values for Petr Cech in gameweek 9.

Table 6.8 shows the price, transfers in, minutes and saves of Petr Cech ahead of gameweek 9 in the 2017/2018 season. His expected points can be computed as:

$$-3.35+8.33\cdot 10^{-7}\cdot 5.5\cdot 10^{6}+4.5\cdot 10^{-5}\cdot 13884+1.12\cdot 10^{-2}\cdot 90+1.76\cdot 10^{-2}\cdot 20\approx 3.2$$

Note that the example does not capture the fact that the β-values are subject to change due to refitting of the model when new information becomes available.

Defenders

From Figure 6.6 it is apparent that the model with the lowest MSE value includes approximately 47 variables, while the model one standard deviation higher (indicated by the rightmost dotted line), includes 10 variables for defenders. By taking overfitting into consideration, the model corresponding to the log(λ)-value with MSE one standard deviation higher than the minimum MSE is chosen.

None of the variables associated with team were selected, but 4 variables associated with opponent were included. Therefore, team is excluded from the linear regression, but opponent is not. In Table 6.9, a summary of the selected variables can be found. Note that the coefficient corresponding to the variable *home/away* is multiplied by 1 if a player plays at home, and 0 otherwise. Thus, the signs in front of all coefficients except the variable *yellow cards* make intuitive sense. A justification for the positive sign in front of yellow cards is perhaps that players receiving many yellow cards are also frequently involved in matches, thus earning points from other activities. It is worth noting that yellow cards is the variable in Table 6.9 with the lowest p-value, even if they are all significant.

Variable	Coefficient (β)	Std. Error	t-value	p-value
(Intercept)	-3.5E+00	3.9E-01	-9.00	3.4E-19
Price	7.9E-07	7.4E-08	10.70	2.4E-26
Transfers in	2.7E-05	3.2E-06	8.54	1.9E-17
Transfers out	-1.9E-05	3.1E-06	-6.10	1.1E-09
Home/Away	3.9E-01	8.5E-02	4.54	5.8E-06
Minutes	1.2E-02	1.1E-03	11.20	1.3E-28
Yellow Cards	7.9E-02	2.0E-02	3.88	1.1E-04

Table 6.9: Summary of the linear regression for defenders.

Midfielders

From Figure 6.6, it is apparent that the model with the lowest MSE value includes 28 variables for midfielders, while the model one standard deviation away includes 5. Intuitively, it is assumed that the model with best predictive power includes more than 5 variables. Thus, the model with 28 variables is chosen. A substantial amount of the levels corresponding to both team and opponents are included, so both variables are carried forward. The coefficients from the linear regression are presented in Table 6.10. All signs in front of the coefficients are consistent with their intuitive interpretation. Furthermore, they are all significant at a significance level of 5%.

Variable	Coefficient (β)	Std. Error	t-value	p-value
(Intercept)	-2.05E+00	3.0E-01	-6.94	4.3E-12
Price	4.76E-07	3.6E-08	13.10	2.4E-38
Transfers In	9.85E-06	1.4E-06	7.18	8.0E-13
Transfers Out	-8.03E-06	1.5E-06	-5.47	4.8E-08
Home/Away	2.82E-01	6.7E-02	4.19	2.9E-05
Minutes	1.01E-02	9.5E-04	10.70	2.9E-26
Goals	7.57E-02	2.2E-02	3.51	4.5E-04
Penalties Missed	-3.98E-01	1.7E-01	-2.31	2.1E-02
Clean Sheets	6.78E-02	1.6E-02	4.21	2.6E-05
Assists	5.23E-02	2.0E-02	2.61	9.1E-03

Table 6.10: Summary of the linear regression for midfielders.

Forwards

From Figure 6.5, it might appear as if there is a correlation between total points and clean sheets for forwards. However, forwards are not rewarded points for clean sheets. Therefore, the variable clean sheet is not taken into consideration for forwards. Figure 6.6 shows that for forwards, the model with the lowest MSE value includes 25 variables, while the model one standard deviation away includes 4. Again, it is assumed that the model with best predictive power includes more than 4 variables. Hence, the model with minimum MSE is selected. Enough levels of the variables team and opponent are included to make it justifiable to select them. The coefficients from the linear regression are presented in Table 6.11. All the signs in front of the coefficients appear consistent with their intuitive interpretation, except for the variable *penalties missed*. This can perhaps stem from that fact that if a player misses a penalty, he might also frequently take the penalties for his team, thus earning points when he scores. Note, however, that both home/away and penalties missed are insignificant in the linear regression model for a significance level of 5%. As the lasso regression initially selected them, they are included. However, from a strictly theoretical perspective, it would perhaps have been more appropriate to exclude them.

Variable	Coefficient (β)	Std. Error	t-value	p-value
(Intercept)	-1.75E+00	6.7E-01	-2.62	9.0E-03
Price	3.89E-07	6.7E-08	5.79	8.4E-09
Transfers In	1.39E-05	1.9E-06	7.24	6.8E-13
Transfers Out	-5.40E-06	1.8E-06	-3.03	2.5E-03
Home/Away	1.52E-01	1.4E-01	1.10	2.7E-01
Minutes	9.02E-03	2.3E-03	3.99	6.9E-05
Goals	6.54E-02	2.6E-02	2.55	1.1E-02
Penalties Missed	5.52E-01	3.0E-01	1.87	6.2E-02

Table 6.11: Summary of the linear regression for forwards.

Remarks

An interesting observation from the variable selection is that team and clean sheets are not selected for goalkeepers or defenders. Forcing the variable into the model could be an interesting approach in order to investigate what difference it makes, but is not considered in this thesis. Further, it is interesting to note that the effect of form appears to be limited. By not considering the most previous events only, the Regression method works as a complement to the Modified Average method, where only the points in the track-length are considered. However, using totals, for instance for saves, goals and assists, constitute a drawback in the sense that observations are quite similar for all players in the first gameweeks. One way to account for this is to include values from last season in more than the first gameweek. This approach would, on the other hand, complicate the process of including players with a lack of data even more. Another approach would be to only fit the model on data from the first gameweeks of previous seasons. However, that would be more sensible if data for multiple seasons were used in the training set.

Since only data for one season is used to train the model, it is considered infeasible also to perform a parameter study. Therefore, the penalty term, R, is set to 4, as stated by the game-rules. The sub-horizon is set to equal to the sub-horizon in the Modified Average approach, namely to 3.

6.2.3 Odds

Necessary data related to odds have been obtained in cooperation with Sportradar. Sportradar have provided team odds including result outcomes as well as individual odds for goals and assists. Odds are provided for each Premier League match from gameweek 1 until gameweek 31 in the 2017/2018 season. Sportradar exported CSV files containing all the necessary data for each match marked with a match-id. In order to link the match-id to actual matches, Sportradar's application programming interface (API) was utilized. Then, the data was processed and converted into forecasts in the statistical program, R.

A few noteworthy comments must be made about the data used. First, compared to the other two solution approaches, the player list is a bit different for the Odds method, as it is based on Sportradar's database of English Premier League players when the data was received. Thus, a minor drawback is that players that have been transferred or lent to other clubs during the 2017/2018 season are not considered. The reduced player list is not expected to affect the results of the model by much, as only a handful of players such as Philippe Coutinho and Michy Batshuayi contributed with a sizable amount of fantasy points before transferring away.

Secondly, goalkeepers are not assigned with an odds of scoring goals or providing assists. Thus, the only manner in which points are forecast for goalkeepers is from expected points earned by keeping clean sheets. Thus, there is no way to distinguish the expected points of a team's first, second and third choice as goalkeeper. It is decided to circumvent the issue by only awarding points to goalkeepers who played in at in at least 3 previous gameweeks. For the first 3 gameweeks of the 2017/2018 season, the last gameweeks of 2016/2017 season is used. For the other players, as described in Section 5.2.3, it is assumed that their probability of starting is incorporated in their probabilities of scoring a goal and providing an assist.

Thirdly, a detailed record of the odds for bookings was not kept by Sportradar. Thus, it is decided to neglect the effect of bookings. As yellow cards only account for a deduction of 1 point in the FPL, it is not considered a severe drawback.

Finally, there exists a more problematic limitation related to the data in the Odds method. Bookmakers usually only provide odds for the upcoming gameweek. Hence, the odds can only be used in order to predict the fantasy points one gameweek ahead. That is, the sub-horizon must be set to 1. Furthermore, odds-data for the 2016/2017 season was not obtainable. Therefore, a parameter tuning concerning the penalty term, R, could not be conducted. Hence, R is set to 4 as stated in the game rules.

6.3 Gamechips

The gamechips are played according to the strategies suggested in Section 5.3. By examining the fixture list of this year's Premier League season, the following decisions are made. The model is given the opportunity of playing the first Wildcard in gameweek 9. It can either choose to play it that particular gameweek or disregard it. Further, the Triple Captain can be played in gameweek 22, as this is the first double gameweek of the season. Playing the Bench Boost is allowed in the second double gameweek of the season, gameweek 34. Thus, the model is given the opportunity of playing the second Wildcard in gameweek 33, in order to prepare for the Bench Boost in the following gameweek. Finally, the model can play the Free Hit in gameweek 31, which is a blank gameweek were only 8 teams are featured.

Notice that the decisions for which gameweeks to play the gamechips are not fixed ahead of the season. Instead, they are incorporated as the fixture list is updated with regards to blank and double gameweeks. Hence, the gamechips are first considered when a double or blank gameweek arise in the sub-horizon. As stated in Section 6.1, it is assumed that whether a gameweek is blank or double is known as many gameweeks in advance as the length of the sub-horizon, i.e., at most 3 gameweeks ahead. If this was not the case, the assumption could give the model an unfair advantage over FPL managers. However, this was never the case in the 2017/2018 season.

6.4 Risk Handling

To estimate the variance of each player and the correlation between players, their historical performances are considered. To obtain reliable estimations, data for the entire 2016/2017 season is used. In addition, data from the 2017/2018 season is used when it becomes available. Therefore, it is infeasible also to find a suitable variance threshold using data on the 2016/2017 season. As a consequence, it is decided to do an ex-post analysis of how the risk handling constraints affect a solution obtained in the 2017/2018 season, thus providing a basis for future research.

6.4.1 Player Variance Estimation

Each player's variance is calculated as the empirical variance in points obtained. All available data from the 2016/2017 and the 2017/2018 season are used. That is, for gameweek 10 of the 2017/2018, data from all rounds of the 2016/2017 as well as the first nine gameweeks of the 2017/2018 season is utilized. This method of estimation is sensible for the Modified Average method, as it does not aim to explain the variance. However, for the Regression method, the explanatory variables are assumed to explain a part of the variance. Therefore, for the Regression method, it would be more suitable to estimate the variance only as the variance of the residuals. Thus, in the study presented in Chapter 7, only the Modified Average method is considered.

Special cases

Lack of historical data
As previously discussed, complete historical data will not be available for all players for reasons such as promotion or transfers. If only one data point is available for a player (i.e., he has only played one match), it is impossible to calculate his variance. In these cases, the variance is set equal to the average variance of all other players. It is worth noting that this will only be an issue for one match, as more than one data point will exist afterwards. The lack of data points is a limitation of the accuracy of the variance estimation and constitutes a significant source of uncertainty. As risk handling is a method undertaken to reduce overall risk, another approach might consider only picking players with a history of for example 10 gameweeks.

Zero Variance
Some players have empirical variance equal to 0. For instance, players that have never played can still be selected. Thus, they may have a low number of expected points and a low price. These players might be favorable to select to fulfill the selected squad constraints. However, it is not desirable to set their variance to 0, as their future performance is not deterministic. That is, they are *not* analogous to the risk-free asset in a portfolio optimization setting. Therefore, their variance is set equal to the lowest non-zero variance obtained by other players.

6.4.2 Correlation Estimation

To determine the correlation between two players on the same or opposing teams, historical data for the 2016/2017 season is considered. Note that data is aggregated such that only the correlation coefficient between different positions overall is considered. The correlation is not calculated at the level of individual players or teams. Therefore, the correlation between players in the same position on the same team is set equal to 1. The correlation between players that are not on the same team, nor facing each other in a gameweek, is assumed to be 0.

Table 6.12 and 6.13 show the correlation coefficient between different positions. In the tables, goalkeepers, defenders, midfielders and forwards are abbreviated GLK, DEF, MID and FWD, respectively. The tables also include the p-value from a Pearson's Product-Moment Correlation test where the alternative hypothesis is a correlation not equal to 0, for players of the same and opposing teams, respectively. In the cases where the correlation is not significant at a significance level of 5% (GLK FWD and DEF FWD of the same teams and GLK GLK of opposing teams), the correlation coefficient is set equal to 0. Otherwise, the correlation coefficients, η, presented in the tables are used.

Position	Position	η	p-value
GLK	DEF	0.689	2.20 E-16
GLK	MID	0.274	1.24 E-10
GLK	FWD	0.0288	5.19 E-01
DEF	MID	0.368	2.20 E-16
DEF	FWD	0.0578	1.76 E-01
MID	FWD	0.238	1.72 E-08

Table 6.12: Correlation coefficient η and p-value from significance test for the correlation between players of the same team.

Position	Position	η	p-value
GLK	GLK	0.0162	7.59E-01
GLK	DEF	- 0.106	3.74E-02
GLK	MID	- 0.312	2.78E-10
GLK	FWD	- 0.336	3.85E-16
DEF	DEF	- 0.319	2.35E-11
DEF	MID	- 0.447	2.20E-16
DEF	FWD	- 0.292	3.35E-09
MID	MID	- 0.254	1.00E-07
MID	FWD	- 0.136	6.32E-03
FWD	FWD	- 0.119	2.18E-02

Table 6.13: Correlation coefficient η and p-value from significance test for the correlation between players of opposing teams.

Chapter 7

Results

In this chapter, the results of applying the forecast-based optimization model outlined in Chapter 5 in the 2017/2018 season of the FPL, is presented. The model is implemented with the numerical values for parameters computed in Chapter 6. Additionally, the results from solving the mathematical model outlined in Chapter 4 with realized points is presented. First, in Section 7.1, the parameters in the mathematical model are initialized. Then, in Section 7.2, the optimal ex-post decisions are presented. Furthermore, in Section 7.3, the performance of the forecast-based optimization model for the FPLDP with the three different forecasting methods is presented. Also, results from running the model with forecasts generated by the method suggested by Bonomo et al. (2014) and the Modified Average method are compared. The model is first solved without the implementation of gamechips. Subsequently, the model is solved when gamechips are included, and a discussion of the effect of the gamechips follows. Finally, in Section 7.4, the risk handling constraints are included and the results are discussed.

The mathematical model is written in the modeling language Mosel and implemented in FICO® Xpress Optimization Suite 8.3, using a HP EliteDesk 800 G3 DM 65W computer with Intel® Core™ i7-7700 3.6 GHz processor and 32 GB RAM. The operating system in use is Windows 10 Education 64-bit. The input data is structured and pre-processed in the statistical programming language R. Further, the results obtained from the mathematical model include decisions on the selected squad, starting line-up, substitution priority, number of penalized transfers, captain, vice-captain and whether a gamechip is used in a gameweek. These are written to CSV files, which are exported back to R to calculate how many points the team obtained. In the end, R is also used to make the data more presentable in the form of tables and plots.

As mentioned in Section 6.1, the model is solved for the first 35 gameweeks of the 2017/2018 season. In general, it is important to note that for the forecast-based optimization model, no overly categorical statements should be made due to the limited available data. However, insightful observations can be derived.

7.1 Initialization of Parameters

Table 7.1 and 7.2 display the initialization of parameters and sets used when the mathematical model is solved for both realized points and expected points. Notice that the realized points, ρ, is substituted with $\hat{\rho}$ when the model is run with expected points. Moreover, the penalty term, R, is set according to the values decided in Section 6.2.

Set		
$\mathcal{T}$	-	35 gameweeks.
$\mathcal{P}$	-	625 players.
$\mathcal{C}$	-	20 teams.
$\mathcal{L}$	-	$\{1, 2, 3\}$, where 1 is first priority.
$\mathcal{T}_{FH}$	-	Gameweek 1 to 21 is in the first half of the season.
$\mathcal{T}_{SH}$	-	Gameweek 22 to 35 is in the second half of the season.

Table 7.1: Initialization of sets.

Parameters		
ρ_{pt}	-	Realized points for a player p in a gameweek t.
ϵ	-	Set to 0.1.
$\kappa_1, \kappa_2, \kappa_3$	-	Set to 0.01, 0.001 and 0.0001, respectively.
C^S_{pt}, C^B_{pt}	-	Value collected from FPL's homepage.
R	-	4 points deducted for each penalized transfer.
M^K	-	2 goalkeepers required in the selected squad.
M^D	-	5 defenders required in the selected squad.
M^M	-	5 midfielders required in the selected squad.
M^F	-	3 forwards required in the selected squad.
M^C	-	3 players allowed to have from the same team.
E	-	11 players required in the starting line-up.
E^K	-	1 goalkeepers required in the starting line-up.
E^D	-	3 defenders required in the starting line-up.
E^M	-	3 midfielders required in the starting line-up.
E^F	-	1 forward required in the starting line-up.
B^S	-	£100m as starting budget.
β	-	Set to 1.
$\bar{\alpha}$	-	Set to 14.
ϕ	-	3 players are substitutes.
ϕ^K	-	1 goalkeeper among the substitutes.
$\overline{Q}$	-	2 free transfers possible to accumulate over gameweeks.
$\underline{Q}$	-	1 free transfer given every gameweek.

Table 7.2: Initialization of parameters.

All the parameters specifically stated in the rules of FPL are set accordingly. These include the number of players in the different positions in both selected squad and starting line-up, maximum players from the same team, the starting budget and restrictions on free transfers. To tighten the formulation as much as possible, the parameters β and α are set so to the smallest, but sufficiently high, value. β is used in constraints (4.30) and (4.31) in Chapter 4 which concern the substitution priority. Considering that the variable adopted is binary, setting the value to 1 is reasonable. α is used in constraints (4.42), and is set to the highest possible value for the number of penalized transfers. As 1 free transfer is given each gameweek and the selected squad consists of 15 players, it is not possible to have more than 14 penalized transfers. Hence, the parameter α is set accordingly.

7.2 Solution With Realized Points

Before the performance of the forecast-based optimization model for the FPLDP is examined, it is interesting to study the optimal solution when running the mathematical model with ex-post data. For one, the solution serves as a benchmark of the performance of the forecast-based optimization model for the FPLDP. Further, the optimal solution is compared to that of the top-performing manager to see how close the best manager is to the optimal solution. This section is divided into three parts. First, a brief discussion of the problem size is given. Then, a discussion of the results obtained by the mathematical model using realized points for the 2017/2018 season is presented. Finally, the solution is more thoroughly examined and the sensibility of the decisions are briefly discussed. Additionally, the solution is compared to the performance of the top manager.

As mentioned in Chapter 4.2, a simplification is made by assuming that a player's sell price equals his value. However, the assumption is not expected to have a substantial impact on the optimal solution. For one, even if the model does not account for the concept of a sell-on fee, variations in a player's value are incorporated. Hence, when the model sells a player from the selected squad, it only obtains a slightly higher price than what it would do in reality. Furthermore, most players do not fluctuate with more than £0.2m during a season. Actually, only 3 players had a price change exceeding £1m during the 2017/2018 season. Also, among the players with the 10 largest price increases, the average change was £0.63m, earning a profit of £0.3m.

7.2.1 Problem Size

Table 7.3 displays the problem size of the FPLDP with input data from the 2017/2018 season. The problem size is given before and after the function Presolve is run. Presolve is a function integrated in Xpress Optimizer and its objectives are to reduce redundant variables, eliminate redundant constraints and remove linearly dependent constraints, which consequently reduces the complexity of the problem and improve the computing time. From the table, it can be observed that the Presolve function is successful in eliminating variables, but the reduction of constraints is not that severe. A reason for this could be that most of the constraints are linearly independent, and that there is a limited presence of empty rows in the FPLDP. In addition, a large part of the variables x_{pt}, $x_{pt}^{freehit}$ and y_{pt}

are redundant, as many of them are defined even though only a few of them are used in the starting line-up and selected squad constraints. Thus, to reduce the solution time, an effort could have been put into reducing redundant variables in the procedure of implementation. However, since the model only has to be solved once, this is disregarded.

	Rows(Constraints)	Columns(Variables)
Original Problem Statistics	172 034	254 425
Xpress Presolve Statistics	170 750	206 689

Table 7.3: Problem size of the model run with realized points.

7.2.2 Results of Running the Model With Realized Points

During testing, it was observable that nearly optimal solutions were found after 1000 seconds. Hence, the maximum running time was set to 86 400 seconds, or 1 day. This is illustrated in Figure 7.1.

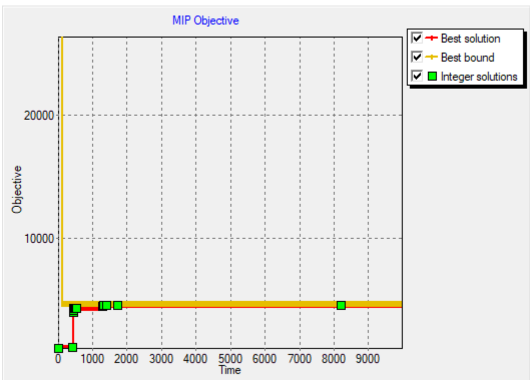

Figure 7.1: Graph of when the solutions are found in Mosel Xpress-MP.

To test the model's scalability, the model was run for different numbers of gameweeks. The results are summarized in Table 7.4, and the computational time is illustrated in Figure 7.2. It is notable that the solution time increases significantly when solved for more than 20 gameweeks. In Table 7.4, it is observable that the optimal solution for 30 and 35 gameweeks was not found within maximum running time. However, the solution found for 35 gameweeks only had a gap of 0.75 %, which was considered sufficient to serve as a benchmark against the top manager. Note that in Table 7.4, the objective value is not directly transferable to total points obtained in FPL. This is due to parameters in objective function such as the ϵ used to select vice-captain.

Number of gameweeks	5	10	15	20	25	30	35
Objective value	773.44	1380.95	2002.69	2660.42	3327.12	3956.76	4553.03
Gap	-	-	-	-	-	0.27 %	0.75 %
Time to optimality	29 s	295 s	1 103 s	2 010 s	23 692 s	N/A	N/A

Table 7.4: Model run with different number of gameweeks.

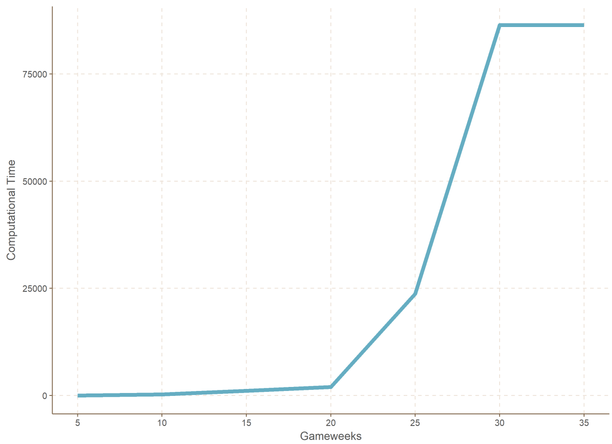

Figure 7.2: Solution time when model is solved for different number of gameweeks.

7.2.3 Performance in Fantasy Premier League

In the following, the optimal solution for 35 gameweeks is examined and discussed. Figure 7.3 gives an overview of how many points the model obtained in each gameweek. The colored dots represent gameweeks where the different gamechips were used.

As seen in Figure 7.3, the first and second Wildcards were used in gameweek 3 and 22, respectively. Thus, the model played the second Wildcard at the earliest stage possible in the second half of the season. Further, the Triple Captain was played in gameweek 31, which is reasonable as Liverpool's Mohammed Salah scored 4 goals and had 1 assist in this particular gameweek, yielding a score of 29 points. This was the maximum number of points obtained in a gameweek by any Premier League player. The Free Hit was played in gameweek 26, a "normal" gameweek, while the Bench Boost was played in gameweek 34, a double gameweek. Thus, the only gamechip that was used consistent with the implementation in the forecast-based optimization model was the Bench Boost.

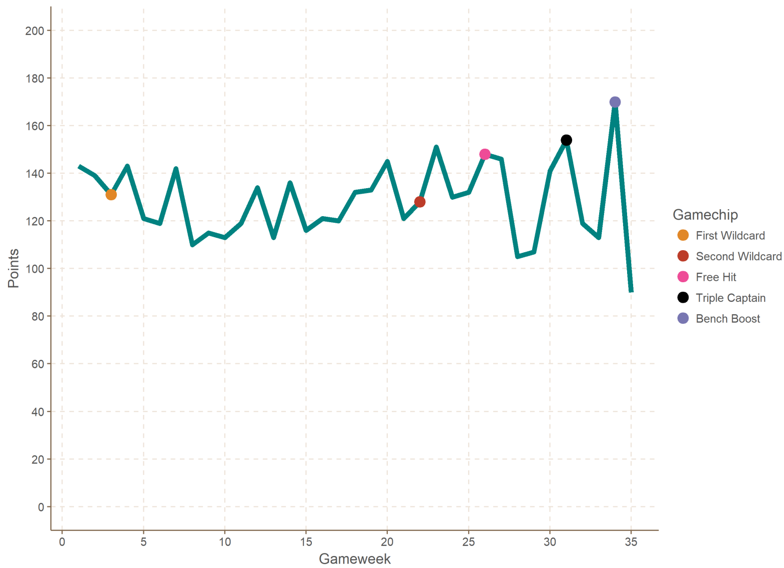

Figure 7.3: Points per gameweek with realized points.

In Figure 7.4, the optimal solution is compared gameweek for gameweek against the weekly average points obtained among all managers and the points obtained by the top manager. It is also indicated when the top manager played the gamechips. The top manager never played a particular gamechip in the same gameweek as the optimal solution. The first Wildcard was played in gameweek 15 and the second Wildcard was played in gameweek 32. The Triple Captain was used in gameweek 22, a double gameweek. Moreover, the Free Hit was played in gameweek 35, a blank gameweek. Thus, the usage of the Triple Captain and Free Hit is consistent with the reasoning made for the forecast-based optimization model. Note that the use of the Bench Boost is not indicated, as the top manager had not played it until gameweek 35.

From Figure 7.4, it is clear that the optimal solution substantially out-performs both the weekly average among managers and the top manager each gameweek. In Table 7.5, the performance of the top manager, as well as the performance required to reach different rankings among all managers, are compared against the optimal solution. A complete overview of the rankings after gameweek 35 was played can be found in Appendix C. It is clear that even the top manager only managed to achieve approximately 50% of the points obtained by the optimal solution. One implication is that it is immensely difficult to compete optimally in FPL, given the nature of the uncertainty in the problem. Moreover, the result can motivate the use of optimization-based tools, as a potential for improvement indeed is found when comparing the top manager's performance to the optimal solution.

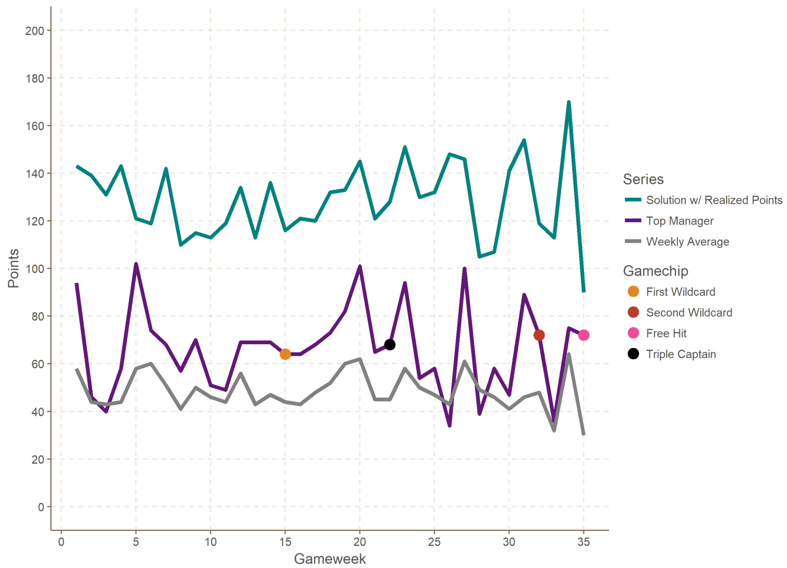

Figure 7.4: Weekly average, top performer and solution with realized points for every gameweek.

	Mean	Percentage of optimal solution
Optimal solution	128.57	100.00 %
Top Manager	66.54	51.75 %
Top 5 %	55.83	43.42 %
Top 10 %	54.30	42.23 %
Top 20 %	52.20	40.60 %
Top 30 %	50.40	39.20 %
Top 40 %	48.50	37.72 %
Top 50 %	46.50	36.17 %

Table 7.5: Comparing managers to optimal solution.

A curious observation made from in Table 7.5 is that the difference in mean between the top 50% and the winner is only slightly above 20 points per round, while the difference between the optimal solution and the winner is above 60 points. Further, in mean, the difference between the winner and top 5% is 10.71 points per gameweek. In comparison, the difference between finishing among the top 10% and top 5% is only an increased mean of 1.43 points per gameweek. A plausible explanation for the small difference is that many managers select the same players. For instance, at some point, Mohammed Salah was selected by more than 63% of the managers. Thus, for the forecast-based optimization model, a slight increase in mean score can be enough to improve the performance in terms of overall ranking significantly.

7.3 Forecast-Based Solutions

In this section, the results of the forecast-based optimization model run with forecasts generated by the three different forecasting methods are presented. First, the performances when the gamechips are disregarded are presented and discussed. Then, the performance of each method when gamechips are included is presented. This serves two purposes. First, it is done in order to compare the methods against each other without the complicating inclusion of gamechips. Secondly, it allows a detailed examination of the effect of gamechips. Finally, the results are summarized, and plausible explanations for the observations made are discussed. To evaluate the performance of the methods, their performances are compared to that of the optimal solution with realized points, the top manager and the weekly average among all managers. Concerning computational time, the model is solved in a matter of seconds for all gameweeks.

7.3.1 Results Disregarding Gamechips

In Figure 7.5, the points obtained by the different forecasting methods are plotted for each gameweek. Table 7.6, summarizes the results. Notice that for the Odds method, results are presented for 31 gameweeks, as data only was available for these gameweeks.

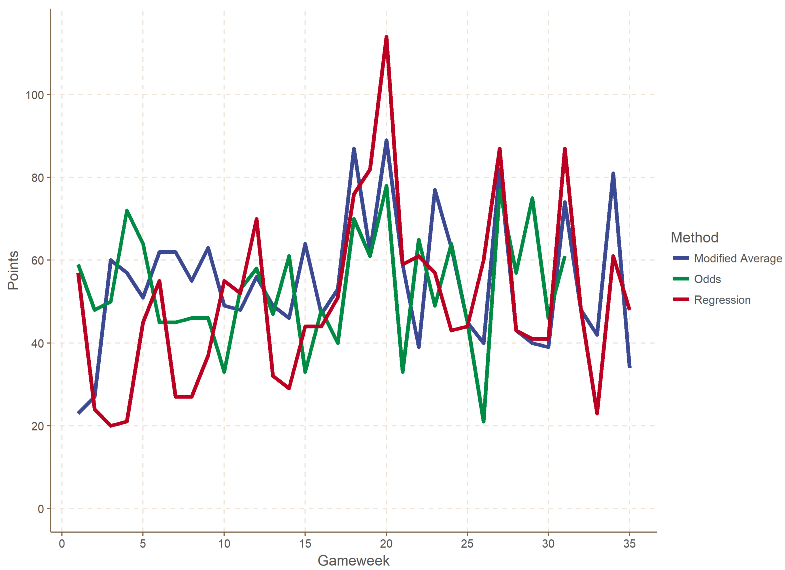

Figure 7.5: Results of the different forecasting methods.

Comparisons of overall rankings are made in terms of mean points. This is done because 31 gameweeks are considered for the Odds method and 35 gameweeks for the two other methods. From Table 7.6, it is clear that the Modified Average and the Odds method outperform the Regression method. Moreover, it is clear the mean number of penalized transfers vary considerably between the methods. The difference between the Modified Average and Regression is particularly apparent. Furthermore, the Regression method has the highest volatility measured in standard deviation. The standard deviations presented is empirical standard deviation calculated gameweek by gameweek.

Solution method	Mean	St.Dev.	Overall ranking	Mean penalized transfers
Modified Average	54.74	15.85	Top 8%	0.23
Regression	50.43	20.73	Top 30%	1.80
Odds (31 gameweeks)	53.23	13.94	Top 13%	0.81

Table 7.6: Results disregarding gamechips.

Modified Average

From Figure 7.5, it is evident that the Modified Average performs poorly in the first two gameweeks. A reasonable explanation is that the forecasts are based on the previous season. Also, several players are not taken into consideration in the first gameweeks, for instance due to promotions and international transfers. In fact, the Modified Average selects Gary Cahill and Cesc Fabregas, whom both received a red card, in the first gameweek. With Cahill selected as captain, they deducted a total of -7 points.

It is observable that the weekly results tend to stabilize from gameweek 3 to gameweek 18. This can stem from the fact that track length now only includes realized points from previous gameweeks in the 2017/2018 season. Additionally, players from outside the Premier League, including players from newly promoted teams, are now assigned expected points.

From gameweek 19 to gameweek 35, the Modified Average improves its performance compared to that in the first part of the season. Again, a plausible explanation is that the forecasts are better, by for instance being able to capture players that are performing consistently well.

Regression

The Regression method displays the highest fluctuations in terms of standard deviation and reaches the highest weekly realization. However, the method is outperformed by the two other methods in terms of mean. The weakest weekly performances are obtained at the beginning of the season. This is probably explained by the fact that at the beginning of the season, each player is listed with low values for variables such as total goals, total assists, total saves etc., making it hard to distinguish the players from each other. Additionally, some players are not considered for the first gameweek due to transfers and promotions. The method achieves its highest weekly score in gameweek 20, earning 114 points. By comparison, the maximum number of points obtained by a manager this gameweek was

156 points. However, that manager played the Triple Captain. Adjusting for the gamechip use, the best manager earned 137 points, leaving a difference of only 23 points.

Odds

With a mean of 64 points in the 5 first gameweeks, the Odds method greatly outperforms the other methods in this interval. By comparison, the Modified Average earns a mean of 51 points and the Regression method a mean of 45 points. Furthermore, it has the overall lowest standard deviation. One explanation is that the forecasts of points are solely based on bookmakers predictions. Since bookmakers are professionals, it is not expected that the accuracy of their odds will change much over the season. Of course, bookmakers will adjust their calculations based on recent performance of players and teams. However, these adjustments will not be as critical as the adjustments for the two other solution methods.

The Odds method is beaten by the Modified Average method. An explanation is that the method has a limitation since forecasts only can be made one gameweek ahead. That is, the sub-horizon must be equal to 1. Note that the performance of the Odds method is not a direct indication of bookmakers' abilities to set probabilities as they are profit-seekers.

Comparing Modified Average with the solution suggested by Bonomo et al. (2014)

The reader is reminded that the Modified Average method is inspired by the forecasting approach suggested for the Argentinian Fantasy League by Bonomo et al. (2014). However, it was developed with the aim of improving the aforementioned approach. Therefore, it is interesting to compare the performance of the two different methods. As mentioned in Section 6.2, Bonomo et al. (2014) set the track-length to 3 gameweeks and the sub-horizon to 1 gameweek. However, the Argentinian Fantasy League does not penalize transfers. Thus, in order to make the comparison as fair as possible, both methods are run with a penalty term, R, set equal to 4. Hence, the optimal combination of track-length and sub-horizon with a penalty 4 is chosen for the Modified Average. These values correspond to a track-length of 5 gameweeks and sub-horizon of 1 gameweek, as can be confirmed in Table 6.2. From Figure 7.6 and Table 7.7, it is clear that the Modified Average outperforms the approach suggested by Bonomo et al. (2014) and that for the 2017/2018 season the Modified Average method is indeed an improvement of the method suggested by Bonomo et al. (2014).

Solution method	Mean	St.Dev.	Overall ranking	Mean penalized transfers
Modified Average	47.77	21.07	Top 43%	1.86
Bonomo et al. (2014)	42.31	17.43	Top 69%	2.23

Table 7.7: Performance of Modified Average and Bonomo et al. (2014).

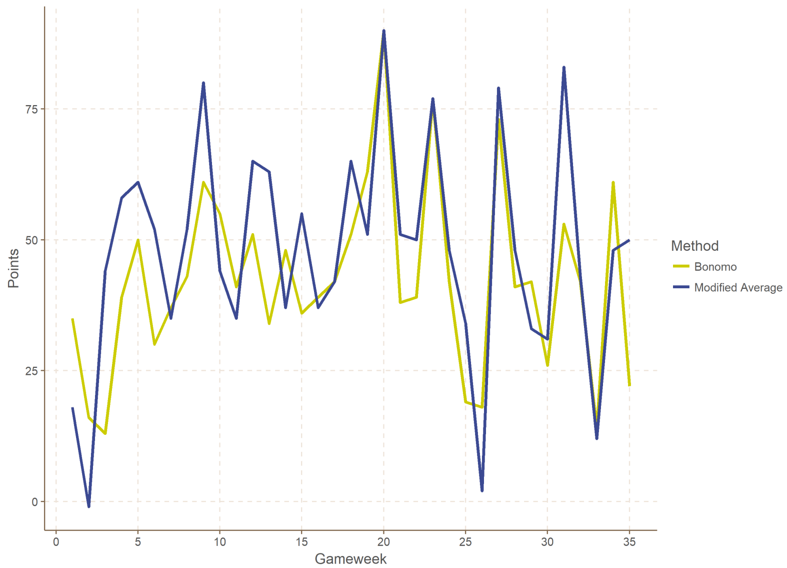

Figure 7.6: Results of Modified Average and the solution approach suggested by Bonomo et al. (2014).

7.3.2 Results Including Gamechips

In Table 7.8, a summary of the performance of each model when gamechips are included is presented. For the sake of legibility, the abbreviation GC is occasionally used for gamechips in tables and figures. In general, by comparing the mean obtained with and without gamechips, it is apparent that the effect of gamechips is not consistent across methods. The performance of the Regression method is improved, while both the Modified Average and Odds method perform worse.

Solution method	Mean	Std.Dev.	Overall ranking	Mean penalized transfers
Modified Average	53.74	16.62	Top 12%	0.23
Regression	54.31	20.32	Top 10%	1.20
Odds (31 gameweeks)	50.23	15.17	Top 30%	0.48

Table 7.8: Results including gamechips.

In Figure 7.7, each method is plotted with and without gamechips. In the following, it is discussed how each gamechip have affected the solutions. Note that for each method, points obtained are similar for the first 6 gameweeks. As mentioned in Section 6.3, the model was allowed to use a Wildcard in gameweek 9. Hence, this decision was first considered in gameweek 7 due to the sub-horizon of 3 gameweeks. Therefore, different decisions are made from that point on.

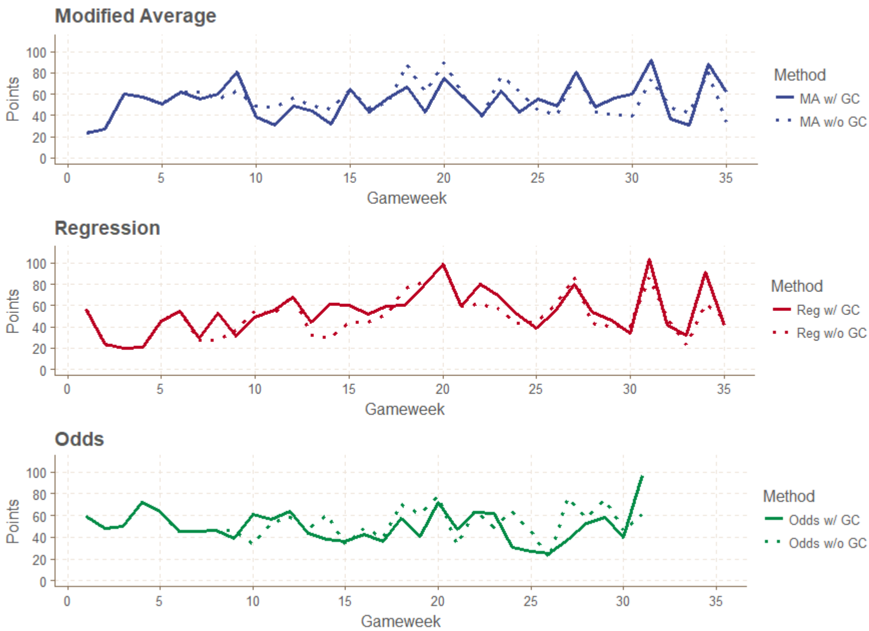

Figure 7.7: Performance of forecasting methods with gamechips.

To evaluate the impact of the first Wildcard, comparison of performances are made from gameweek 7, as this is the first gameweek were decisions regarding the Wildcard are introduced. Moreover, only the performance until gameweek 19 is measured, as decisions regarding the Triple Captain are first introduced in this gameweek. As expected, Table 7.9 shows that all methods perform a substantially higher number of transfers when the Wildcard is played. However, the effect of these transfers appears to be rather limited for the Modified Average and Odds method. The Regression, on the other hand, performs considerably better when the Wildcard is used.

Method	Mean GW 7-19	Transfers GW 9
Modified Average w/GC	51.08	11
Modified Average w/o GC	57.00	1
Regression w/GC	53.92	12
Regression w/o GC	48.15	5
Odds w/GC	46.62	11
Odds w/o GC	49.31	1

Table 7.9: Performance from gameweek 7 to 19.

As can be seen in Table 7.10, all methods select either Harry Kane or Son Heung-min as Triple Captain in gameweek 22. First, it is worth noting that all players play for Tottenham, one of two teams that had a double gameweek. Furthermore, the Regression and Odds method had both Son and Kane in their starting line-ups. Thus, both the Regression and Odds had a higher forecast for Son than for Kane. A curious observation is that it was speculated in whether Kane would play both matches or not.[1] Such events can be caught by the Regression method by the variables *transfers in* and *transfers out*, and by the odds set by the decisions of bookmakers. For the Modified Average method, however, there is no way to pick up such signals. By not being able to incorporate any qualitative information such as line-up rumors, the Modified Average has a disadvantage when it comes to making the critical decision as whom to select for Triple Captain.

Method	Triple Captain	Points
Modified Average w/GC	Harry Kane	9
Modified Average w/o GC	Harry Kane	6
Regression w/GC	Son Heung-min	36
Regression w/o GC	Son Heung-min	24
Odds w/GC	Son Heung-min	36
Odds w/o GC	Son Heung-min	24

Table 7.10: Performance of Triple Captain in gameweek 22.

From Table 7.11, it is apparent that all methods perform better in the blank gameweek 31, the gameweek the Free Hit is used. Furthermore, by examining Table 7.11, it is clear that difference is 17 points or more for all methods. Furthermore, it is notable that only the Odds method makes a substantial amount of transfers ahead of the blank gameweek. This can be explained by the sub-horizon of 1, as the other methods can plan 3 matches ahead. It is worth noting that all methods obtain a relatively high score in gameweek 31. This is mostly explained by the fact that Mohammed Salah earned 29 points for the particular gameweek. As all methods had him as captain, they earned a score of 58 points on him alone.

Method	Points GW31	Transfers	Players playing
Modified Average w/GC	92	14	11
Modified Average w/o GC	74	0	5
Regression w/GC	104	13	11
Regression w/o GC	87	2	9
Odds w/GC	97	15	11
Odds w/o GC	61	10	11

Table 7.11: Performance of Free Hit in gameweek 31.

[1] See article ”Priceless Harry Kane a doubt for Tottenham’s game with Swansea” at https://www.independent.co.uk/sport/football/premier-league/tottenham-spurs-transfer-news-harry-kane-no-price\-illness-swansea-doubt-team-news-a8133251.html

As the second Wildcard and Bench Boost are played in cooperation, they are discussed together. The second Wildcard was played in gameweek 33, preparing for the double gameweek 34. Data for these gameweeks are only available for the Modified Average and Regression method, so the discussion is limited to these methods. Since the Bench Boost is played in gameweek 34, it is considered by the model from gameweek 32. Gameweek 31 is the gameweek when the Free Hit is played. Therefore, only the gameweeks 32-34 are considered when evaluating the effect of the second Wildcard and Bench Boost. Table 7.12 shows that in terms of mean points obtained, the effect of the Wildcard and Bench Boost for gameweek 32-34 is positive only for the Regression method. In particular, the Modified Average performs worse in the gameweek the Wildcard is played, while the Regression performs better. Furthermore, both methods improve in the gameweek the Bench Boost is played. However, as 15 players are awarded points, this is expected.

Method	GW 32	GW 33 (W)	GW 34 (BB)	Mean
Modified Average w/GC	37	31	88	52
Modified Average w/o GC	48	42	81	57
Regression w/GC	42	32	91	55
Regression w/o GC	48	23	61	44

Table 7.12: Performance of Wildcard and Bench Boost in GW32-34.

7.3.3 Summary

Solution method	Mean	Std.Dev.	Overall ranking	Mean penalized transfers
Modified Average w/ GC	53.74	16.62	Top 12%	0.23
Modified Average w/o GC	54.74	15.85	Top 8%	0.23
Regression w/ GC	54.31	20.32	Top 10%	1.20
Regression w/o GC	50.43	20.73	Top 30%	1.80
Odds (31 gameweeks) w/ GC	50.23	15.17	Top 30%	0.48
Odds (31 gameweeks) w/o GC	53.23	13.94	Top 13%	0.81
Realized points	128.57	16.50	N/A	4.23
Weekly Average	48.53	7.99	Top 40%	-
Top Manager	66.54	18.28	Top 1	0.11

Table 7.13: Summary of the results.

In Table 7.13, the results of all different methods with and without gamechips are summarized. Furthermore, the solution with realized points is given as a benchmark. However, the model with realized points significantly beats the forecast-based optimization model.

Therefore, as a more realistic benchmark for top performance, the performance of the top manager is included. As a benchmark of average performance, the weekly average is used.

As the season consists of 38 gameweeks, comparing overall performance is somewhat problematic when only gameweek 1-35 is considered. For instance, the top manager still has a gamechips left. A different number of gameweeks is also an issue when comparing the Odds method with the two other methods. Nevertheless, the results should be indicative as more than 5.9 million manager participates in FPL. Once again, it is stressed that as testing is only performed on one season, overly categorical statements should not be made. Nonetheless, with either the Modified Average method without gamechips, or the Regression based approach with gamechips, the forecast-based optimization model appears to have the potential to reach the top 10% of managers in the FPL. Furthermore, a position in the top 30% is achieved by all methods. Thus, beating the average weekly performance indeed appears to be obtainable. On the other hand, all models are considerably out-performed by the top manager. Without further improvement, winning the FPL with the forecast-based optimization model appears to be unobtainable.

The Odds method performs best over the first five gameweeks. As the Odds method is based on bookmakers odds, it is not as heavily influenced by past season's performance as the Modified Average and Regression method. In fact, from the previous season's Dream team, i.e., the 11 players that collected most points, only one player had a score above two points in the first gameweek. In addition, the Odds method does not exclude players due to promotion or transfers to Premier League ahead of the season. The observation is substantiated by the fact that in gameweek 5, the gameweek when entire track-length is first founded in the 2017/2018 season, the Modified Average method starts to outperform the Odds method. Hence, the Odds method appears to be the method best suited for forecasting in the earliest gameweeks of the season.

As evident from Table 7.13, the number penalized transfers are considerably different between the methods. In particular, it is high for the Regression method. It can be argued that the higher penalty term incorporated for the Modified Average method is only a compensation for inaccurate forecasts. Thus, it appears as if the Regression method in particular also could benefit from a parameter tuning approach concerning the penalty term R. However, this was constrained by the availability of data.

In general, the methods exhibit different variability. Based on the standard deviations presented in Table 7.13, the Regression method is most volatile, while the Odds method is the least. One explanation is that the Regression method does not account for recent performance in as a direct manner as the Modified Average method does by averaging points in the last rounds. Moreover, odds are set by the standard procedures of professional bookmakers, while the other methods are solely based on data and will to a higher degree incorporate extreme and unexpected past events. Additionally, the performances for all methods appear to fluctuate more in the second half of the season. One sensible reason is the emergence of blank and double gameweeks.

The effect of gamechips is not consistent for all methods. The Triple Captain and Free Hit appear to boost performance for all methods. The Second Wildcard in combination with Bench Boost appear to have a favorable effect on the Regression model, but not for the Modified Average. The same goes for the first Wildcard. This can be due to several reasons. First, for the Regression method in Section 6.2.2, it is found that time-series effect of points is rather limited. Hence, for the Regression, only variables such as total amount of goals, assists, saves etc. are considered. When playing the Wildcard, the entire selected squad can be changed, and the new selected squad is carried forwards. If the Wildcard is played and forecasts are based on the Regression method, the performances of players over all the previous gameweeks are considered. For the Modified Average, on the other hand, the forecasts are only based on the last 5 matches. When considering to change the entire selected squad, i.e., playing the Wildcard, taking all previous gameweeks into consideration seems to be the preferable approach. Secondly, the parameter tuning is done without gamechips in mind. That is, neither track length nor sub-horizon are optimized with respect to the Wildcard. As all gamechips are meant to provide opportunities that benefit managers, it is counter-intuitive that the results are worse when they are included. In the end, the varying results of the model run with gamechips is due to the implementation of the gamechips. In fact, if the gamechips are used after the model is run without first taking them into consideration, they can be guaranteed not to have a negative influence on performance. The Wildcards and Free Hit can be used only in order to ”pay” for penalized transfers after decisions are made, while the Triple Captain and Bench Boost can be played after the captain and the selected squad are decided. Thus, the gamechips will only have a positive or zero effect. As seen, however, this is not guaranteed when implementing the gamechips with the strategies suggested in this thesis. A further possible explanation of the unsatisfactory performance when gamechips are included is that an implementation only allowing gamechips to be played in particular gameweeks is not ideal for an optimization model. Since the use of gamechips are fixed to certain gameweeks, the solution space is decreased and the solutions found are sub-optimal. The implementation of gamechips is regarded as one of the areas with most potential for improvement in the forecast-based optimization model.

7.4 Risk Handling

In this section, the effect of risk handling in the forecast-based optimization model is examined. That is, constraints (5.12) - (5.16) are included, and the model is run with changing thresholds. However, since the performance is measured in the unit points, the thresholds are from now on denoted as standard deviations rather than variances. Further, it is emphasized that a *threshold* refers to a value chosen before a model is run. *Realized standard deviation*, on the other hand, refers to an empirical value computed after the model is run for all gameweeks. Forecasts from the Modified Average method are used and gamechips are included. Including the risk handling constraints increase the computational time significantly. For the lowest threshold considered, the model is solved in approximately 20 minutes for all gameweeks.

In Sections 3.4 and 5.4, it is pointed out that the FPLDP resembles a classic portfolio optimization problem. Therefore, it is interesting to investigate if there exists a similar expected return/risk trade-off in the FPLDP as it does in portfolio optimization. Accordingly, it is studied how varying the thresholds affects expected points for the first sub-problem. The mathematical model maximizes a team's expected points. Hence, one would assume that when the threshold is decreased, the expected points of the team also decreases. This would coincide with the *efficient frontier* found in portfolio optimization. The efficient frontier consists of the optimal set of portfolios that gives the highest expected points for a certain level of risk. In Figure 7.8, the expected points, i.e., the objective value, obtained by the optimization model for the first sub-problem is plotted against different thresholds. The curve exhibits similar behavior as the efficient frontier but is not smooth. This can be explained by the fact that in the FPLDP one can either include a player or not. That is, the decision is binary. In other words, opposed to the classic portfolio optimization problem, one cannot hold a fraction of an asset in the FPLDP.

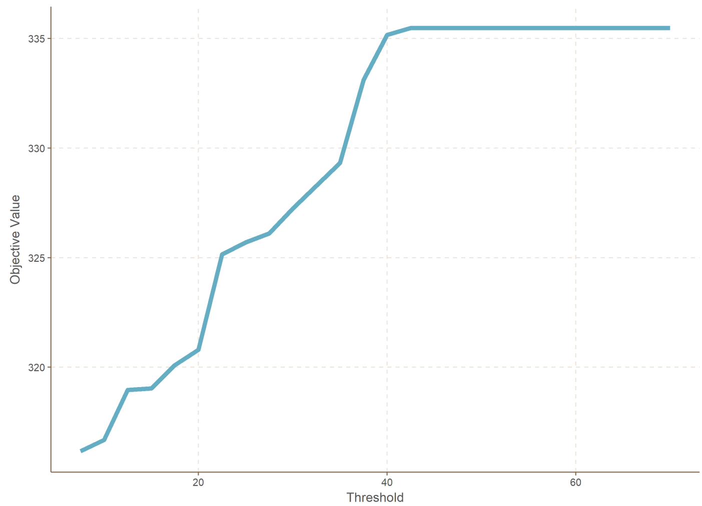

Figure 7.8: Threshold plotted against objective value in the mathematical model for the first sub-problem.

In Figure 7.9, the mean points obtained in the 2017/2018 season when the model is run with different thresholds are plotted. From Figure 7.9, it is observable that the plot deviates from that of the efficient frontier. However, some similarities can be found. Higher thresholds are for instance associated with higher mean scores, and lower thresholds are associated with lower means. When the constraints are non-binding, the mean stabilizes at 54.34 points. Intuitively, the mean should have been 53.74 points, the same as when the model is run without the constraints. However, after examining the objective values, it is clear that there are different solutions with the same expected value for the first sub-problem. Thus, there is no guarantee that the solver selects the same solution. A possible way around is a different implementation of variance constraints. A *trade-off approach* (Mansini et al., 2015) could have been an interesting alternative, where instead of bounding the variance it is taken into consideration by moving it to the objective function.

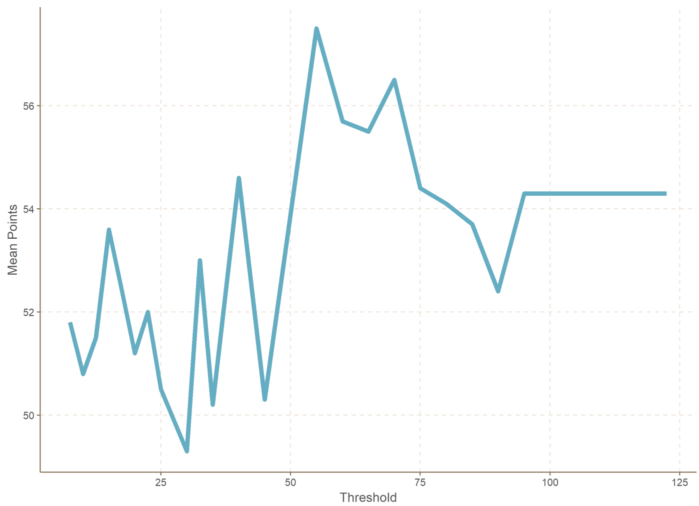

Figure 7.9: Performance with varying thresholds.

As seen in Figure 7.9, the best mean score is not reached for an unlimited threshold, but for a threshold of 55 points. Here, the model achieved a mean of 57.54 points, corresponding to a spot among the top 1.5% of all managers. In general, the solutions outperform the solution with an unlimited threshold in the area with a threshold between 55-70 points. In Figure 7.10, the solutions with threshold 55 and 7.5 are plotted against the solution without a threshold. It is noticeable that the variance appears to be higher than for the unlimited case when the threshold is 55, but not 7.5. In Table 7.14, the mean and realized standard deviation for thresholds from 7.5-12.5 and 55-65 are shown. Based on the table, it is clear that the goal of achieving lower realized standard deviation in the solution is reached with lower thresholds. Furthermore, in these cases, the mean is reduced, analogous to the case of portfolio optimization. Observe, however, that when the thresholds are between 7.5-12.5, they are unable to keep the realized standard deviations below the thresholds. The realized standard deviation for the thresholds in the region 55-65 exceeds the realized standard deviation of the unlimited case. For threshold between 55-65, however, the realized standard deviation is well below the threshold. Yet, as the threshold is approximately as high as the mean points obtained, it is counter-intuitive that such a threshold should be binding. Thus, the positive effect in mean is probably related to the nature of the problem, not to the variance reduction per se. That is, having a factor decide how many players of the same and opposing teams to include in the selected squad can be a beneficial strategy in terms of improving performance, even if it does not necessarily reduce the realized standard deviation of points obtained.

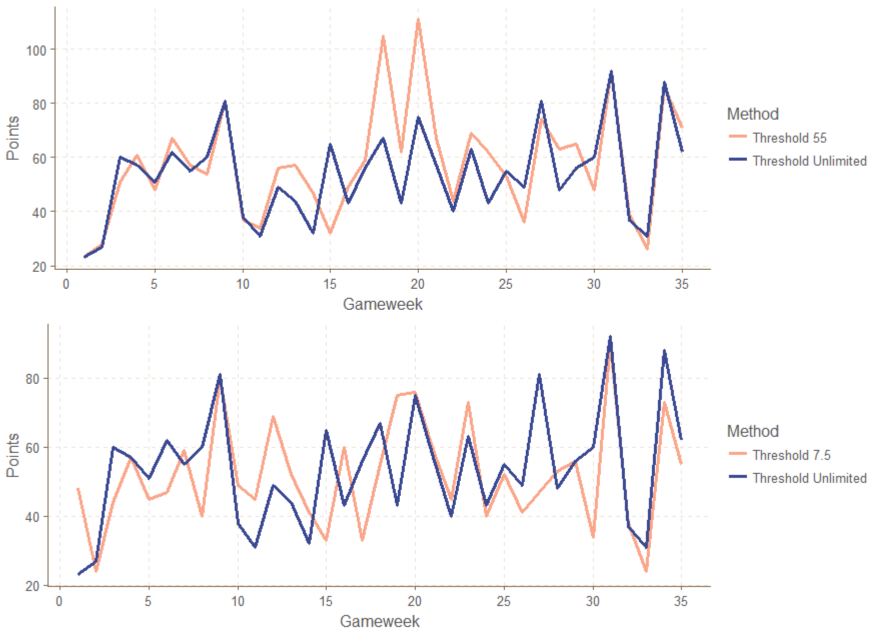

Figure 7.10: Performance with threshold 7.5 and 55 compared against performance with no threshold..

Threshold	Mean	Realized St.Dev.
7.5	51.77	15.7
10	50.77	14.6
12.5	51.49	15.9
Unlimited	53.74	16.8
55	57.54	20.8
60	55.71	18.1
65	55.51	18.8

Table 7.14: Results for the high and low thresholds.

To further examine the effect of the risk threshold, the starting line-up with a threshold of 7.5 and 55 are compared after the first Wildcard is played, i.e., in gameweek 9. The starting line-up for threshold 7.5 and 55 are given in Table 7.15 and Table 7.16, respectively. By comparing the teams, it is clear that players from 9 different teams were in the starting line-up when the threshold was 7.5, while only 6 different teams were featured when the threshold was 55. Furthermore, with a threshold of 55, three midfielders from the same team, Manchester City, were selected. Additionally, no players in the starting line-up faced each other. With a threshold of 7.5, on the other hand, four pairs of players faced each other. In two of these four pairs, a defender faced a midfielder. As seen in Table 6.13, defenders and midfielders of opposing teams are the relationship with the highest negative correlation. Thus, the two cases illustrate the impact risk handling has on the starting line-up.

Player	Position	Team	Opponent
Hart	GLK	WHU	BHA
Ward	DEF	BUR	MCI
Azpilicueta	DEF	CHE	WAT
Valencia	DEF	MUN	HUD
Groß	MID	BHA	WHU
Bakayoko	MID	CHE	WAT
Sane	MID	MCI	BUR
Choupo-Moting	MID	STK	BOU
Richarlison	MID	WAT	CHE
Diouf	FWD	STK	BOU
Kane	FWD	TOT	LIV

Table 7.15: Starting line-up in gameweek 9 when threshold is set to 7.5.

Player	Position	Team	Opponent
Hart	GLK	WHU	BHA
Monreal	DEF	ARS	EVE
Lascelles	DEF	NEW	CRY
Fonte	DEF	WHU	BHA
Bruyne	MID	MCI	BUR
Sane	MID	MCI	BUR
Sterling	MID	MCI	BUR
Choupo-Moting	MID	STK	BOU
Eriksen	MID	TOT	LIV
Diouf	FWD	STK	BOU
Kane	FWD	TOT	LIV

Table 7.16: Starting line-up in gameweek 9 when threshold is set to 55.

To summarize, a threshold between 55-70 points appears to have a positive effect on performance in terms of mean, but not in terms of reducing realized standard deviation. Furthermore, a threshold of 7.5-15 points appears to be able to lower realized standard deviation at the cost of mean obtained. An interval where both the realized standard deviation is decreased, and the performance is increased, is not found. Again, it is stressed that all results are based on data from only two seasons and are indicative at best.

Chapter 8

Concluding Remarks

This thesis describes an optimization model for the Fantasy Premier League Decision Problem by suggesting an approach for optimizing Fantasy Premier League decisions. To the author's knowledge, the mathematical model presented is the first of its kind for the Fantasy Premier League. Also, to the best of the author's knowledge, the concept of modeling gamechips have never been considered in Fantasy Sports, and risk handling has never been considered in the Fantasy Premier League.

The mathematical model has been solved with realized points in order to obtain the optimal solution. Furthermore, the mathematical model has been solved by the use of a rolling horizon heuristic with forecasts of player points as input. The forecasts are generated by three different methods: the Modified Average method, the Regression method and the Odds method. Moreover, the gamechips are implemented using qualitative strategies. The strategies are mostly developed by considering whether a gameweek is blank or double. In addition, risk handling constraints are added and their impact is studied.

Solving for realized points reveals a considerable potential for improvement in performance in the Fantasy Premier League. Even the top manager among 5.9 million users only manages to achieve approximately 50% of the optimal solution. This speaks in volumes of the nature of uncertainty in the FPLDP. Furthermore, it motivates researching the problem, as the top performing solutions are far from optimal. For the forecast-based optimization model developed, achieving a position among the top 30% of all managers appears obtainable by all marks. Moreover, forecasts based either on the Regression or Modified Average method are deemed most promising. For these methods, a top 10% finish appears to be within reach.

The effect of adding risk handling constraints related to variance has been analyzed. The results are promising and indicate that a strategy that considers which players face each other in a gameweek can either be used to limit variance or boost performance, depending on a risk preference threshold. Obtaining both effects simultaneously, however, appears to be impossible. In the best cases, the model reaches approximately the top 1.5% of managers when risk handling is introduced alongside forecasts from the Modified Average method. Winning the Fantasy Premier League, i.e., finishing as the top manager, appears to be unachievable with the methods suggested in this thesis.

Even though the results are encouraging, there exist some drawbacks worth addressing. Most notably, the implementation of the gamechips is unsatisfactory in the sense that with forecasts from the Modified Average and Odds method, the performances are aggravated when gamechips are introduced. Another drawback is the sparsity of data. As data for only two seasons are available, the results must be interpreted carefully.

As 5.9 million managers compete in the Fantasy Premier League, competition is fierce. In fact, a mean of only 1.9 points per gameweek separates a top 10% finish from a top 20% finish in 2017/2018 season. As of now, the mathematical model developed can already be used as a decisions support tool. For one, it can be run ex-post to answer insightful questions like, ”What would have been the optimal transfer last gameweek?” or, ”What is the optimal team so far?”. Furthermore, the forecast-based optimization model can be used as a complementary support tool for FPL managers. For instance, it can propose the k best set of transfers for a manager's team in each round. This can be done by running the model k times, in each case constraining the optimal solutions previously generated. The growing attention from commercial actors and the monetary figures seen in other Fantasy Sports, such as in Daily NFL Fantasy Sports, indicate that improving average performance by only a small amount may be of substantial worth in the future. Thus, further development of the forecast-based optimization model, as well as more research on the FPLDP in general, has great potential value.

Chapter 9

Recommendations for Further Research

The FPLDP is not an in-depth researched area. In fact, only two articles on Fantasy Football were found (Matthews et al. (2012); Bonomo et al. (2014)) and only one of them specifically addressed the FPLDP (Matthews et al., 2012). To the author's knowledge, the mathematical model presented is the first of its kind for the Fantasy Premier League, and the concept of modeling gamechips have never been considered in Fantasy Sports before. Thus, the novelty of the FPLDP makes recommendations for future research topics a vital issue to address.

When the model is run with realized points, the results demonstrate that there is a significant gap between the optimal solution and the top manager. Moreover, considering the sample size of 5.9 million participating managers, it illustrates that there is much uncertainty in the FPLDP. Thus, instead of taking the deterministic solution method adopted in the thesis, using a stochastic solution method or a method based on machine learning are ideas for further research. Modeling the FPLDP as a Markov Decision Problem is also possible, as done by Matthews et al. (2012).

A potential for improvement of the forecast-based optimization model lies in the implementation of the gamechips. As seen in the thesis, the implementation does not guarantee an increased performance when the gamechips are played. Thus, the qualitative strategies suggested do not seem complementary to the forecast-based optimization model. The Wildcard, in particular, could benefit from another form of implementation. Thus, examining the effect of different implementations is an exiting topic for further research. One interesting approach is to let the model decide when to use the gamechips in the rolling horizon heuristic, but punish the objective function for the use. This will keep the model from using all gamechips in the first sub-problems. Another approach would be to implement the gamechips analogous to the exercise of a financial option. Other approaches related to the Triple Captain include using a risk-adjusted measure analogous to the Sharpe Ratio (Sharpe, 1994) or an approach based on the optimal solution of the Secretary Problem (Lindley (1961); Bruss (1984)).

Multiple approaches with a potential of improving forecasts of player points exist. Some are solely related to the fact that in the future, more data will be available. For example, a parameter tuning can be conducted for the Regression method. Furthermore, with more data available, one can fit the regression model by only considering matches at a similar stage of previous seasons. For instance, total goals can have a high predictive power of realized points in the second part of the season, but perhaps not in the first couple of gameweeks. Methods for handling players without data due to for instance transfers or promotions are also of interest. These could, for instance, be based on the assumption that players of the same price are expected to perform similarly.

Other topics for future research related to the forecasts are not founded in the scarcity of data. Instead, they are connected to discoveries made in the thesis. For instance, it is clear that the Odds method performs best at the beginning of the season. Thus, combining forecasting methods is an interesting approach to improve performance. For example, the Odds method could be used for the first 5 gameweeks before switching to the Modified Average method for the rest of the season. In combination, the model could be allowed to use the Wildcard in the gameweek forecasting methods are changed. Alternatively, odds could be used as an explanatory variable in the regression. Finally, forecasting of future price could be used as a measure to profit from trading players, thus obtaining a higher budget.

An analysis of risk handling was presented in this thesis. The analysis indicated that the risk handling could be used to improve the performance or alternatively to reduce risk. However, further research is required before more conclusive statements can be made.

Bibliography

Adkins, R., Paxson, D., 2013. Deterministic models for premature and postponed replacement. Omega 41, 1008–1019.

Alexander, C., 2008. Market Risk Analysis Volume 1, 1st Edition. John Wiley & Sons.

Bar-Eli, M., Avugos, S., Raab, M., 2006. Twenty years of ”hot hand” research: Review and critique. Psychology of Sport and Exercise 7, 525–553.

Bartsch, T., Drexl, A., Króger, S., 2006. Scheduling the professional soccer leagues of Austria and Germany. Computers & Operations Research 33, 1907–37.

Becker, A., Sun, X. A., 2016. An analytical approach for fantasy football draft and lineup management. Journal of Quantitative Analysis in Sports 12, 17–30.

Bertoni, S., 2015. Fantasy Sports Site FanDuel Scores $1 Billion Valuation. URL: https://www.forbes.com/sites/stevenbertoni/2015/07/14/fantasy-sports-site-fanduel-scores-1-billion-valuation/#dec3ba9719e2, Online; accessed 22.04.2018.

Bjertnes, V. R., Nørstebø, O., Vabo, E., 2016. Valuing Individual Player Involvements in Norwegian Association Football. Master's thesis, Norwegian University of Science and Technology.

Bonomo, F., Cardemil, A., Durán, G., Marenco, J., Sabán, D., 2012. An application of the traveling tournament problem: The argentine volleyball league. Interfaces 42, 245–259.

Bonomo, F., Durán, G., Marenco, J., 2014. Mathematical programming as a tool for virtual soccercoaches: a case study of a fantasy sport game. International Transactions in Operational Research 21, 399–414.

Boon, B., Sierksma, G., 2003. Team formation: Matching quality supply and quality demand. European Journal of Operational Research 148, 277–292.

Boulier, B., Stekler, H., 2003. Predicting the outcomes of national football league games. International Journal of Forecasting 19, 257–270.

Boulier, B., Stekler, H., Amundson, S., 2006. Testing the efficiency of the national football league betting market. Applied Economics 38, 279–284.

Bruecker, P. D., den Bergh, J., Belien, J., Demeulemeester, E., 2015. Workforce planning incorporating skills: State of the art. European Journal of Operational Research 243, 1–16.

Bruss, F. T., 1984. A unified approach to a class of best choice problems with an unknown number of options. Annals of Probability 12, 882–891.

Burke, D., Meek, B. P., Norwood, J. M., 2016. Exploring the legality of the lucrative world of fantasy sports. Journal of Legal, Ethical and Regulatory Issues 19, 38–55.

Chand, S., McClurg, T., Warda, J., 2000. A model for parallel machine replacement with capacity expansion. European Journal of Operational Research 21, 519–531.

Chang, T.-J., Meade, N., Beasley, J., Sharaiha, Y., 2000. Heuristics for cardinality constrained portfolio optimisation. Computers & Operations Research 27, 1271–1302.

Clarke, S., Dyte, D., 2000. Using official ratings to simulate major tennis tournaments. International Transactions in Operational Research 7, 585–594.

Clarke, S., Norman, J., 1995. Home Ground Advantage of Individual Clubs in English Soccer. Journal of the Royal Statistical Society. Series D (The Statistician) 44, 509–521.

Coleman, J. B., 2012. Identifying the "Players" in Sports Analytics Research. Interfaces 42, 109–118.

Constantinou, A., Fenton, N., 2012. Determining the level of ability of football teams by dynamic ratings based on the relative discrepancies in scores between adversaries. Journal of Quantitative Analysis in Sports 9, 37–50.

Crowder, M., Dixon, M., Ledford, A., Robinson, M., 2002. Dynamic modelling and prediction of English football league matches for betting. The Statistician 51, 157–168.

Davis, P., Richards, A., Ollerton, J., 2007. Determining the composition and benefit of the pre-hospital medical response team in the conflict setting. Journey of the Royal Army Medical Corps 153, 269–273.

Della, F. C., Oliveri, D., 2006. Scheduling the professional soccer leagues of Austria and Germany. Computers & Operations Research 33, 1907–37.

Dixon, M. J., Coles, S. G., 1997. Modelling Association Football Scores and Inefficiencies in the Football Betting Market. Applied Statistics 46, 265–280.

Durán, G., Guajardo, M., Miranda, J., Sauré, D., Souyris, S., Weintraub, A., Wolf, R., 2007. Scheduling the Chilean soccer league by integer programming. Interfaces 37, 539–52.

Durán, G., Guajardo, M., Wolf-Yadlin, R., 2012. Operations research techniques for scheduling Chile's second division soccer league. Interfaces 42, 273–285.

Elo, A., 1978. The rating of chessplayers, past and present. Arco Publishing.

Fantasy Sports Trade Association, 2017. Industry Demographics: Actionable Insights & Insightful Data. URL: https://fsta.org/research/industry-demographics/, Online; accessed 05.02.2018.

FIFA, 2012. Men's Ranking Procedure. URL: http://www.fifa.com/fifa-world-ranking/procedure/men.html, Online; accessed 15.05.2018.

Fry, M. J., Lundberg, A. W., Ohlmann, J. W., 2007. A Player Selection Heuristic for a Sports League Draft. Journal of Quantitative Analysis in Sports 3, 5–5.

Gareth, J., Witten, D., Hastie, T., Tibshirani, R., 2017. An Introduction to Statistical Learning, 8th printing. Springer.

Gibson, M. R., Fry, M. J., Ohlmann, J. W., 2007. An agent-based stochastic ruler approach for a stochastic knapsack problem with sequential competition. Computers and Operations Research 37, 598–609.

Goossens, D., Beliän, J., Van Reeth, D., 2017. Optimization modelling for analyzing fantasy sport games. Information Systems and Operational Research 55, 275–294.

Goossens, D., Spieksma, F., 2009. Scheduling the Belgian soccer leagues. Interfaces 39, 109–18.

Graham, I., Stott, H., 2008. Predicting bookmaker odds and efficiency for uk football. Applied Economics 40, 99–109.

Heitner, D., 2017. The Hyper Growth Of Daily Fantasy Sports Is Going To Change Our Culture And Our Laws. URL: https://fsta.org/research/industry-demographics/, Online; accessed 21.04.2018.

Henz, M., 2001. Scheduling a major college basketball conference revisited. Operations Research 49, 163–168.

Hopp, W., Nair, S., 1994. Markovian deterioration and technological change. IIE Transactions 26, 74–82.

Hvattum, L., Arntzen, H., 2009. Using elo ratings for match result prediction in association football. Norwegian University of Science and Technology.

Hyndman, R. J., Athanasopoulos, G., 2018. Forecasting: Principles and Practice. Monash University, Australia.

Jabali, O., Gendreau, M., Laporte, G., 2012. A continuous approximation model for the fleet composition problem. Transportation Research Part B: Methodological 46, 1591–1606.

Kendall, G., 2008. Scheduling English football fixtures over holiday periods. Journal of the Operational Research Society 59, 743–755.

Kendall, G., Knust, S., Ribeiro, C. C., Urrutia, S., 2010. Scheduling in sports: An annotated bibliography. Computers & Operations Research 37, 1–19.

Kirshner, S., 2011. Optimizing an NBA Team's Approach to Free Agency Using a Multiple-Choice Knapsack Model. MIT Sloan Sports Analytics Conference.

Knust, S., 2010. Scheduling non-professional table-tennis leagues. European Journal of Operational Research 200, 358–367.

Komarudin, D., Guerry, M.-A., Feyter, T., Vanden Berghe, G., 2013. The roster quality staffing problem a methodology for improving the roster quality by modifying the personnel structure. European Journal of Operational Research 230, 551–562.

Kostuk, K., Willoughby, K., 2012. A decision support system for scheduling the Canadian Football League. Interfaces 42, 286–295.

Kristiansen, B. K., Gupta, A., 2017. Developing a model for the fantasy team composition problem. Project Report, Norwegian University of Science and Technology.

Leitner, C., Zeileis, A., Hornik, K., 2010. Forecasting sports tournaments by ratings of (prob)abilities: A comparison for the euro 2008. International Journal of Forecasting 26, 471–481.

Lindley, D., 1961. Dynamic programming and decision theory. Applied Statistics 10, 39–51.

Lund, A., Lund, M., 2017. Pearson Product-Moment Correlation. URL: https://statistics.laerd.com/statistical-guides/pearson-correlation-coefficient-statistical-guide.php, Online; accessed 22.05.2018.

Maher, M., 1982. Modelling Football Association Scores.

Mansini, R., Ogryczak, W., Speranza, M. G., 2015. Linear and Mixed Integer Programming for Portfolio Optimization. Springer International Publishing.

Markowitz, H. M., 1952. Portfolio selection. The Journal of Finance 7, 77–91.

Matthews, T., Ramchurn, S. D., Chalkiadakis, G., 2012. Competing with Humans at Fantasy Football: Team Formation in Large Partially-observable Domains. In: Proceedings of the Twenty-Sixth AAAI Conference on Artificial Intelligence. AAAI'12. AAAI Press, pp. 1394–1400.
URL http://dl.acm.org/citation.cfm?id=2900728.2900925

McCracken, G., 2012. Culturematic: How Reality TV, John Cheever, a Pie Lab, Julia Child, Fantasy Football, Burning Man, the Ford Fiesta Movement, Rube Goldberg, NFL Films, Wordle, Two and a Half Men, a 10,000-year Symphony, and ROFLcon Memes Will Help You Create and Execute Breakthrough Ideas. Harvard Business Press.

Nevill, A., Holder, R., 1999. Home advantage in sport. Sports Medicine.

Newell, S., Easton, T., 2017. Optimizing Tiered Daily fantasy sports – Mathematically Modeling DraftKings NFL Millionaire Maker Tournament. In: Proceedings of the Industrial and Systems Engineering Conference. pp. 289–294.

Ozceylan, E., 2016. A mathematical model using ahp priorities for soccer player selection: A case study. South African Journal of Industrial Engineering 27, 190–205.

Pantuso, G., 2017. The football team composition problem: A stochastic programming approach. Journal of Quantitative Analysis in Sports 13, 113–129.

Pollard, R., Pollard, G., 2005. Long-term trends in home advantage in professional team sports in north-america and england. Journal of Sports Sciences, 1876–2003.

Purdum, D., 2017. Future of sports betting: fantasy sports. URL: http://www.espn.com/chalk/story/_/id/19309213/chalk-why-fantasy-sports-sports-betting-collision-course-us, Online; accessed 14.03.2018.

Rajagopalan, S., Soteriou, A. C., 1994. Capacity acquisition and disposal with discrete facility sizes. Management Science 40, 903–917.

Rasmussen, R. V., 2008. Scheduling a triple round robin tournament for the best Danish soccer league. European Journal of Operational Research 185, 795–810.

Rue, H., Salvesen, Ø., 2000. Prediction and retrospective analysis of soccer matches in a league. Norwegian University of Science and Technology.

Sæbø, O., Hvattum, L., 2015. Evaluating the efficiency of the association football transfer market using regression based player ratings. Norsk Informatikkonferanse.

Sæbø, O., Hvattum, L., 2017. Modelling the financial contribution of soccer players to their clubs. Journal of Sports Analytics xx (20xx) xxx.

Saur, M. C., Starr, K., Husted, M., Newman, A. M., 2012. Scheduling soft- ball series in the rocky mountain athletic conference. Interfaces 42, 296–309.

Schönberger, J., Mattfeld, D., Kopfer, H., 2004. Memetic algorithm timetabling for non-commercial sport leagues. European Journal of Operational Research 153, 102–16.

Shapiro, A., Philpott, A., 2007. A tutorial on stochastic programming.

Sharpe, W. F., 1994. The sharpe ratio. Journal of Portfolio Management 21, 49–58.

Skysports, 2017. Chelsea's N'Golo Kante wins PFA Players' Player of the Year award. URL: http://www.skysports.com/football/news/11095/10848240/chelsea8217s-ngolo-kante-wins-pfa-players-player-of-the-year-awardl, Online; accessed 16.05.2018.

Smith, B., Sharma, P., Hooper, P., 2006. Decision making in online fantasy sports communities. Interactive Technology and Smart Education 4, 347–360.

Summers, A. E., Swartz, T. B., Lockhart, R. A., 2005. Optimal Drafting in Hockey Pools. Department of Mathematics & Statistics working paper.

Tavana, M., Azizi, F., Azizi, F., Behzadian, M., 2013. A fuzzy inference system with application to player selection and team formation in multi-player sports. Sport Management Review 16, 97–110.

Trick, M. A., Yildiz, H., Yunes, T., 2012. Scheduling major league baseball umpires and the traveling umpire problem. Interfaces 42, 232–244.

Trninić, S., Papić, V., Trninić, V., Vukičević, D., 2008. Player selection procedures in sports games. Acta Kinesiologica 2, 24–28.

Tversky, A., Gilovich, T., 1989. The cold facts about the hot hand in basketball. Psychology of Sport and Exercise 2, 16–21.

van Voorhis, T., 2002. Highly constrained college basket- ball scheduling. Journal of the Operational Research Society 53, 603–609.

Wright, M., 2005. Scheduling fixtures for new zealand cricket. IMA Journal of Management Mathematics 16, 99–112.

Wright, M., 2006. Scheduling fixtures for basketball new zealand. Computers & Operations Research 33, 1875–93.

Yang, Y., 2015. Predicting regular season results of nba teams based on regression analysis of common basketball statistics. University of California at Berkeley.

Zenios, S. A., 2005. Practical Financial Optimization. The Wharton Financial Institutions Center.

Appendices

Appendix A

Bonus Points System

The Bonus Points System (BPS) used to allocate bonus points in the Fantasy Premier League is presented in this appendix. The BPS utilizes several statistics supplied by Opta, a sport analytics company, to create a performance score for a player. The players with the top three BPS scores in a given match receive bonus points - three points to the highest-scoring player, two to the second best and one to the third. The information in Table A.1 is obtained from www.premierleague.com.

Action	**BPS**
Playing 1 to 60 minutes	3
Playing over 60 minutes	6
Goalkeepers and defenders scoring a goal	12
Midfielders scoring a goal	18
Forwards scoring a goal	24
Assists	9
Goalkeepers and defenders keeping a clean sheet	12
Saving a penalty	15
Save	2
Successful open play cross	1
Creating a big chance (a chance where the receiving player should score)	3
For every 2 clearances, blocks and interceptions (total)	1
For every 3 recoveries	1
Key pass	1
Successful tackle (net*)	2
Successful dribble	1
Scoring the goal that wins a match	3
70 to 79% pass completion (at least 30 passes attempted)	2
80 to 89% pass completion (at least 30 passes attempted)	4
90%+ pass completion (at least 30 passes attempted)	6
Conceding a penalty	-3
Missing a penalty	-6
Yellow card	-3
Red card	-9
Own goal	-6
Missing a big chance	-3
Making an error which leads to a goal	-3
Making an error which leads to an attempt at goal	-1
Being tackled	-1
Conceding a foul	-1
Being caught offside	-1
Shot off target	-1

Table A.1: Bonus Points System.

Appendix B

The Elo System

The Elo system introduced in Section 3.3.2 is a rating system used for measuring the relative strength level of sport teams and individual athletes. Within football, the Elo system can be used in order to determine a club's Elo rating. The great advantage of the Elo system lies in its simplicity: there is only one value per club at each point in time, the higher the better.

The performance of a team is not measured absolutely, but inferred from wins, losses and draws against other teams. If team A has a Elo rating of R_A and team B has a Elo rating of R_B, team A has an expected score of

$$E_A = \frac{1}{1 + 10^{\frac{R_B - R_A}{400}}} \tag{B.1}$$

when facing team B. Similarly, team B has an expected score of

$$E_B = \frac{1}{1 + 10^{\frac{R_A - R_B}{400}}} \tag{B.2}$$

When teams play each other and win or lose, they exchange points. Hence, the club's Elo rating is updated once a match has been played. Using the expected scores from Equation B.1, team A's updated Elo rating is calculated as

$$R_A^{'} = R_A + (S_A - E_A) \times k \tag{B.3}$$

where S_A is the result (1 for win, 0.5 for draw and 0 for loss). Further, k is a factor which determines the adjustments of the Elo rating. A higher k-value will increase the changes in the rating and the Elo ratings will suffer from more variation. For smaller k-values, more stable Elo ratings are created. In chess, the World Chess Federation suggest that a value of $k = 20$ should be used for players with an Elo rating below 2400.

Similarly, team B's updated Elo rating is given by

$$R_B^{'} = R_B + (S_B - E_B) \times k \tag{B.4}$$

In the following, an illustrative example shows how Elo ratings of two teams are updated once a match between them has been played:

Assume that team A has an Elo rating of 1881 and that team B has a rating of 1650. If team A faces team B, team A has an expected score according to Equation B.1:

$$E_A = \frac{1}{1 + 10^{\frac{1650-1881}{400}}} = 0.791$$

while team B has an expected score of

$$E_B = \frac{1}{1 + 10^{\frac{1881-1650}{400}}} = 0.209$$

Further, if team A won the match, its Elo rating will increase to a rating according to equation B.3:

$$R_A^{'} = 1881 + (1 - 0.791) \times 20 = 1885$$

As for team B, its Elo rating decreases to

$$R_B^{'} = 1650 + (0 - 0.209) \times 20 = 1646$$

In the following, the Elo ratings used for 2017/2018 season is presented in form of tables.

Team	GW1	GW2	GW3	GW4	GW5	GW6	GW7	GW8	GW9
Man City	1864	1873	1870	1874	1900	1906	1922	1932	1947
Liverpool	1835	1842	1851	1862	1851	1842	1854	1850	1863
Tottenham	1885	1894	1886	1879	1903	1892	1912	1915	1929
Man Utd	1855	1864	1876	1880	1884	1888	1914	1917	1929
Chelsea	1907	1895	1906	1911	1923	1916	1944	1934	1921
Arsenal	1845	1853	1843	1832	1841	1845	1860	1863	1858
Leicester	1714	1715	1721	1717	1717	1711	1713	1713	1716
Burnley	1626	1646	1638	1645	1656	1659	1663	1674	1679
Everton	1749	1760	1768	1763	1756	1735	1740	1728	1733
Newcastle	1621	1619	1612	1625	1641	1645	1646	1650	1659
Bournemouth	1651	1649	1635	1631	1632	1635	1639	1640	1644
Crystal Palace	1640	1620	1620	1607	1607	1596	1601	1599	1619
Swansea	1647	1654	1646	1659	1653	1658	1656	1650	1661
West Ham	1670	1668	1664	1651	1661	1660	1665	1671	1679
Southampton	1689	1690	1697	1694	1683	1688	1691	1685	1688
Watford	1601	1609	1627	1623	1645	1633	1652	1654	1672
Brighton	1581	1580	1577	1581	1598	1589	1606	1603	1611
Stoke	1660	1659	1672	1674	1683	1673	1674	1680	1683
West Brom	1642	1653	1664	1662	1655	1651	1656	1654	1664
Huddersfield	1471	1499	1510	1513	1514	1514	1529	1525	1611

Table B.1: Elo ratings for 2017/2018 season.

Team	GW10	GW11	GW12	GW13	GW14	GW15	GW16	GW17	GW18
Man City	1947	1963	1968	1974	1964	1966	1959	1968	1972
Liverpool	1853	1857	1865	1869	1859	1866	1883	1877	1870
Tottenham	1933	1938	1939	1932	1914	1902	1903	1907	1909
Man Utd	1910	1920	1913	1903	1895	1900	1918	1909	1911
Chelsea	1921	1912	1919	1929	1919	1920	1923	1909	1912
Arsenal	1867	1860	1854	1866	1848	1851	1844	1841	1837
Leicester	1721	1729	1730	1724	1713	1724	1733	1740	1754
Burnley	1673	1679	1689	1696	1680	1689	1686	1682	1698
Everton	1708	1691	1697	1695	1653	1665	1684	1690	1698
Newcastle	1662	1655	1646	1642	1616	1618	1619	1612	1605
Bournemouth	1652	1649	1658	1666	1656	1647	1649	1650	1648
Crystal Palace	1611	1610	1609	1609	1607	1609	1614	1614	1621
Swansea	1650	1647	1637	1629	1618	1616	1614	1621	1616
West Ham	1656	1656	1649	1639	1629	1618	1620	1634	1638
Southampton	1691	1691	1681	1675	1676	1675	1679	1682	1667
Watford	1667	1657	1651	1660	1665	1659	1667	1661	1654
Brighton	1628	1628	1638	1637	1624	1622	1618	1605	1603
Stoke	1668	1678	1678	1678	1659	1651	1660	1656	1650
West Brom	1655	1651	1642	1635	1632	1630	1631	1624	1631
Huddersfield	1538	1536	1546	1537	1525	1523	1521	1534	1531

Table B.2: Elo ratings for 2017/2018 season.

Team	GW19	GW20	GW21	GW22	GW23	GW24	GW25	GW26	GW27
Man City	1981	1983	1985	1978	1980	1971	1972	1974	1968
Liverpool	1878	1878	1882	1885	1890	1899	1885	1889	1888
Tottenham	1901	1909	1912	1912	1909	1914	1908	1918	1919
Man Utd	1914	1911	1904	1897	1904	1907	1911	1901	1903
Chelsea	1914	1910	1912	1916	1915	1909	1915	1889	1866
Arsenal	1839	1839	1844	1839	1840	1828	1834	1817	1824
Leicester	1732	1736	1727	1723	1729	1735	1741	1732	1729
Burnley	1697	1689	1696	1693	1689	1681	1677	1677	1683
Everton	1704	1708	1708	1699	1693	1688	1685	1693	1685
Newcastle	1603	1614	1612	1610	1620	1619	1617	1617	1620
Bournemouth	1640	1638	1636	1644	1645	1657	1659	1684	1690
Crystal Palace	1643	1644	1639	1646	1656	1663	1658	1660	1657
Swansea	1610	1609	1605	1615	1611	1613	1627	1644	1648
West Ham	1654	1643	1645	1645	1658	1670	1668	1666	1654
Southampton	1665	1661	1657	1664	1654	1655	1660	1657	1666
Watford	1632	1625	1634	1624	1622	1622	1616	1618	1641
Brighton	1604	1611	1609	1611	1610	1601	1595	1599	1611
Stoke	1634	1643	1642	1638	1628	1625	1632	1630	1624
West Brom	1628	1619	1620	1624	1618	1627	1631	1628	1620
Huddersfield	1553	1557	1558	1560	1555	1543	1536	1533	1531

Table B.3: Elo ratings for 2017/2018 season.

Team	GW28	GW29	GW30	GW31	GW32	GW33	GW34	GW35
Man City	1983	1993	1993	1989	1989	1979	1958	1971
Liverpool	1913	1917	1916	1902	1906	1931	1935	1938
Tottenham	1935	1938	1923	1922	1922	1936	1939	1920
Man Utd	1896	1902	1904	1886	1886	1893	1905	1894
Chelsea	1876	1870	1864	1853	1853	1846	1840	1851
Arsenal	1828	1819	1816	1818	1818	1837	1841	1830
Leicester	1732	1727	1722	1727	1727	1740	1729	1718
Burnley	1682	1680	1685	1693	1693	1703	1712	1712
Everton	1700	1692	1683	1685	1693	1692	1697	1698
Newcastle	1641	1643	1639	1645	1645	1653	1665	1676
Bournemouth	1678	1676	1676	1664	1668	1673	1671	1662
Crystal Palace	1656	1653	1647	1638	1648	1648	1651	1657
Swansea	1662	1649	1659	1651	1651	1651	1651	1651
West Ham	1670	1667	1652	1631	1631	1646	1653	1650
Southampton	1667	1669	1665	1645	1645	1637	1635	1634
Watford	1639	1647	1651	1639	1635	1638	1630	1622
Brighton	1620	1633	1642	1627	1627	1621	1618	1618
Stoke	1629	1634	1634	1625	1617	1616	1614	1616
West Brom	1623	1612	1604	1586	1581	1577	1578	1594
Huddersfield	1557	1568	1565	1559	1548	1547	1551	1560

Table B.4: Elo ratings for 2017/2018 season.

Appendix C

Fantasy Premier League 2017/2018 Season Rankings

In this appendix, Table C.1 provides information about the performance that was required to obtain different ranks in Fantasy Premier League after gameweek 35 of the 2017/2018 season had been. The information is retrieved from FPL's homepage, https://fantasy.premierleague.com.

Overall	Total Points	Ranking (%)	Mean
1	2329	-	66.54
2	2311	-	66.03
3	2297	-	65.63
4	2296	-	65.60
5	2294	-	65.54
6	2293	-	65.51
7	2292	-	65.49
8	2288	-	65.37
9	2287	-	65.34
500	2197	0.01%	62.77
1 000	2181	0.02%	62.31
73 529	2034	1.25%	58.11
122 729	2007	2.08%	57.34
190 386	1982	3.23%	56.63
289 729	1955	4.92%	55.86
307 269	1951	5.21%	55.74
615 009	1897	10.44%	54.20
895 898	1860	15.20%	53.14
1 228 630	1822	20.85%	52.06
1 515 599	1790	25.72%	51.14
1 791 435	1761	30.40%	50.31
2 094 985	1728	35.55%	49.37
2 419 676	1691	41.06%	48.31
2 678 000	1660	45.45%	47.43
2 999 773	1621	50.91%	46.31
3 156 602	1602	53.57%	45.77
3 526 620	1556	59.85%	44.46
4 011 099	1492	68.07%	42.63

Table C.1: Fantasy Premier League 2017/2018 season rankings after gameweek 35 was played.