Impact of Penalties on Score Differentials and Drive Outcomes in the NFL

by

Gurashish Bagga

B.Sc., University of Texas at Arlington, 2020

Project Submitted in Partial Fulfillment of the
Requirements for the Degree of
Master of Science

in the
Department of Statistics and Actuarial Science
Faculty of Science

SIMON FRASER UNIVERSITY
Fall 2023

Declaration of Committee

Name:	**Gurashish Bagga**
Degree:	**Master of Science**
Thesis title:	**Impact of Penalties on Score Differentials and Drive Outcomes in the NFL**
Committee:	**Chair:** Yi Lu Professor, Statistics and Actuarial Science **X. Joan Hu** Supervisor Professor, Statistics and Actuarial Science **Tim Swartz** Committee Member Professor, Statistics and Actuarial Science **Haolun Shi** Examiner Assistant Professor, Statistics and Actuarial Science

Abstract

This project investigates the influence of offensive and defensive penalties on score differentials and drive outcomes in the NFL, incorporating various factors that could affect game dynamics. Employing linear regression, we initially look into the specific impacts of penalties on score differentials. Building upon this analysis, a linear regression model with random intercepts for teams and seasons was employed to further refine our understanding of these interactions. The study then delves into drive outcomes, utilizing logistic regression to examine the distinct effects of penalties and predict drive success. Additionally, a Random Forest Algorithm is implemented for the same purpose, allowing for a comparative assessment of the predictive capabilities of logistic regression and the Random Forest method. Through this comprehensive approach, the project indentifies effective methods for predicting drive outcomes, shedding light on the intricate dynamics of penalties in NFL games and providing insights for both fans and analysts.

Keywords: Linear Regression; Logistic Regression; Random Forests; Sports Analytics; Offensive/Defensive Penalties; Prediction

Acknowledgements

I would like to express my immense gratitude to my supervisor, Dr. Joan Hu. Thank you for your constant support and encouragement, putting up with me and invaluable feedback, and general advice. Your unwavering support and astute advice have been instrumental in shaping my growth as a scholar.

I am grateful to Dr. Tim Swartz for his valuable suggestions and for generously sharing his expertise in sports analytics and serving on my committee. A sincere thank you to Dr. Yi Lu, and Dr. Haolun Shi for serving on my committee. I would also like to thank all the professors I met over my time here for making this experience memorable.

To my parents and my sister, whose boundless love and unwavering belief in my potential have been the bedrock of my resilience and determination; this accomplishment is as much yours as it is mine. Your sacrifices and encouragement have been the silent forces propelling me forward, and for that, my gratitude knows no bounds. Brendan - thanks for introducing me to sports analytics and to tidyverse. Additionally, I am thankful for your movie recommendations which have been as delightful as your culinary choices have been adventurous. Kim - I cherish our beer-fueled conversations about everything from the divine to daily gossip, and I'm grateful for the joy of your company and the patience of your listening, even as I rambled on. Eli - thank you for being such a source of entertainment and unique in your ways. Mandy - thank you for showing me the potential of peanut butter as peculiar as it maybe. Renny - thank you for your insights on topics beyond modeling and the chats about world problems that we cannot solve. Robyn - Your readiness to chauffeur us through Vancouver has been the key to some of my fondest memories, just as your welcoming spirit has made them all the more special. Alice, and Lucas - thanks for sharing your Vancouver expertise and different views on things.

Finally, thanks to Nick, Tim, Ali and so many more friends — for making life more interesting along the way.

Table of Contents

List of Tables

List of Figures

Chapter 1

Introduction

Sports analytics have become a cornerstone not only in the realm of athletic performance but also in the broader business and operational aspects of sports organizations. The strategic application of data analytics extends to ticket sales, an area that has historically relied on less sophisticated methods. In the era of digital ticketing, analytics offers a wealth of opportunities to enhance the customer experience. The transition from paper tickets to digital platforms has revolutionized the way organizations identify and understand their customers, allowing for a more personalized approach to marketing and sales [14]. The growth of the sports market has not been uniform across all revenue streams, with ticket revenues lagging behind in some instances. However, with the adoption of analytics, sports entities now have the capacity to glean insights from data, leading to an optimization of pricing strategies, improved fan engagement, and ultimately, an increase in overall revenue streams [22]. A study published in the "International Journal of Sports Marketing and Sponsorship" explored the factors influencing spectators' adoption of digital ticketing [12].

Beyond sales, analytics profoundly impacts the scouting and hiring processes. It provides a data-driven approach to evaluate player performance potential, identify undervalued athletes, and optimize recruitment strategies. Similarly, when hiring coaching staff, analytics can be used to analyze historical performance data, thereby informing decisions that align with the team's strategic objectives. This analytical approach ensures that investments in talent are made with a higher degree of confidence, grounded in empirical evidence rather than intuition alone.

These transformative effects of sports analytics underscore the shift towards a more data-centric approach in sports management. Teams and organizations equipped with robust analytics capabilities are better positioned to make informed decisions, whether in the context of enhancing team performance, engaging with fans, or driving business outcomes. The integration of analytics in sports is not just a trend but a fundamental change in how sports entities operate and compete in a data-driven era.

1.1 The Evolution of Data-Driven Decision-Making in Sports

In recent years, the field of sports analytics has seen remarkable growth, profoundly changing the way teams and individuals approach the game. This rise in data-driven decision-making has been fueled by the belief that statistical analysis can significantly enhance the chances of victory, a notion supported by the work of leading institutions such as MIT Sloan and prominent platforms like Forbes and the National Football League (NFL).

Venture capital has also taken notice, pouring investments into the sports tech market, indicative of the field's growing importance and potential for impact. Since 2014, thousands of global deals and funding rounds have been initiated, signaling a strong belief in the transformative power of analytics in sports [8].

Artificial Intelligence (AI) and Machine Learning (ML) have been at the forefront of this transformation. Over the last two decades, these technologies have not just altered but fundamentally reinvented the ways in which sports are both consumed and analyzed. The deployment of AI and ML in sports has paved the way for enhanced decision-making and forecasting, drawing significant attention from academics, industry professionals, and policymakers alike [4].

An example of the practical application of sports analytics can be seen in cycling. Data analytics has been leveraged to optimize training by analyzing physiological metrics, environmental conditions, and race histories. This has led to the creation of highly tailored training routines and nutrition plans, which are pushing the limits of human performance and endurance [17].

The integration of analytics into sports has, therefore, not only changed the way athletes and managers view their sports but has also opened up new horizons for performance, strategy, and player development. As these technologies continue to evolve, their impact on sports is expected to deepen, expanding the boundaries of what athletes can achieve and offering teams a competitive edge that is grounded in empirical evidence and strategic foresight.

Moreover, the utilization of analytics has transcended beyond mere performance enhancement to influence the economic and managerial aspects of sports. Sports franchises and athletic organizations are now employing sophisticated models to assess player value, optimize ticket pricing, and enhance fan engagement. These analytical models have also played a pivotal role in injury prevention and career longevity, by enabling teams to make informed decisions regarding player health and game strategies. The infusion of analytics into sports has undeniably established a new paradigm, one where data is as valuable as physical prowess on the field.

The impact of sports analytics on performance extends across various sporting disciplines, transforming how players, teams, and coaches strategize and execute their play. Taking a decision-making lens, analytics have become crucial in the decisions made by general managers and coaches in major team sports. This evolution in strategy is underpinned by the rigorous analysis of both quantitative and qualitative data, promoting more informed decisions and challenging conventional wisdom in sports [6]. The seminal work "Moneyball" highlighted the potential value of data analytics in sports, particularly in baseball, where Sabermetrics provided a framework for team assembly and in-game strategy based on rigorous data analysis. This approach has since influenced numerous sports, leading to a broader adoption of data analytics for decision-making in areas such as player evaluation, game strategy, health, and fitness. Moreover, the disciplined application of analytics in player evaluation and team building has shown that teams with lower payrolls can achieve sustained success, challenging the traditional correlation between high payrolls and winning performance as seen in examples like the Oakland Athletics and Tampa Bay Rays in Major League Baseball [6].

In hockey, the application of analytics has significantly influenced the game, an example is understanding how external factors like crowd presence affect penalties. A study conducted during the COVID-19 pandemic, when National Hockey League games were played without spectators, provided a unique opportunity to analyze referee decision-making. The research revealed that referees tended to award more penalties to the away team compared to the home team when crowds were present. However, in the absence of spectators, the disparity in penalties between home and away teams diminished. This finding suggests that social pressure from fans impacts referees' decisions, contributing to the home advantage phenomenon in professional hockey [7]. A similar study was conducted for penalties in the NFL [5].

Soccer, too, has seen a significant impact from the integration of analytics. The sport has shifted from intuition-based decision-making to an evidence-based approach underpinned by statistical analysis. Data analytics in soccer has grown from simple goal counts to advanced predictive models that inform strategic decisions and player evaluations. The introduction of GPS tracking, video analysis, and wearable tech has been instrumental in collecting complex data, which, in turn, is used to simulate games and predict future performance, providing a competitive edge to teams. Today, data analytics is not just a tool but a critical component in soccer, used by clubs worldwide to improve player performance and develop strategies, with artificial intelligence and machine learning poised to push the boundaries even further [20].

These examples are illustrative of the broader trend across sports where analytics is becoming increasingly central to the decision-making process, affecting not just player and team performance but also aspects like fan engagement and media coverage. The future of sports is one where every decision, on and off the field, is informed by a wealth of data, with the potential to revolutionize how games are played, teams are managed, and talent is developed.

1.2 Penalties and Analytics in NFL

Penalties in the NFL are infractions of the rules, punishable by yardage losses, automatic first downs, or even reversal of plays, which can alter the course of a game. Their purpose is to ensure fair play, safety, and sportsmanship among competitors. However, penalties carry a dual nature; while primarily seen as setbacks, they can also be wielded as strategic instruments. For instance, a defense may take a deliberate offside penalty to stop the clock or a savvy offense might induce a penalty to gain advantageous positioning. The strategic employment of penalties can thus be as critical as any play call, with teams often training to both avoid unnecessary penalties and to leverage them to their benefit when the opportunity arises. This dual nature turns penalties into a chess match within the broader contest, where understanding and manipulating the rules can be as impactful as executing a game-winning touchdown.

The strategic nature of penalties in the NFL is exemplified by the decision-making process surrounding whether to accept or decline a penalty. For instance, during a game, if Team A executes a successful play but Team B commits a penalty, Team A has the option to decline the penalty to preserve the advantageous result of the play, such as retaining a touchdown or a significant yardage gain[3]. This decision-making can significantly influence the scoring opportunities and overall strategy, as effectively managing penalties allows a team to maximize their chances of scoring while limiting their opponent's opportunities. The wisdom in such choices is underpinned by a keen understanding of the game's situation and the potential consequences of each penalty, demonstrating the strategic depth and the dual nature of penalties in the game of football.

In the realm of NFL game outcomes, penalties have emerged as a significant factor with a measurable impact on the probability of winning. Analytical models, such as the one developed by Pelechrinis and Papalexakis[15], which boasts an 84% cross-validation accuracy using box score statistics, underscore the strategic implications of penalties. The model's strength was its simplicity and transparency, enabling it to perform comparably with sophisticated systems like ESPN's FPI and Microsoft's Cortana[15]. A comprehensive examination of over 64,000 penalties called between 1999 and 2020 reveals that while penalties like False Starts and Defensive Offsides can be stable and reflective of team discipline, penalties overall present as an unstable metric for predicting team performance. This instability is echoed in the fluctuating nature of Expected Points Added (EPA) when accounting for nullified plays due to penalties [9]. Statistically, teams with fewer than 40 yards of penalties tend to have a higher win percentage, underscoring the importance of disciplined play and control over game flow. These findings highlight the multifaceted impact of penalties, not just on the scoreboard, but as a strategic element that shapes the ebb and flow of NFL games.

Complementing this, a more recent endeavor published by Dartmouth Sports Analytics developed a predictive model incorporating four key variables: yards gained, yards allowed, turnovers lost, and turnovers recovered. This model, featuring a pseudo R^2 value of 0.4150, takes into account turnover differentials, a crucial aspect of football analytics. It was rigorously tested against the 2021 NFL season's games, accurately predicting outcomes with impressive precision, correctly forecasting 132 out of 150 games[13].

These studies underline the intricate nature of modeling in sports analytics, where even the most seemingly straightforward statistics can yield profound insights into game outcomes. The predictive power of such models is not only a testament to the quantitative rigor but also to the nuanced understanding of the game's strategic elements. The convergence of statistical analysis and football acumen has indeed enhanced the ability to forecast game results with a high degree of confidence, underscoring the evolving synergy between data science and sports.

The objective of this project is to investigate the extent to which offensive and defensive penalties influence the score differential in NFL games. This inquiry is rooted in the persistent debate among football enthusiasts and analysts regarding the comparative impact of these penalties on the game's outcome. The purpose of studying the impact of offensive and defensive penalties on score differentials using regression models with fixed effects and random intercepts in our project is to quantitatively understand how various types of penalties influence the outcome of NFL games. By employing these statistical methods, you can isolate the effects of specific penalties, on the score differential. Regression models enable the examination of relationships between variables, while models witj random intercepts account for the nested structure of the data (e.g. games within seasons, teams within games). This approach allows for a more nuanced analysis, considering both fixed effects (like the type of penalty) and random effects (such as variability across games or seasons). This study aims to provide insights into the impact of penalties in football, offering potentially valuable information for teams, coaches, and analysts regarding game strategy and rule enforcement's impact on game outcomes.

1.3 Organization of the Project

This project systematically investigates the influence of penalties on score differentials using aggregated game-level data. Chapter 2 presents the data alongside a descriptive analysis, setting the foundation for the study. Chapter 3 delves into the models employed, with a particular focus on interpreting key coefficients. Chapter 4 extends the discussion to a more detailed drive-level data analysis, enhancing the granularity of our statistical approach. In Chapter 5, we synthesize our findings and discuss the model's limitations. The ultimate goal of this research is to discern the extent to which penalties can affect score differentials and drive outcomes. Despite the complexities inherent in modeling score differentials without highly detailed tracking data, this project aims to shed some light on the implications of penalties from various analytical perspectives.

Chapter 2

Overview of the Data

2.1 nflfastR Package

We used a well known R package nflfastR [1] to access the data required for this project. The nflfastR package was officially introduced to the public through a tweet by Ben Baldwin on April 27, 2020. This package, designed to scrape NFL play-by-play data, was created to allow end-users to access this information more swiftly than previous methods. Ben Baldwin, along with contributors such as Tan Ho, Sebastian Carl, and others within the nflverse community, played a pivotal role in developing the nflfastR package. It is part of the broader nflverse, a collection of data and R-based packages enabling in-depth access to NFL data dating back to 1999. The nflfastR project has its roots in the nflscrapR project, initiated by Maksim Horowitz and Sam Ventura of Carnegie Mellon University, and continued by Ron Yurko. The transition from nflscrapR to nflfastR was marked by a tweet from Yurko on September 14, 2020, signifying the beginning of a new era in NFL data analytics. The nflfast R package led to peer-revied work by Yurko et al titled "nflWAR: A Reproducible Method for Offensive Player Evaluation in Football"[21].

The nflfastR dataset serves as a comprehensive repository for NFL play-by-play data, providing an extensive array of variables pertinent to games, players, and plays. Originating from the nflfastR project, which aims to facilitate the analysis of NFL data, this dataset encompasses a wide demographic of professional American football teams and players spanning multiple seasons. The dataset's inclusivity of various teams and players across the league ensures a representative sample for robust statistical analysis. The nflfastR project is renowned for its meticulous data compilation and has been extensively utilized within sports analytics research, thereby offering a reliable foundation for rigorous academic investigation. We will using data from 2011-2022 season for our analysis.

The nflfastR dataset encompasses a broad spectrum of variables ranging from basic game information to intricate play-level details. This includes but is not limited to, play outcomes, player statistics, game contexts (e.g., down and distance), rush and pass statistics and advanced metrics such as expected points added (EPA) and win probability. The dataset's granularity enables the construction of multifaceted statistical models capable of evaluating the impact of penalties on game outcomes. The ability to parse data from 2011 to 2022 allows for a longitudinal analysis, affording insights into trends and changes over time, if any. The depth and breadth of the dataset ensure that any statistical analysis conducted can account for a multitude of factors that might influence game dynamics.

The utilization of the nflfastR dataset is driven by the need to understand the statistical impact of penalties on the score differentials within NFL games. Penalties can serve as pivotal moments in football, with the potential to significantly alter the course of a game. By analyzing penalty data in conjunction with score differentials, the study aims to quantify the extent to which penalties contribute to game outcomes, providing valuable insights for teams, coaches and the viewers. nflffastR is an open-source package and and can be accessed directly in R[1].

2.2 Descriptive Analysis

2.2.1 Aggregated Game Level Data

This dataset encompasses a comprehensive collection of football game statistics from the 2011 to 2022 seasons, totaling 3,241 games. Within this collection, 138 games represent postseason play, while the remainder constitutes regular season matchups. A significant volume of penalty data is included, with 40,792 penalties recorded in total. Analysis of these penalties reveals a nearly balanced distribution between defensive penalties, which account for 19,097 instances (47%), and offensive penalties, totaling 21,695 (53%). Offenive penalties are defined as the penalties made by the team in possession of the ball and defenisve penalties are defined as the penalties made by the team not in possession of the ball.

A notable transition occurred in the 2021 season when the league amended its regular season structure, increasing the number of games from 16 to 17 per team. This modification is mentioned for its potential relevance, although its impact on the overall analysis is mitigated by the fact that only two seasons within the dataset, specifically 2021 and 2022, operate under this new regime.

Additionally, the dataset includes 197 games (6% of the total) that extended into overtime. While the dynamics of a game may shift in overtime play, for the purposes of this analysis, such games are treated equivalently to non-overtime games, with the final score differential and total penalty count at the conclusion of the game being the primary metrics of study.

The following table presents a yearly breakdown of the penalties assessed

Season	Total Penalties	Offensive Penalties	Defensive Penalties
2011	3375	1867	1508
2012	3318	1750	1568
2013	3241	1669	1572
2014	3522	1843	1679
2015	3663	1927	1736
2016	3545	1905	1640
2017	3520	1767	1753
2018	3565	1891	1674
2019	3568	1921	1647
2020	2986	1486	1500
2021	3336	1873	1463
2022	3153	1796	1357

Table 2.1: Number of Penalties observed in each season over 2011-2022

An intriguing anomaly presented in the dataset is the discernible decline in penalty frequency during the year 2020. This trend may be attributable to the unique circumstances of the 2020 season, wherein contests were conducted with either greatly reduced or entirely absent stadium audiences—a measure necessitated by the global health crisis. It posits the hypothesis that crowd dynamics potentially exert a tangible influence on the adjudication of penalties, a factor that was significantly mitigated during this period. The impact of crowds on NFL penalties has been a subject of research, particularly in the context of the COVID-19 pandemic when many games were played without the usual fan presence. According to the paper "False Start? An Analysis of NFL Penalties With and Without Crowds,"[5] researchers capitalized on the unique circumstances to examine the effect of crowd absence on the home advantage and game outcomes. Drawing on play-by-play data from NFL games, the paper investigates whether quarterbacks could better manipulate opposition defenses in the absence of crowds. The findings suggest that the lack of crowd noise may indeed give quarterbacks an advantage in orchestrating the game, potentially due to clearer communication and less pressure from the external environment. Also, each year there are consistently more offensive penalties than defensive ones.

The evaluation of penalties within a game can be quantified through three primary metrics. The initial measure is the straightforward enumeration of penalties incurred. The second metric, penalty yardage, is deemed more informative for several reasons. It not only reflects the severity of an infraction—highlighting the disparity between penalties of 5 and 15 yards—but also incorporates instances where penalties may be declined, an element not directly captured by a mere count. However, penalty yardage alone does not consider the contextual gravity of the penalty. For example, a 5-yard infraction against the defense assumes greater significance if the offense is at the 15-yard line compared to the 50-yard line. Thus, the third metric employed is Expected Points Added (EPA), which integrates situational variables such as down, distance to first down, and field position, estimating the expected point impact of a specific play. EPA is a statistical metric widely used in the analysis of football, including the NFL. EPA quantifies the impact of each play by measuring the change in expected point outcome before and after a specific play occurs. The concept revolves around the idea that the value of a down-and-distance situation on the field can be expressed in terms of expected points, representing the likely point outcome for the possessing team. For instance, if a team gains significant yardage on a play, their expected points for that drive increase. Conversely, if they lose yards or face an unsuccessful play, the expected points decrease. EPA provides a nuanced evaluation of player and team performance, allowing analysts and coaches to assess the effectiveness of plays in terms of their impact on scoring potential. Positive EPA values indicate plays that contribute favorably to a team's scoring chances, while negative values signify plays that hinder those chances. This metric has become a valuable tool for understanding the strategic and statistical dynamics of football games, helping teams optimize their decision-making processes for enhanced on-field performance.

Typically, the EPA for defensive penalties is positive, suggesting an advancement and tactical benefit. Conversely, the EPA for offensive penalties is generally negative, reflecting lost yardage or a retreat down the field. Although penalties are not a direct indicator of game outcomes, the subsequent analysis will present a scatter plot correlating the differentials of these three metrics with the score differential, specifically home minus away scores, while distinguishing between offensive and defensive penalties.

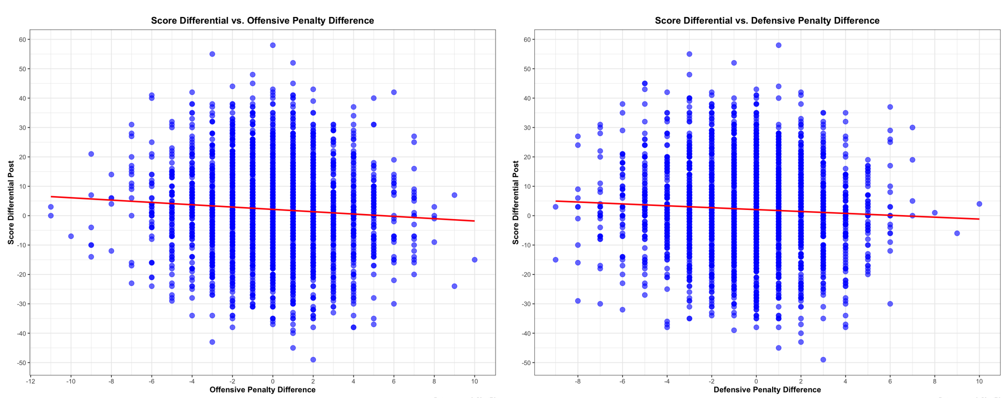

Figure 2.1: Left: Scatter plot of offensive penalty count and score differentials. Right: Scatter plot of defensive penalty count and score differentials.

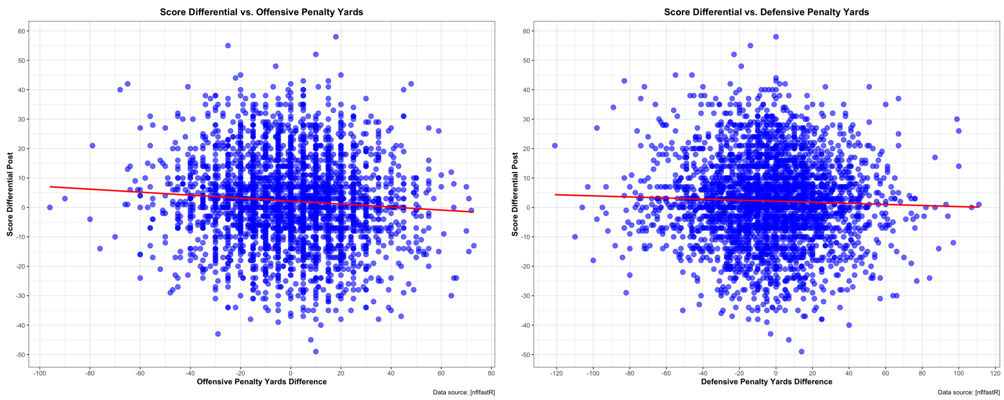

Figure 2.2: Left: Scatter plot of offensive penalty yards and score differentials. Right: Scatter plot of defensive penalty yards and score differentials.

Figure 2.3: Left: Scatter plot of offensive penalty EPA and score differentials. Right: Scatter plot of defensive penalty EPA and score differentials.

The lines in the figures above represent the regression fit for each penalty metric. This helps in observing any patterns in a more succint manner. The graphical representation in Figure 2.3 elucidate that Expected Points Added (EPA) exhibits a more pronounced correlation with score differential in comparison to the other two metrics in Figure 2.1 and 2.2. This observation is congruent with anticipations, given that EPA incorporates situational factors—a dimension of the game that is essential for a comprehensive analysis. The importance of these situational factors is further substantiated by Figure 2.4. It reveals a higher incidence of penalties during the third and fourth downs. This pattern may be indicative of an escalated level of aggression in play during these critical junctures, where offensive teams are striving for either a score or a renewal of downs, and defensive teams are equally determined to thwart their advances.

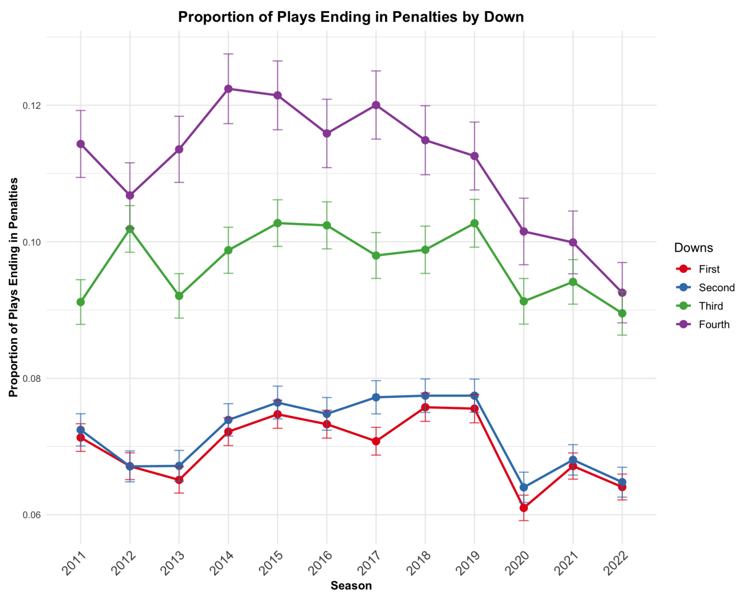

Figure 2.4: Proportion of plays ending up in penalty for each down

Figure 2.5 elucidates the proportion of plays culminating in penalties within each quarter. To provide a more detailed analysis, this is disaggregated into two distinct visual representations. Figure 2.6 specifically delineates the proportions of penalties categorized as either offensive or defensive across each quarter. We observe in Figure 2.5 that more plays end up in penalties during the second quarter compared to the first one. A proposed hypothesis for this observation might be because the teams try to play it safe and trying to understand each others strategies. On the other hand during the second quarter the first half is about to end the possession of the ball changes in the second half. This often calls for a change in strategies of both teams during the second quarter.

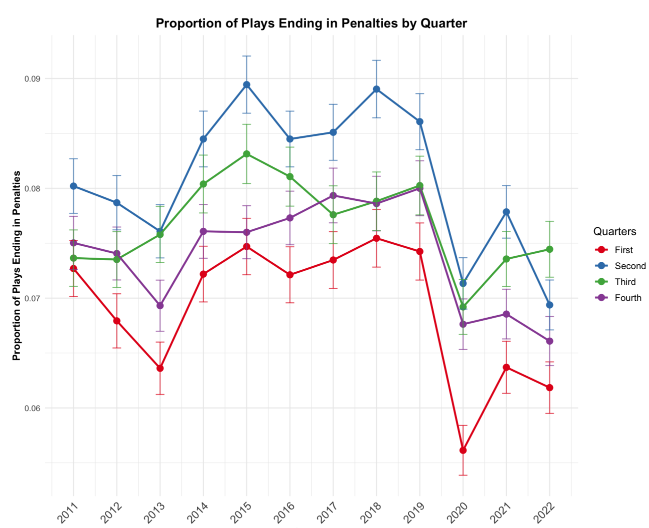

Figure 2.5: Proportion of plays ending up in penalty for each Quarter

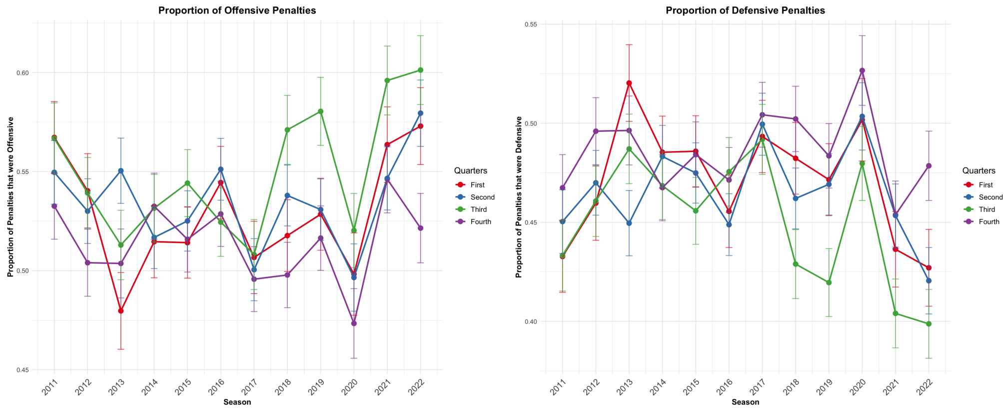

Figure 2.6: Left: Proportion of Penalties that are Offensive. Right: Proportion of Penalties that are Defensive.

Now, we present a meticulously constructed correlation matrix, which encapsulates the interrelationships among key variables critical to our analysis. This matrix serves as a foundational tool for unraveling the intricate dynamics that underpin the game's strategic and operational aspects. By scrutinizing the interdependencies among these variables, we gain valuable insights that help us avoid multi-collinearty issues.

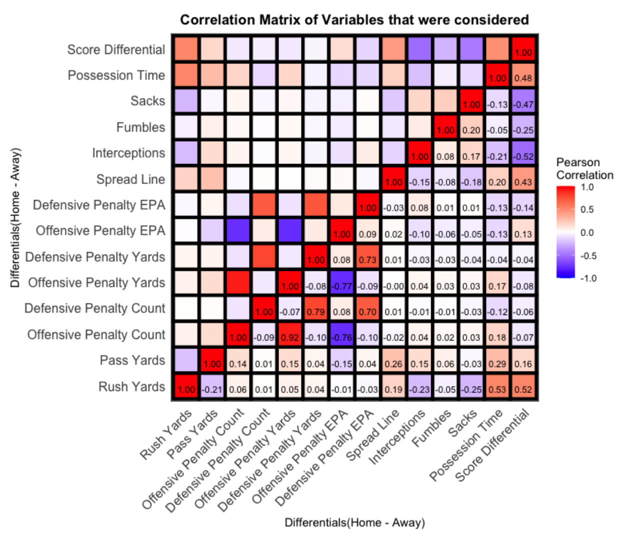

Figure 2.7: Correlation matrix for variables considered

The correlation matrix above is based on the observed data from games. In Figure 2.7, we notice that penalty-related variables exhibit a high degree of correlation, a finding that aligns with intuitive expectations. For instance, an increase in defensive penalties is logically associated with a corresponding rise in penalty yards conceded. This pattern of interrelation among penalty-related metrics is a predictable aspect of the data.

Conversely, certain variables, while only moderately correlated, hold significant relevance for the analysis. A notable example is the negative correlation observed between rush yards and interceptions. This relationship suggests that an increase in rushing plays tends to coincide with a decrease in interceptions. This inverse relationship is evident in the data and is consistent with the understanding that rushing plays reduce the opportunities for passing plays, which are typically more prone to interceptions. However, this correlation is not as apparent at the individual game level, indicating that it emerges more clearly through the aggregation of data across multiple games.

In future model development, these slightly correlated variables will be included, and their impact will be carefully evaluated. To ensure the robustness of the model and to address potential concerns of multicollinearity, the Variance Inflation Factor (VIF) will be calculated. This step is crucial to ascertain that the model's variance is not unduly inflated due to the presence of multicollinearity, thereby preserving the integrity and reliability of the model's findings.

Now, we will present the plot that will inform about the distribution of offensive and defensive penalty EPA.

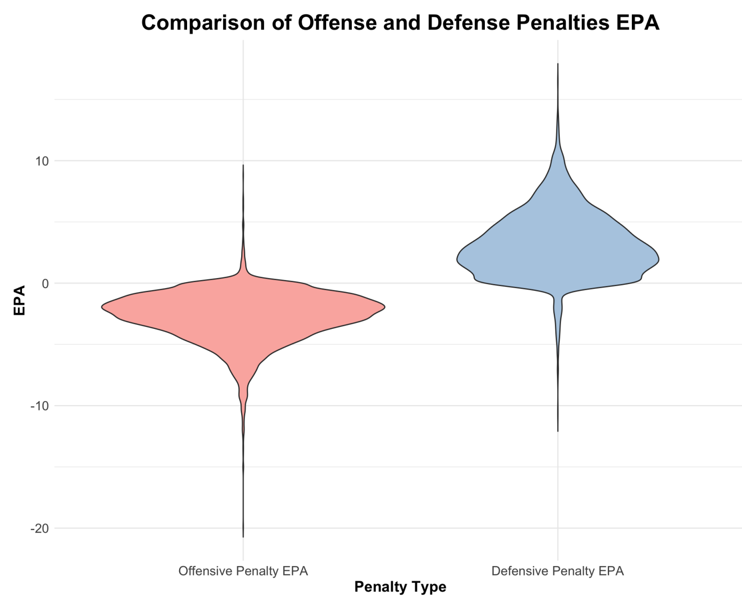

Figure 2.8: Violin plot for offensive and defensive penalty EPA

The analysis of the plot in Figure 2.8 reveals that the Expected Points Added (EPA) for offensive penalties is predominantly negative, a logical outcome considering that penalties against the offense typically result in a loss of expected points due to the incurred disadvantage. In contrast, defensive penalties generally yield a positive EPA, reflecting the advantage conferred to the opposing team.

Notably, there are exceptional instances where an offensive penalty may result in a positive EPA, and conversely, a defensive penalty may lead to a negative EPA. These outlier occurrences, though rare, are integral to the analysis. The inclusion of these data points is essential for a comprehensive understanding, as the analysis is based on aggregated game-level data, which offers a holistic view of the impact of penalties.

Further examination of the data indicates that the distribution tail for defensive penalty EPA is more pronounced. This observation suggests that, on an individual basis, defensive penalties can have a more detrimental impact on defensive teams compared to the effect of offensive penalties on offensive teams. This insight underscores the varying implications of penalties depending on whether they are incurred by the offense or the defense, and it highlights the importance of considering the context and consequences of each penalty type in the analysis.

We will now examine the histogram of the score differential to gain insight into its distribution. This visual analysis will allow us to observe the range and frequency of score differentials, providing a clearer understanding of their variance within the dataset. The distribution of scores in sporting events has been extensively studied, with various statistical models applied to predict outcomes and analyze the dynamics of different games. In soccer, the double Poisson model is one of the most established methods [11]. Initially developed in 1982, it remains a popular choice for predicting football scores despite the advent of numerous other techniques. This model assumes that goals scored by each team are Poisson distributed with means depending on their respective attacking and defensive strengths. The double Poisson model's predictions were highly accurate for the Euro 2020 football tournament, as they won the Royal Statistical Society's prediction competition[16]. NFL scoring does not follow a Poisson process as closely as in those sports since there are various ways to score like touchdowns, field goals and safety.

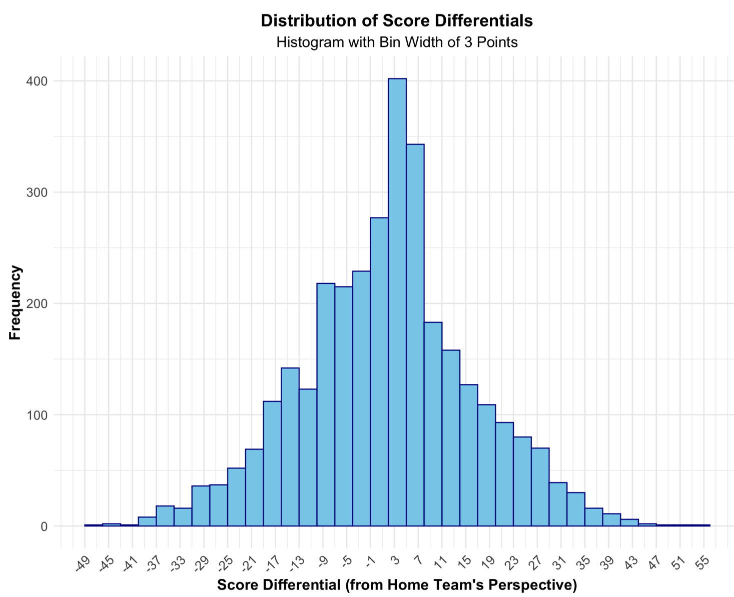

Figure 2.9: Histogram of score differentials

Based on the histogram provided, which displays the frequency of score differentials (home score minus away score) for NFL games, the distribution appears to be roughly symmetric around zero and has a single peak, suggesting that the differentials are almost normally distributed. This is characterized by the “bell curve“ shape of the histogram. In the context of NFL score differentials, where the values can range on both sides of zero (indicating home team wins when positive and away team wins when negative), we can use normal distribution. We must acknowledge that there is room for exploration here since the score differential is not truly a continuous variable.

The normal distribution is a good choice here because we assume that the score differentials are influenced by many small, independent effects, which is a reasonable assumption for NFL games. It is also practically useful because it is well-understood and has many statistical properties that faciliate analysis. However, if you observe that the tails of the distribution are heavier than would be expected in a normal distribution (indicating more extreme outcomes than the normal distribution would predict), you might consider a distribution that can handle such “fat tails“, like a t-distribution too. The tails on both ends suggest that there are fewer instances of very large or very small score differentials, which aligns with the properties of a normal distribution where extreme values are possible but less likely.

In Figure 2.10 we also present the boxplot for score differential by season. We observe that there are no major differences in the distribution of score differential during each season. In Chapter 3 we will still add a random intercept for season in our model, to quantify the dependence(if any) of games within a season.

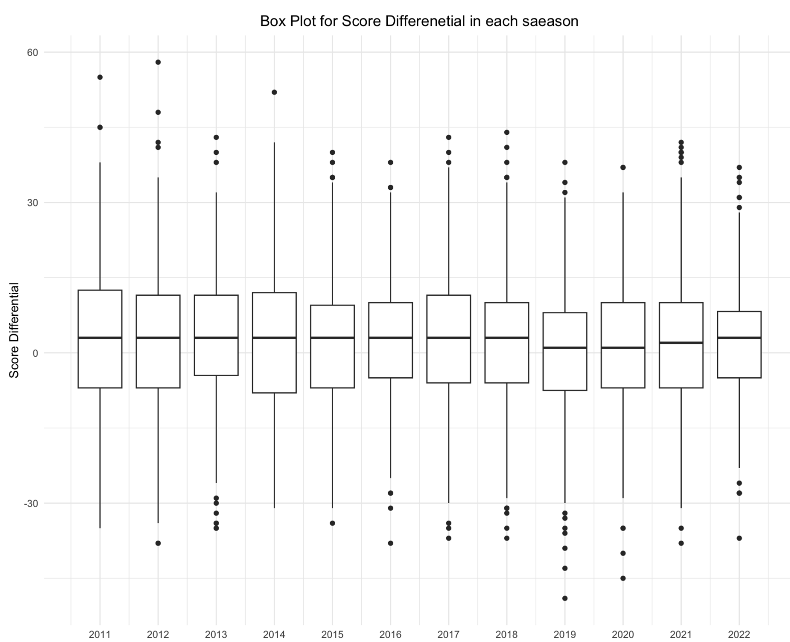

Figure 2.10: BoxPlot for Score Differentials in each season

In our dataset, we identified 64 distinct types of penalties. A detailed distribution of these penalties is provided in Appendix B. Notably, the two most common penalties are of an offensive nature. This suggests that, despite the relatively lower expected point loss associated with offensive penalties, their higher frequency of occurrence contrasts with the less frequent but more impactful defensive penalties in terms of Expected Points Added (EPA).

2.2.2 Drive Level Data

In our dataset, a total of 72,945 unique drives have been identified, following the exclusion of overtime drives. Overtime drives are omitted from this analysis due to potential variations in strategies that are typically employed during overtime periods. Furthermore, we incorporate 'quarter' as one of our covariates. However, it would be prudent to address quarters 5 and 6 separately in future analyses for a more nuanced understanding.

Of these drives, 27,929 (representing 38.2%) involved at least one penalty. This statistic underscores the prevalence of penalties in regular gameplay and their potential impact on game outcomes.

Additionally, 25,907 drives (amounting to 35.5%) concluded with either a field goal or a touchdown. The remaining drives culminated in safeties, turnovers, or opposing team touchdowns. In Figure 2.11, we delve into the distribution of offensive and defensive penalty Expected Points Added (EPA) for drives culminating in touchdowns and field goals, compared to those that did not. This comparison aims to elucidate the influence of penalties on the efficacy of scoring drives.

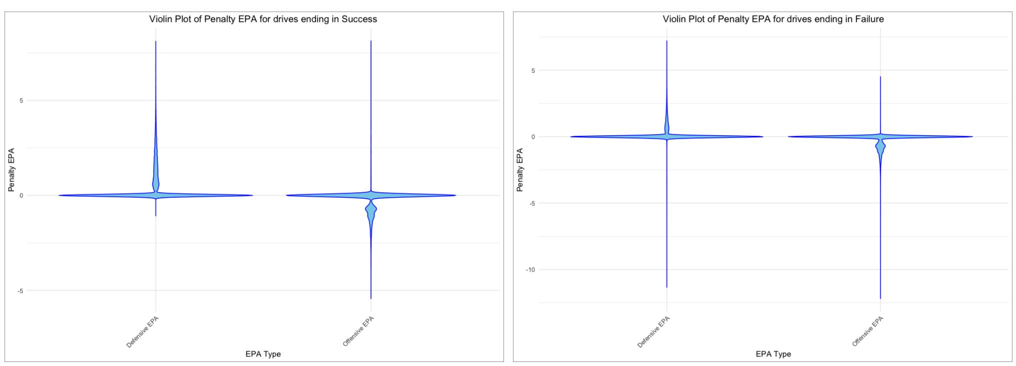

Figure 2.11: Left: Penalty EPA for successful drives. Right: Penalty EPA for unsuccessful drives.

Figure 2.12 presents the drives observed across each of the four quarters. This includes not only the total number of drives per quarter but also a comparative breakdown of their outcomes. Specifically, it distinguishes between drives culminating in touchdowns or field goals versus those resulting in other outcomes. This delineation allows for a detailed examination of drive success rates and patterns within each quarter, contributing to a deeper understanding of game across different stages of the match.

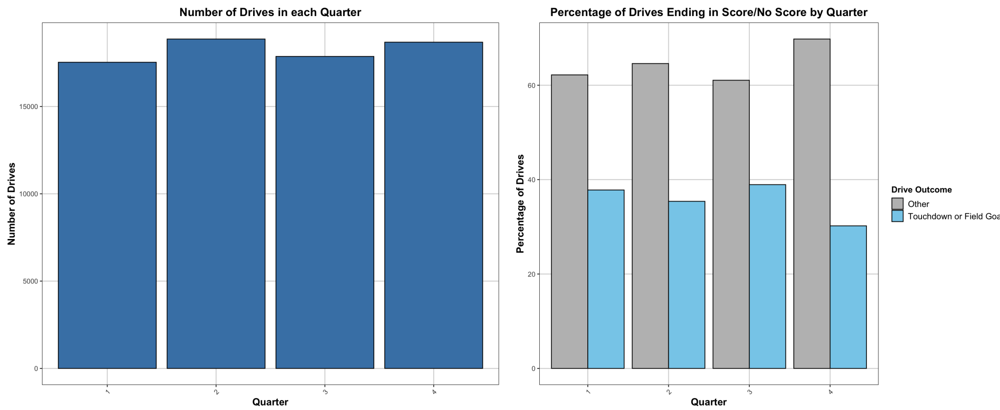

Figure 2.12: Left: Number of Drives in each Quarter. Right: Success rate of Drives.

Observations from the data reveal a notable disparity in the number of drives between different quarters, with the second and fourth quarters exhibiting a higher frequency of drives compared to the first and third. This variation could be attributed to a quicker pace of play during the second and fourth quarters or it might be a consequence of 2-minute warning and an attempt to score last in the two halves. Furthermore, a distinct decrease in the proportion of drives resulting in touchdowns and field goals is evident in the fourth quarter. This trend suggests an increased difficulty in scoring towards the end of the game. Potential explanations for this phenomenon might include a more robust defensive strategy in the latter stages or the impact of fatigue on offensive performance.

2.3 Covariates of Interest

To initiate our analysis, it is imperative to delineate and scrutinize the variables that are pivotal to our study. The focus will be primarily on the differential factors, characterized as the distinctions between home and away conditions. These variables in Table 2.2 are imperative to the aggregated game-level analysis. The Table 2.3 has covariates that are imperative to drive level analysis. All the differentials in Table 2.2 are home minus away.

Table 2.2: Description of variables for aggregated game-level data

Covariate(Differentials)	**Description**
`season`	Season of the game
`Offensive Penalty`	Difference in count of offensive penalties between the teams in a game.
`Defensive Penalty`	Difference in count of defensive penalties between the teams in a game.
`Offensive Penalty Yards`	Difference in penalty yards lost due to offensive penalties between the teams in a game.
`Defensive Penalty Yards`	Difference in penalty yards given up due to defensive penalties between the teams in a game.
`Offensive Penalty EPA`	Difference in expected points added (EPA) due to offensive penalties between the teams in a game.
`Defensive Penalty EPA`	Difference in EPA due to defensive penalties between the teams in a game.
`Spread Line`	The closing spread line for the game. A positive number means the home team was favored by that many points, a negative number means the away team was favored by that many points. (Source: Pro-Football-Reference [2])
`Fumbles`	Differential count of fumbles between the teams.
`Interceptions`	Differential count of interceptions between the teams.
`Passing Yards`	Difference in passing yards between the teams, standardized (scaled) for the analysis.
`Rushing Yards`	Difference in rushing yards between the teams.
`Sacks`	Differential count of sacks between the teams.
`Possession Time`	Differential of possession time between the teams.
`Score Differential`	Final score differential.

Table 2.3: Description of the Variables for the drive-level data

Covariate	**Description**
`Yardline 100`	The distance in yards from the opponent's end zone. A higher value indicates the offensive team is farther from scoring.
`Offensive Penalty EPA`	Expected Points Added (EPA) by the offensive penalties. It measures the impact of offensive penalties on scoring potential.
`Defensive Penalty EPA`	Expected Points Added (EPA) by the defensive penalties. It measures the impact of the defensive penalties on the scoring potential.
`Score Differential`	The point difference between the teams at the beginning of the drive. A positive value indicates the offensive team is leading.
`Drive Inside20`	Indicator of whether the drive reached within 20 yards of the opponent's end zone (red zone).
`ha drive`	Home or Away indicator for the offensive team during the drive.
`Posteam Timeouts Remaining`	The number of timeouts remaining for the offensive team during the drive.
`Drive Time of Possession`	The duration of the drive in terms of possession time.
`Fumble`	Indicator of whether a fumble occurred during the drive.
`Qtr`	The quarter of the game during which the drive took place.
`Interaction 1`	The interaction between offensive penalty epa and possesion time.
`Interaction 2`	The interaction between defensive penalty epa and possesion time.

Chapter 3

Score Differentials: Regression Analysis and Results

In this chapter, we will explore linear regression models, including both fixed effects and models with random intercepts, with a focus on score differential as the dependent variable. The dataset for this analysis has been aggregated at the game level, utilizing home-away differentials as predictive factors. For instance, a predictor might be the net yards gained, calculated as the total yards gained by the home team minus those gained by the away team. Given that the data is aggregated on a game-by-game basis, this approach provides a less granular perspective. However, it remains valuable in discerning whether offensive and defensive penalties exert significantly different impacts on the score differential. This method aims to yield insights into the broader patterns and effects in the realm of game strategy and outcomes. It is important to note that we have opted not to explore a model that includes only teams as a random effect. The rationale behind this decision stems from the dynamic nature of teams in professional sports; teams undergo significant changes each season, including alterations in player rosters and coaching staff. As such, treating games played by a single team across multiple seasons as a homogeneous group may not accurately capture the evolving nature of these teams.

3.1 Linear Regression with Fixed Effects

In this study, we approach the analysis of NFL games through a linear regression model, initially incorporating three predictor variables: offensive penalty count, defensive penalty count, and spreadline. The spreadline, a significant factor in this model, is designed to reflect overall team quality. This metric is typically derived by betting companies using proprietary algorithms, taking into account a myriad of factors including player injuries, the performance capability of quarterbacks, and other relevant variables like home team advantage.

The inclusion of the spreadline variable is crucial. It essentially acts as a comprehensive indicator of pre-game conditions and team dynamics. Given that spreadlines are established before the games and incorporate extensive pre-game information, the standard intercept in the regression model becomes less meaningful. In typical regression models, the intercept represents the average score differential when no information about the game is available. However, with the constant availability of spreadline data prior to games, this intercept loses its significance in our model.

Furthermore, to enhance our model's robustness, we will also explore the integration of two additional penalty-related metrics: penalty EPA (Expected Points Added) and penalty yards. These variables offer a deeper insight into the impact of penalties on game outcomes and may provide valuable predictive power in conjunction with the previously mentioned variables.

Our objective is to construct a model that not only captures the immediate effects of observable variables but also incorporates the nuanced aspects of game dynamics as represented by the spreadline. This approach allows for a more comprehensive understanding of the factors influencing the outcomes of football games.

Table 3.1: Regression Analysis Results for Analyzing Score Differential

	Model 1	*Model 2*	*Model 3*
Variables			
Intercept	$-0.091(0.244)$	$-0.049(0.243)$	$-0.117(0.238)$
Offensive Penalty Count(D)	$-0.377^{***}(0.086)$		
Offensive Penalty Yards(D)		$-0.052^{***}(0.010)$	
Offensive Penalty EPA(D)			$0.697^{***}(0.079)$
Defensive Penalty Count(D)	$-0.378^{***}(0.091)$		
Defensive Penalty Yards(D)		$-0.024^{**}(0.008)$	
Defensive Penalty EPA(D)			$-0.557^{***}(0.061)$
Spread Line	$1.025^{***}(0.038)$	$1.028^{***}(0.038)$	$1.011^{***}(0.037)$
Model Statistics			
Adjusted R^2	0.208	0.211	0.238
Residual Std. Error	12.986 (df = 3226)	12.981 (df = 3226)	12.756 (df = 3226)
F Statistic	284.212^{***} (df = 3)	287.993^{***} (df = 3)	336.389^{***} (df = 3)

Note: $^{*}p<0.1$; $^{**}p<0.05$; $^{***}p<0.01$
The parenthesis besides estimates contain the standard errors.
(D) = The difference of that predictor(Home - Away)
p-values are associated with the null hypothesis that the corresponding coefficients are equal to zero.

In our analysis encompassing three distinct models, we consistently observe that penalty variables emerge as significant predictors for the score differential, with the spreadline variable also demonstrating notable significance. Each model in Table 3.1 uses a different penalty metric, for example Model 1 uses penalty counts, Model 2 uses penalty yards and Model 3 uses penalty EPA. Intriguingly, the spreadline variable consistently exhibits a positive association, which confirms the fact that the spreadline is the prediction of the game outcome from home team viewpoint. If spreadline increases by one unit, the score differential increases by $\beta_{spreadline}$ which is positive. The null hypothesis for all the coefficients is $\beta_i = 0$.

Having established the penalties as significant factors influencing game outcomes, it becomes imperative to account for additional variables that could potentially affect the dynamics of the game. This includes, but is not limited to, defensive metrics such as sacks and interceptions, and offensive metrics like fumbles. These variables represent critical aspects of team performance and strategy, and their inclusion in the model is essential to provide a more comprehensive and accurate understanding of the factors that drive score differentials in football games.

The objective is to refine our model to encompass a wider array of variables, thereby capturing a more nuanced view of the game. By integrating these additional elements, our analysis can more effectively isolate the specific impacts of penalties and other variables on game outcomes, leading to more nuanced and insightful conclusions about the nature of competitive football.

In the initial phase of our project, in Chapter 3 Figure 2.7, we observed a high degree of correlation among the various penalty variables. Based on this finding, our subsequent analysis will primarily utilize the penalty EPA (Expected Points Added) metric. This decision stems from the EPA's inherent ability to account for situational factors and its direct correlation with score differentials, making it a particularly insightful variable for our study.

While the penalty EPA will be our primary focus, we will not entirely disregard penalty yards, which will still be examined and discussed. For those interested in more details of penalty yards, we will include the estimates in the appendix section of our report. This approach ensures that while our primary analysis is streamlined and focused, comprehensive data is still available for those seeking more insights.

In the next iteration of our model, we aim to expand our analytical scope by incorporating additional predictors that influence game outcomes. Specifically, we will include the differentials in interceptions and sacks, which are indicative of defensive strength. Additionally, we will consider the differentials in rushing and passing yards, as well as fumbles, to serve as proxies for the offensive strength of the teams. These additional variables are selected to provide a more rounded and nuanced understanding of the dynamics influencing game outcomes.

The forthcoming Table 3.2 will present the estimates derived from this enhanced model, showcasing the impact of these diverse predictors alongside the spreadline and penalty metrics.

Table 3.2: Regression Analysis Result for Analyzing Score Differential

Variables	Estimates(Standard Error)
Intercept	0.247(0.148)
Spread Line	$0.393^{***}(0.025)$
Offensive Penalty EPA(D)	$0.649^{***}(0.050)$
Defensive Penalty EPA(D)	$-0.452^{***}(0.038)$
Fumbled(D)	$-1.409^{***}(0.089)$
Intercepted(D)	$-3.765^{***}(0.103)$
Sacked(D)	$-1.295^{***}(0.059)$
Rushing Yards(D)	$0.071^{***}(0.002)$
Passing Yards(D)	$0.042^{***}(0.002)$
R^2	0.706
Adjusted R^2	0.706
Residual Std. Error	7.925 (df = 3221)
F Statistic	938.885^{***} (df = 8; 3221)

Note: $^{*}p<0.1$; $^{**}p<0.05$; $^{***}p<0.01$
(D) = The difference of that predictor(Home - Away)
p-values are associated with the null hypothesis that the corresponding coefficients are equal to zero.

In our analysis presented in Table 3.2, we observe that all coefficients in the model are highly significant. Notably, the adjusted R-squared value of this model is substantially higher than those of the previous models, which only included penalty and spreadline predictors. This enhancement in the model's explanatory power underscores the importance of incorporating a broader range of variables to capture the complexity of football game dynamics.

A key observation is the particularly large coefficient associated with interceptions. This aligns well with our understanding of the game, as interceptions often lead to a loss of possession or, in more detrimental cases, result in a defensive touchdown for the opposing team. In comparison, the coefficients for sacks and fumbles are also significant but to a lesser extent. The relatively smaller impact of sacks and fumbles can be rationalized by the fact that, while sacks result in a loss of yards, the team maintains possession of the ball. Similarly, fumbles offer a chance for the team to recover possession, unlike interceptions where possession is definitively lost.

Our primary focus, however, lies in the coefficients for offensive and defensive penalty EPA. We rigorously compared the coefficients for offensive penalty EPA and defensive penalty EPA by conducting an F-test. The null hypothesis for this test was that the sum of these two coefficients is zero. The resulting p-value from this test was less than 0.05, indicating that the difference between the offensive and defensive penalty EPA coefficients is statistically significant. This finding emphasizes the distinct impacts of offensive and defensive penalties on the game's score differential, providing valuable insights into the nuanced effects of penalties within the context of American football.

Interpreting these coefficients, we find that for each unit increase in the EPA differential due to an offensive penalty, the average score differential increases by approximately 0.65, assuming all other variables remain constant. However, it is important to contextualize this finding. Our analysis, informed by Figure 2.8, implies that losing one expected point due to an offensive penalty is generally more challenging than gaining an expected point from defensive penalties. This is because defensive penalties tend to impose more severe penalties on the defense. Consequently, even though the absolute coefficient of the offensive penalty EPA is larger than that of the defensive penalty EPA, this does not necessarily imply that offensive penalties are more detrimental than defensive penalties. This poses a question for the future work on this project about how many penalties on average does it take to lose an expected point by offensive penalties.

Our primary objective remains to discern whether offensive and defensive penalties have equivalent impacts or if a significant difference exists. To further investigate this, we plan to fit one more model, incorporating all covariates but focusing on penalty yards. The estimates from this model will be detailed in Appendix A for thorough examination.

We also conducted similar F-test for the coefficient estimates of offensive and defensive penalty yards. In this test, the null hypothesis was that the difference between the offensive and defensive penalty yards coefficients is zero. We found a significant difference, with the coefficient for offensive penalty yards being larger than that for defensive penalty yards. This suggests that losing a yard due to an offensive penalty is more detrimental than gaining a yard due to a defensive penalty, although the difference is not substantial. This finding is somewhat surprising, given the prevalent belief in the football community that defensive penalties are more harmful than offensive ones. Our results indicate that this belief may not hold true at an aggregated game level, inviting further exploration into the nuanced dynamics of penalties in football.

3.2 Linear Regression with Random Intercepts

In light of the insights gleaned from our linear regression analysis, particularly regarding the significance of penalties, it is pertinent to explore additional dimensions of the data. Specifically, the potential correlation of games within the same season and the correlation of games played by the same teams merit investigation. To address these aspects and examine the variability both seasonally and across teams, we will employ linear models with random intercepts. This approach will enable us to assess the significance of random effects attributable to seasons and teams, as well as the variability inherent in these factors.

3.2.1 Season Varying Intercept Model

The model in this section incorporates only the season as a random effect, thereby accounting for intra-seasonal correlations. The model is defined as follows.

$$y_{ij} = \beta_0 + \beta_1 x_{1,ij} + \beta_2 x_{2,ij} + \cdots + \beta_k x_{k,ij} + u_j + \epsilon_{ij} \tag{3.1}$$

$$\text{where } u_j \sim N(0, \sigma_u^2), \epsilon_{ij} \sim N(0, \sigma^2)$$

- y_{ij}: The score differential for the i-th game in the j-th season.
- β_0: The intercept of the model, representing the baseline score differential when all covariates are at their reference levels.
- $\beta_1, \beta_2, \ldots, \beta_k$: The fixed effect coefficients for each of the covariates. These coefficients measure the average change in the score differential for a one-unit change in the respective covariate, holding all other covariates constant.
- $x_{1,ij}, x_{2,ij}, \ldots, x_{8,ij}$: The covariates for the i-th game in the j-th season. These correspond to spread line, fumbles differential, interceptions differential, sacks differential, passing yards differential, rushing yards differential, offensive penalties EPA differential, and defensive penalties EPA differential.
- u_j: The random effect for the j-th season. This term captures any season-specific effects on the score differential that are not explained by the fixed effects.
- ϵ_{ij}: The residual error term for the i-th game in the j-th season. This term captures the variability in the score differential that is not explained by the fixed or random effects.

Table 3.3 below give the estimates from the model with only season as a random intercept.

Table 3.3: Estimates for Linear Regression with Random Intercept (Model 3.1)

Variables	Estimate	Std. Error	t value
Intercept	0.250	0.182875	1.368
Spread Line	0.391163	0.025343	15.435
Fumbled(D)	-1.413770	0.089329	-15.827
Intercepted(D)	-3.765639	0.102643	-36.687
Sacked(D)	-1.294291	0.058561	-22.102
Passing Yards(D)	0.041920	0.001570	26.703
Rushing Yards(D)	0.071210	0.001894	37.592

Continued on next page

Table 3.3 – Continued from previous page

Variables	Estimate	Std. Error	t value
Offensive Penalty EPA(D)	0.650350	0.050055	12.993
Defensive Penalty EPA(D)	-0.452989	0.037842	-11.971
Random Intercept			
Season (σ_u^2)	0.1387	0.372	

In Table 3.3, we observe that while all fixed effects maintain their significance, the impact of the random effects estimates appears to be minimal. To further investigate this, Table 3.4 presents the results of a Likelihood Ratio Test (LRT). This test examines the null hypothesis, which posits that models with and without random effects do not exhibit significant differences.

Our findings indicate that the p-value exceeds the 0.05 threshold, leading us to fail to reject the null hypothesis. This suggests that the variation in score differentials across seasons is not statistically significant. It is noteworthy that the estimates related to penalty variables are nearly identical to those in the model without random effects. This result agrees with our exploratory analysis in Figure 2.10 that didn't show any major variation in score differentials across the seasons.

Another explanation for the insignificance of seasonal variation could be the inclusion of the spread line variable in our model. It is conceivable that the primary sources of variation across seasons are changes in rules, strategies, and team dynamics. The spread line effectively encapsulates these effects, thereby diminishing the apparent variability among seasons. This highlights the comprehensive nature of the spread line in accounting for a wide array of influencing factors in the model.

Model	npar	AIC	BIC	logLik	deviance	Chisq	Df	Pr($>\chi^2$)
Fixed Effects	9	22542	22597	-11262	22524			
Random Effects	10	22542	22603	-11261	22522	1.9661	1	0.1609

Table 3.4: Model Comparison Statistics

Future research endeavors could beneficially incorporate a predictor that effectively quantifies the quality of the team besides the spreadline. We acknowledge that the spreadline is one of the best variables to account for team quality since gambling markets are efficient. Other variables might give some different insights. Such an addition would offer a more comprehensive understanding of the team's influence on game outcomes, particularly in light of the substantial transformations teams undergo from season to season. This consideration would undoubtedly enrich the model and provide more nuanced insights into the dynamics of team sports.

3.2.2 Season and Team Varying Intercept Model

This model is more complex, integrating both team and season as random effects, with an additional consideration of the nested structure of team within season. This layered approach allows us to disentangle the individual and combined influences of season and team on the outcomes.

$$y_{ijk} = \beta_0 + \beta_1 x_{1,ijk} + \beta_2 x_{2,ijk} + \cdots + \beta_n x_{n,ijk} + w_{jk} + \epsilon_{ijk} \tag{3.2}$$

$$\text{where } w_{jk} \sim N(0, \sigma_w^2), \epsilon_{ijk} \sim N(0, \sigma^2)$$

- y_{ijk}: The score differential for the i-th game in the j-th season and for the k-th penalty team.
- $x_{1,ijk}, x_{2,ijk}, \ldots, x_{n,ijk}$: The covariates for the i-th game in the j-th season and for the k-th team. These correspond to spread line, fumbles differential, interceptions differential, sacks differential, passing yards differential, rushing yards differential, offensive penalties EPA differential, and defensive penalties EPA differential.
- w_{jk}: The random effect for the k-th home team nested within the j-th season. This accounts for team-specific variations within each season.
- ϵ_{ijk}: The residual error term for the i-th game in the j-th season and for the k-th home team. This term captures the variability in the score differential that is not explained by the fixed or random effects.

Table 3.5 below give the estimates from the model with a random effect for team and a nested random effect for teams within seasons.

Table 3.5: Esimates for Linear Regression with Random Intercept (Model 3.2)

Variables	Estimate	Std. Error	t value
Intercept	0.27	0.18	1.49
Spread Line	0.38	0.03	14.61
Fumbled(D)	-1.41	0.09	-15.78
Intercepted(D)	-3.75	0.10	-36.57
Sacked(D)	-1.30	0.06	-22.18
Passing Yards(D)	0.04	0.002	26.63
Rushing Yards(D)	0.07	0.002	37.51
Offensive Penalty EPA(D)	0.65	0.05	12.92
Defensive Penalty EPA(D)	-0.45	0.04	-11.95
Random Intercept			
Home Team	0.3232	0.5685	
Home Team:Season	0.2848	0.5337	

In Table 3.6, we present the results of a Likelihood Ratio Test (LRT) conducted to assess the impact of random effects in our model. The test is structured around the null hypothesis, which asserts that there is no significant difference between models with and without random effects.

Our analysis reveals that the p-value does not fall below the 0.05 threshold. Consequently, we fail to reject the null hypothesis, suggesting that there is no significant variation among teams or across different seasons in our data set.

This consistency in results may be attributed to the inclusion of the spread line variable in our predictors. The spread line, by its nature, is likely to incorporate a range of factors that influence team performance and seasonal changes. Therefore, its presence in the model could be masking potential variations among teams and seasons, underlining its comprehensive scope in capturing diverse elements that affect the outcomes in our study.

Model	npar	AIC	BIC	logLik	deviance	Chisq	Df	Pr(>χ^2)
Fixed Effects	9	22542	22597	-11262	22524			
Nested Random Effects	11	22543	22616	-11259	22519	3.1	2	0.2122

Table 3.6: Model Comparison Statistics

The preceding analysis elucidates that the variations in score differentials between seasons or teams are not substantial. This finding implies that our linear regression models sufficiently capture the variations in score differentials within the available data set.

A key aspect of our attention remains focused on the coefficients for offensive and defensive penalty Expected Points Added (EPA). It is noteworthy that these coefficients are remarkably consistent with those derived from the linear regression model. This consistency aligns with our observation that random effects do not contribute significantly to explaining the variations in score differentials, either among teams or across different seasons. This reinforces the robustness of the linear regression model in our analysis, especially in the context of the variables under consideration.

3.3 Discussion

In the study presented, we examined the impact of offensive and defensive Expected Points Added (EPA) on the score differentials in football games. Our analysis revealed that both offensive and defensive penalties have a significant effect on these differentials. However, when considering aggregated game-level data, no notable difference was observed in the impact of offensive versus defensive penalties on score differential.

Further investigation identified additional variables that significantly influence score differentials. The estimated values of these variables align with conventional understanding of the game's dynamics, lending credibility to our findings. It is important to note, however, that our model exhibits certain limitations, particularly in its level of detail.

Even though our model is unable to predict score differentials in real-time since that is not the primary purpose of the model. Nevertheless, the model holds potential for predicting score differentials in hypothetical scenarios. This predictive capability could be enhanced by incorporating additional predictor variables, such as Redzone conversion rates. This refers to the frequency with which a team scores when positioned within 20 yards of the opposing team's end zone.

Our analysis also highlights a notable trend: the magnitude of the residuals increases with the size of the score differentials. This observation underscores the inherent unpredictability of game outcomes, especially in instances of exceptionally large score differentials. This finding contributes to the broader discourse on the unpredictable nature of sports outcomes and the factors influencing them.

Overall, while the model demonstrates a significant capacity to explain the variability in score differentials using selected variables, its accuracy and predictive power could be further improved with the inclusion of additional, more granular data points.

Chapter 4

Drive Outcomes: Logistic Regression vs. Random Forest

The dynamics of American football, especially within the NFL, often hinge on split-second decisions and strategic plays. Among these, penalties occupy a particularly debated space. While seemingly straightforward, the question of whether offensive or defensive penalties exert a more substantial negative impact on a team's performance remains contentious. Common football discourse suggests that defensive penalties, particularly those leading to significant offensive gains or automatic first-down conversions, may be more detrimental. However, our analysis seeks to explore this narrative more comprehensively.

In our study, we employed linear regression models with fixed effects and random intercepts to assess the implications of both offensive and defensive penalties on game outcomes, which we delineated as score differentials. To holistically appraise the impact of penalties, we approached them from three distinct perspectives: the total count of penalties, the accumulated yardage resulting from these penalties, and their effect on expected points added (EPA).

To delve deeper and provide a more nuanced understanding of this phenomenon, we transitioned our focus to drive-level data. Here, we sought to model the outcome of individual drives, categorizing them into two principal outcomes: 'Score' and 'No Score'. A 'Score' was defined as any drive culminating in either a touchdown or a field goal. In contrast, 'No Score' encompassed all other possible outcomes, with one notable exception: drives ending in safety were deliberately excluded from our analysis to maintain stability and mitigate potential outliers.

This approach serves a twofold purpose. Firstly, it allows us to glean insights into the roles of offensive and defensive penalties via the coefficients derived from logistic regression. Secondly, it offers an alternative means of evaluating the overall performance of the model, enriching our understanding of its predictive accuracy and robustness. We will account a bunch of other covariates and interaction that might have an impact on the outcome of the drive. This will allow us to isolate the impact of penalties.

4.1 Logistic Regression

In the preliminary phase of our analysis, a rigorous selection method was employed to select the covariates deemed pertinent for the logistic regression model. The Table 2.3 presents the covariates we used to fit this model. These covariates were chosen after careful consideration for the dynamics of the game, preventing any data leakage and multi-collinearity issues. Some higher order terms were also considered and the interaction between quarter and penalty metrics but after a stepwise variable selection method, these were the variable that we decided to keep in the model. The data from season 2011-2021 was used to fit the model and the year 2022 was held out to do predictions and determine predictive accuracy. The function 'glm' from the package 'stats' in R was used to fit the model.

Given the dataset, our response variable, y, indicates whether a drive ended with a score (represented as 1) or not (represented as 0). The logistic regression model relates the probability of a drive ending in a score to various predictor variables. Formally, the probability, $P(y = 1)$, is given by:

$$P(y = 1) = \frac{e^{\beta_0+\beta_1 x_1+\beta_2 x_2+\cdots+\beta_n x_n}}{1 + e^{\beta_0+\beta_1 x_1+\beta_2 x_2+\cdots+\beta_n x_n}} \tag{4.1}$$

Where:

- β_0 is the intercept.
- $\beta_1, \beta_2, \ldots, \beta_n$ are the coefficients of the predictor variables.
- $x_1, x_2, \ldots, x_n$ are the predictor variables.

4.1.1 Results

Table 4.1: Logistic Regression Model Results for Analyzing Drive Outcomes

	Dependent variable: Drive Outcome
(Intercept)	$0.3704^{***}(0.009)$
Yardline 100	$-0.004^{***}(0.000)$
Offensive Penalty EPA	$0.056^{***}(0.004)$
Defensive Penalty EPA	$0.059^{***}(0.003)$
Score Differential	$0.000(0.0001)$
Drive Inside20	$0.554^{***}(0.004)$
ha drive	$0.008^{**}(0.003)$
Posteam Timeouts Remaining	$-0.00003(0.001)$
Fumble	$-0.136^{***}(0.004)$
qtr2	$0.034^{***}(0.004)$

Continued on next page

	Dependent variable: Drive Outcome
qtr3	$0.009^{*}(0.004)$
qtr4	$-0.027^{***}(0.004)$
Drive Time of Possession	$0.0009^{***}(0.00001)$
Offensive penalty EPA:Drive Time of Possession	$0.000(0.00001)$
Defensive Penalty EPA:Drive Time of Possession	$-0.0001^{***}(0.00001)$
Observations	66835

Note: *p<0.1; **p<0.05; ***p<0.01
p-values are associated with the null hypothesis that the corresponding coefficients are equal to zero.

Table 4.1 offers valuable insights into various factors influencing the scoring probabilities in a game. It is essential first to validate these results against our understanding of the game's dynamics.

Starting with the yardline 100 coefficient, its negative value aligns with conventional wisdom: being further from the opponent's zone typically reduces scoring likelihood. Similarly, the positive coefficients for both offensive and defensive penalties are logical, as each incremental point gained via these penalties increases the probability of scoring.

A Wald test was conducted on the coefficients of offensive penalty Expected Points Added (EPA) and defensive penalty EPA. The null hypothesis posited that these coefficients are equal. Given that the p-value exceeded 0.05, there wasn't sufficient evidence to reject this hypothesis, suggesting that offensive and defensive penalties do not significantly differ in their impact on a drive's outcome.

The coefficient for score differential is not significant. The significant positive impact of a drive within the 20-yard line is intuitive, as proximity to the goal line typically correlates with higher scoring chances. Home advantage is also reflected in the positive coefficient of the 'home advantage drive' variable, highlighting the benefits of playing in familiar conditions.

The coefficient estimate for remaining timeouts is not significant. In contrast, the negative coefficient for fumbles reflects the potential for turnovers or momentum disruption.

The negative coefficients for 4 might indicate heightened defensive efforts or the higher offensive team fatigue during this critical phase of the game - the final quarter (excluding rare overtime situations).

The drive time possession coefficient estimate is positive, suggesting that longer drives might associate with a positive outcome. Finally, the interaction coefficients for offensive and defensive penalties merit further investigation. The zero interaction coefficient for offensive penalties and the negative one for defensive penalties suggest there might be underlying dynamics that require additional research for comprehensive understanding.

In our analysis, we applied the developed model to predict the outcomes of drives from the 2022 season. The model demonstrated an accuracy rate of 84%, a notable achievement considering the inherent challenges in predicting touchdowns. It's important to recognize that the majority of drives do not result in touchdowns, making accurate predictions in this context particularly difficult.

The primary objective of this exercise was not merely to predict touchdowns or field goal but to assess the relative impact of offensive and defensive penalties on drive outcomes. Our findings indicate that the differences in impact between offensive and defensive penalties are not statistically significant. This conclusion is reinforced by the p-value, which suggests that our predictions are significantly more accurate than random guesses.

This result is important as it provides a robust validation of our model, highlighting its effectiveness in isolating and evaluating the specific effects of penalties on the progression and outcome of a drive. The model's high accuracy in predicting drive outcomes, despite the complexities involved in such predictions, is a testament to its robustness and the validity of its underlying assumptions and calculations.

Table 4.2: Confusion Matrix

		Predicted	
		Score	No Score
Actual	Score	1535	705
	No Score	234	3636

- Accuracy: [**0.8463**]
- 95% CI : [**(0.837, 0.855)**]
- NIR: [**0.63**]
- P-Value [Acc > NIR] : [$< 2.2e - 16$]

4.2 Random Forest

Random Forest, a sophisticated ensemble learning technique, has gained prominence in various fields, including NFL analytics, for its ability to handle complex classification and regression tasks. This method, which operates by constructing multiple decision trees during training and outputting the mode of the classes (for classification) or mean prediction (for regression) of the individual trees, is particularly effective in handling large datasets with numerous input variables, making it an ideal tool for analyzing the multifaceted nature of football data.

In the realm of NFL analytics, Random Forest has been employed in several innovative ways. One notable application is in predictive analysis models for NFL play-by-play data, where a Random Forest classifier was used to input various play situations—such as time, down, yards to go, and score—and output a prediction of the play type. This approach allows for nuanced understanding of game dynamics and strategies [19].

Another significant application of Random Forest in NFL analytics is in classifying the winner of the Super Bowl. Using data collected during the regular season, Random Forest algorithms, along with other methods, were compared for their efficacy in predicting the championship game's outcome. This not only highlights the algorithm's predictive power but also its utility in high-stakes scenarios where accurate predictions are invaluable [18].

Moreover, Random Forest methods have been utilized to estimate the Win Probability (WP) before each play of an NFL game. By combining pre-play variables, these models provide WP estimates that closely resemble true win probabilities and accurately predict game outcomes, particularly in the later stages of games. This application underscores the algorithm's ability to integrate and analyze real-time data, offering insights that can influence game strategies and decision-making [10].

As we delve deeper into the application of Random Forest in NFL analytics, it is essential to appreciate the versatility and robustness of this method. Its ability to process complex datasets and yield reliable predictions makes it a powerful tool in sports analytics, offering insights that can shape team strategies and enhance understanding of the game.

4.2.1 Results

The process of implementing the Random Forest algorithm in our study commenced with the initiation of parallel processing capabilities. This was achieved through the makeCluster(detectCores()) function, which facilitates the creation of a computing cluster using the maximum number of available CPU cores. Subsequently, registerDoParallel(cl) was invoked to enable the use of this cluster for parallel computations, thereby optimizing computational efficiency.

To ensure the robustness and generalizability of Random Forest, a cross-validation approach was adopted. This was configured using the trainControl method, specifying the cross-validation (cv) technique with five folds (number = 5). This method allows for a comprehensive assessment of the model's performance across different subsets of the data.

A crucial step in optimizing the performance of Random Forest involved tuning the mtry parameter, which represents the number of variables randomly sampled as candidates at each split. A grid of possible mtry values ranging from 1 to 12, incremented by 2, was created. This grid facilitated the exploration of different model configurations to identify the optimal mtry value.

The Random Forest, was then trained on the training dataset using the train function from the 'caret' package. The response variable was regressed on various predictor variables, including yardline_100, offensive EPA, defensive EPA, and others. The model was configured with 500 trees (ntree = 500) and the previously established cross-validation settings.

Upon completion of the training process, the optimal mtry value was identified based on model performance metrics. The best mtry value for accuracy and Cohen's Kappa were extracted from the model results. These values represent the configurations that yielded the highest accuracy and Kappa statistic, respectively, providing insights into the most effective model settings for our predictions. The predictors used in random forest were same as the logistic regression model.

Table 4.3: Confusion Matrix

		Predicted	
		Score	No Score
Actual	Score	1737	503
	No Score	351	3519

- Accuracy: [**0.8602**]
- 95% CI : [**(0.8513, 0.8688)**]
- NIR: [**0.63**]
- P-Value [Acc > NIR] : [$< 2.2e - 16$]

4.3 Model with Game ID as a Random Intercept

The model in this section incorporates game id as a random effect, thereby accounting for intra-game correlations. we presents the results of the model below in Table 4.4.

Table 4.4: Coefficient estimates for Logistic Rgeression with Random Intercept

	Dependent variable:Drive Outcome
	Estimates (Standard Error)
Intercept	-1.482^{***} (0.035)
scale(Yardline 100)	-0.574^{***} (0.013)
Drive Inside20	2.785^{***} (0.028)
ha drive	0.066^{***} (0.023)
binary__fumble	-1.489^{***} (0.048)
qtr2	0.297^{***} (0.033)
qtr3	0.077^{**} (0.033)
qtr4	-0.262^{***} (0.034)

Continued on next page

Table 4.4 – Continued from previous page

	Dependent variable:Drive Outcome
	Estimates (Standard Error)
scale(Offensive EPA)	0.331*** (0.015)
scale(Drive Time of Possession)	0.784*** (0.015)
scale(Defensive EPA)	0.301*** (0.014)
scale(Offensive EPA):scale(Drive Time of Possession)	−0.049*** (0.013)
scale(Drive Time of Possession):scale(Defensive EPA)	−0.129*** (0.012)
Random Intercept	
Game ID	0.018(0.134)
Observations	66,835
Log Likelihood	−25,002.550
Akaike Inf. Crit.	50,033.110
Bayesian Inf. Crit.	50,160.650
Note:	*p<0.1; **p<0.05; ***p<0.01

Based on the results it seems like the variation of drive success between different games is not significant. The coeffecient estimates for rest of the variabels act in a similar manner like the Fixed Effects Logistic Regression Model.

4.4 Comparing the Results and Discussion

The results obtained from our logistic regression analysis indicate that both offensive and defensive penalty Expected Points Added (EPA) significantly influence the outcome of football drives. Notably, the analysis does not reveal a substantial difference between the impacts of offensive and defensive penalties, suggesting that neither type of penalty is inherently more detrimental to the drive outcome than the other. Our study also identified several other covariates that play a significant role in predicting drive outcomes. The estimated coefficients for these variables align with conventional football strategies and outcomes, validating the model's relevance to real-world scenarios. However, it's crucial to acknowledge that our model is not suited for real-time predictions, as the covariates are assessed post-drive.

The model demonstrated a high level of accuracy in predicting drive outcomes for the 2021 football season, indicating its efficacy with the chosen covariates. Yet, there remains room for improvement. The inclusion of more detailed covariates could enhance the model's precision and granularity. A key insight from this study is that despite the inherent complexity of football, a significant portion of the variance in drive outcomes can be explained by a relatively small set of variables. This finding underscores the potential for predictive modeling in sports analytics.

Looking forward, the model could benefit from distinguishing between field goals and touchdowns as separate outcomes. This distinction is particularly relevant in assessing the impact of penalties, as the outcome of a drive resulting in a field goal can sometimes be suboptimal, especially when a touchdown was a feasible alternative. For instance, a team settling for a field goal despite being in the red zone (20 yard line) could indicate a missed opportunity for a more favorable outcome. To address this nuanced aspect of drive outcomes, future iterations of the model might employ a multinomial logistic regression approach, which would offer a more refined analysis of the different types of drive outcomes and their contributing factors. The drives in the same game might also not be independent and adding a random intercept for the game ID is something we could look into in the future.

In the comparision of model performances, the Random Forest marginally outperformed the logistic regression model, achieving an accuracy of 86% as opposed to 84% by the logistic regression model. Despite this slight edge in predictive accuracy, we advocate for the use of logistic regression over Random Forest in contexts where model interpretability is paramount. The inherent transparency and simplicity of logistic regression make it a more suitable choice in scenarios where understanding the influence and relationship of individual predictors is crucial.

However, it is important to acknowledge the potential of the Random Forest algorithm, especially in handling more complex and nuanced datasets. Random Forest's robustness against multicollinearity and its proficiency in capturing complex, non-linear relationships suggest that its predictive accuracy could be significantly enhanced with the inclusion of more sophisticated predictor variables. Random Forest is a valuable tool in scenarios where model interpretability is less critical than predictive accuracy and handling of complex data structures.

Chapter 5

Discussion

In our study, we employed linear regression models with fixed effects and random intercepts to examine the influence of penalties on game outcomes, focusing on score differentials as a key measure. Contrary to the widespread belief that defensive penalties have a more detrimental effect than offensive ones, our analysis did not yield definitive evidence supporting this notion.

The random interept models revealed that the score differential does not exhibit significant variation across seasons or among teams penalized within a season. This model echoed the findings of the linear regression analysis, indicating no substantial difference in the impact of offensive versus defensive penalties.

Subsequently, our project shifted to a drive-level analysis, employing a logistic regression framework to model the outcomes of drives. Consistent with our previous findings, this model also indicated no significant disparity in the effects of offensive and defensive penalties on drive outcomes.

It is important to acknowledge the limitations of our model. One key constraint is its inability to predict outcomes before a drive begins, as it requires in-drive covariates. However, the model remains valuable for analyzing hypothetical scenarios aposteriori.

Looking forward, there are several potential avenues for further research. One such direction includes separately modeling the effects of offensive and defensive penalties, which could yield more nuanced insights into their respective impacts on game dynamics. This approach may help in uncovering subtler patterns and effects not apparent in the current analysis. Also, our analysis was limited to aggregate game-level and drive-level data. The next step would be to look into play by play data and see if one of the offensive or defensive penalties have a more detrimental impact on the value of the play. We also plan to look into how a single player can have an impact based on penalties. In the future we plan to explore this more with Bayesian Hierarchical Models.

Bibliography

[1] nflfastR package (www.nflfastr.com).

[2] Pro football reference (www.pro-football-reference.com).

[3] Scott Bingham. Decline penalty football: A strategic guide for decision-making, Sep 2023.

[4] Nader Chmait and Hans Westerbeek. Artificial intelligence and machine learning in sport research: An introduction for non-data scientists. *Frontiers in Sports and Active Living*, 3, 2021.

[5] Alex Farnell. False start? an analysis of NFL penalties with and without crowds. *Journal of Sports Economics*, 24(6):695–716, 2023.

[6] George Foster, Norm O'Reilly, and Zachary Naidu. Playing-side analytics in team sports: Multiple directions, opportunities, and challenges. *Frontiers in Sports and Active Living*, 3, 2021.

[7] Joël Guérette, Caroline Blais, and Daniel Fiset. The absence of fans removes the home advantage associated with penalties called by national hockey league referees. *PLOS ONE*, 16(8), 2021.

[8] David Jarvis and Kevin Westcott. The use of athlete data analytics, Apr 2022.

[9] Jack Lichtenstein. Exploring stability and predictive power of penalties in the NFL, Jan 2021.

[10] Dennis Lock and Dan Nettleton. Using random forests to estimate win probability before each play of an NFL game. *Journal of Quantitative Analysis in Sports*, 10(2), 2014.

[11] M. J. Maher. Modelling association football scores. *Statistica Neerlandica*, 36(3):109–118, 1982.

[12] Armin Marquez, Beth A. Cianfrone, and Timothy Kellison. Factors affecting spectators' adoption of digital ticketing: The case of interscholastic sports. *International Journal of Sports Marketing and Sponsorship*, 21(3):527–541, 2020.

[13] Isabel Pantle. Analyzing the predictive power of NFL statistics, Jan 2022.

[14] Shankar Parameshwaran. How analytics can boost competitiveness in sports, Jan 2023.

[15] Konstantinos Pelechrinis and Evangelos Papalexakis. The anatomy of American Football: Evidence from 7 years of NFL game data. *PLOS ONE*, 11(12), 2016.

[16] Matthew J. Penn and Christl A. Donnelly. Analysis of a double poisson model for predicting football results in euro 2020. *PLOS ONE*, 17(5), 2022.

[17] Kathryn E. Phillips and William G. Hopkins. Determinants of cycling performance: A review of the dimensions and features regulating performance in elite cycling competitions. *Sports Medicine - Open*, 6(1), 2020.

[18] Yazan F. Roumani. Sports analytics in the NFL: Classifying the winner of the superbowl. *Annals of Operations Research*, 325(1):715–730, 2022.

[19] Veeral Shah. NFL play predictive analysis.

[20] SoccerTAKE. Soccer analytics: How data is changing the game, May 2023.

[21] Ronald Yurko, Samuel Ventura, and Maksim Horowitz. Nflwar: A reproducible method for offensive player evaluation in football. *Journal of Quantitative Analysis in Sports*, 15(3):163–183, 2019.

[22] Övünç Yılmaz, Rob F. Easley, and Mark E. Ferguson. The future of sports ticketing: Technologies, data, and new strategies. *Journal of Revenue and Pricing Management*, 22(3):219–230, 2023.

Appendix A

Regression Results

Table A.1: Regression Models for Analyzing Score Differential(Yards)

Differentials	*Dependent variable: Score Differential*
Spread Line	$0.405^{***}(0.024)$
Offensive Penalty Yards	$-0.078^{***}(0.006)$
Defensive Penalty Yards	$-0.054^{***}(0.005)$
Fumble	$-1.462^{***}(0.090)$
Interception	$-3.948^{***}(0.103)$
Sacks	$-1.297^{***}(0.059)$
Rushing Yards	$0.073^{***}(0.002)$
Passing Yards	$0.042^{***}(0.002)$
R^2	0.702
Adjusted R^2	0.702
Residual Std. Error	7.96 (df = 3233)
F Statistic	955.5^{***} (df = 8; 3233)

Note: $^{*}p<0.1$; $^{**}p<0.05$; $^{***}p<0.01$
p-values are associated with the null hypothesis that the corresponding coefficients are equal to zero.

Appendix B

List of Different Penalties

Table B.1: Penalty Types and Counts

Penalty Type	Count
Offensive Holding	7843
False Start	7035
Defensive Pass Interference	3121
Unnecessary Roughness	2460
Defensive Holding	2438
Defensive Offside	1924
Delay of Game	1703
Neutral Zone Infraction	1602
Illegal Block Above the Waist	1450
Roughing the Passer	1317
Illegal Use of Hands	1084
Face Mask	1050
Offensive Pass Interference	1026
Illegal Formation	672
Illegal Contact	653
Unsportsmanlike Conduct	592
Encroachment	583
Intentional Grounding	423
Illegal Shift	341
Ineligible Downfield Pass	314
Defensive 12 On-field	284
Taunting	266
Horse Collar Tackle	242
Offside on Free Kick	206
Defensive Too Many Men on Field	191
Illegal Blindside Block	129
Illegal Motion	122
Chop Block	108
Illegal Substitution	97
Running Into the Kicker	93

Table B.1: Penalty Types and Counts

Penalty Type	Count
Lowering the Head to Initiate Contact	92
Personal Foul	90
Tripping	89
Disqualification	83
Illegal Forward Pass	82
Player Out of Bounds on Punt	80
Ineligible Downfield Kick	79
Roughing the Kicker	78
Illegal Touch Kick	74
Low Block	70
Fair Catch Interference	69
Illegal Touch Pass	64
Clipping	56
Offensive Too Many Men on Field	56
Offensive 12 On-field	53
Leverage	41
Player Out of Bounds on Kick	41
Illegal Double-Team Block	40
Illegal Crackback	35
Defensive Delay of Game	29
Interference with Opportunity to Catch	28
Offensive Offside	20
Invalid Fair Catch Signal	15
Kick Catch Interference	14
Leaping	10
Illegal Bat	9
Illegal Peelback	9
Delay of Kickoff	6
Illegal Wedge	5
Illegal Kick	1
Illegal Kick/Kicking Loose Ball	1
Illegally Kicking Ball	1
Kickoff Out of Bounds	1
Short Free Kick	1