Impact of Load Management on NBA Player Performance

Dena Chen, Edison Huang, Clement Ou, Sahil Parikh
April 28, 2025

I. Introduction

Fatigue is a critical factor in performance when playing any type of sport. In the NBA season, teams play 82 games in approximately 6 months, with players expected to perform consistently at a high caliber of play. While many believe that rest days for NBA basketball players help reduce injury risk and enhance performance during high-stakes games, the question remains: Does rest actually lead to improved performance in subsequent games, and does it contribute to better team outcomes? The motivation for this project is to provide valuable insights for teams and players to optimize their performance as it can influence strategic decisions such as minute distribution, game scheduling, as well as player rest schedules.

In our analysis, we used different models to explore the relationship between load management factors and their effects on player performance. Specifically, Ordinary Least Squares (OLS) and Linear Mixed-Effects (LMER) models are used to examine the relationship between rest, workload, opponent strength, and plus/minus. In addition, we used Decision Tree-based models (such as XGBoost) and SHAP (SHapley Additive exPlanations) analysis to capture the non-linear relationships and interpret the influence of individual features.

The results indicate that the number of rest days had little to no impact on player performance. Instead, features related to minutes played were more impactful. Interestingly, cumulative minutes played and load management (less than 10 minutes of play in the previous game) both had positive effects in the linear model. Furthermore, SHAP values confirm that minutes played and previous game minutes have a negative relationship with plus/minus, reinforcing the idea that fatigue from heavy usage can reduce a player's effectiveness.

II. Data

Our data comes from the hoopR-data repository's NBA data, accessed with the load_nba_player_box() function. The dataset is a table with 6,696 rows of player box scores, with 57 columns of variables providing relevant information for each row of data.

Key variables in the dataset include:
- Player information: athlete_id, athlete_display_name, athlete_position_name
- Game context: game_date, season, team_name, opponent_team_name, home_away
- Performance metrics: minutes, plus_minus
- Team performance: team_score, opponent_team_score, team_winner

To better utilize the data in our project, we conducted simple data cleaning. This includes removing rows with missing plus/minus or minutes data and calculating rest days between games based on dates. Additionally, binary variables for features like home/away games were created, and rolling averages and standard deviations of minutes played were computed to capture trends. The final dataset includes variables such as rest days, previous game minutes, rolling averages of minutes played, a home/away indicator, team, and opponent performance (win percentage), a load management indicator flagging if a player was used for fewer than 10 minutes their previous game, player positions, and whether a player is in a back-to-back game scenario (i.e., their last game was the previous day).

To explore the data, we conducted some simple exploratory data analysis.

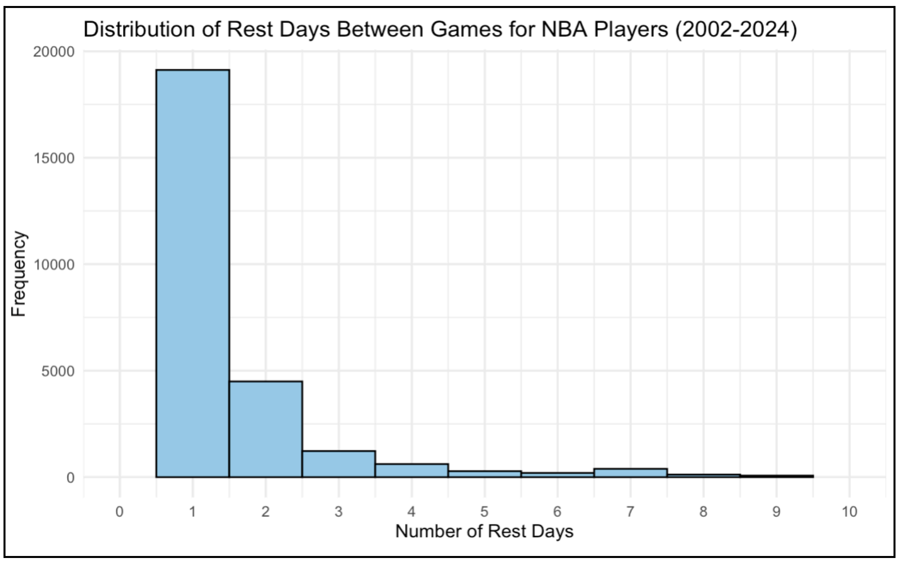

Based on the chart above, we can see that the distribution of rest days between games for NBA players is heavily right-skewed with the peak at 1 rest day. This shows that most NBA players only have 1 rest day between games.

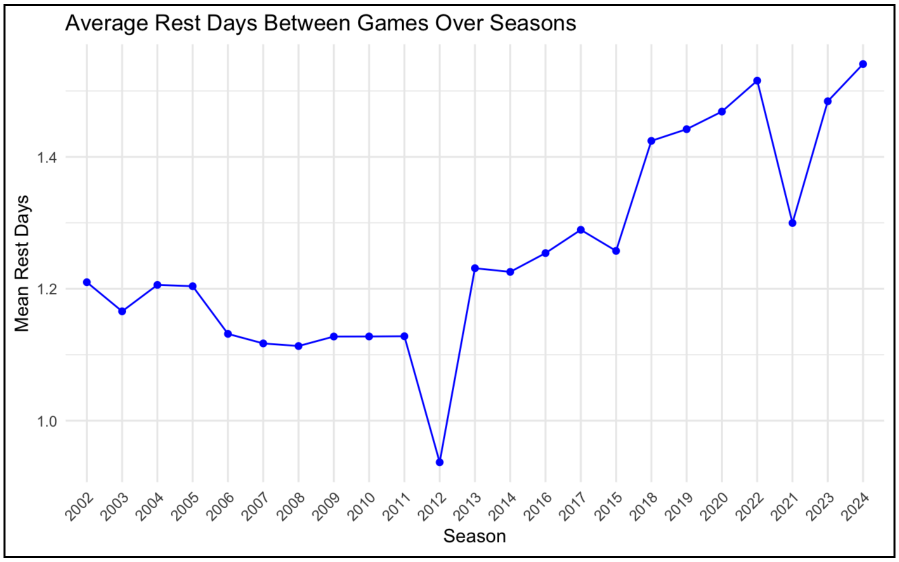

Additionally, as we compare the average rest days over time from 2002 to 2024, we can see that despite the increase in the average number of rest days, the average is still relatively low at 1.5 rest days in 2024. There is a small upward trend in the average number of rest days, but the difference is minimal over the 12 years.

III. Methods

III. 1. Linear Models

In our methodology, we first fit a baseline ordinary least squares regression of plus/minus on four predictors—scaled days since last game, scaled three‑game cumulative minutes, a binary rest indicator, and scaled opponent strength—assuming independent, homoskedastic Gaussian errors so that each coefficient reflects the change associated with a one‑standard‑deviation shift in its covariate. Because our data include repeated observations on the same athletes and teams (violating the OLS independence assumption), we next fitted a two‑level linear mixed‑effects model with random intercepts for athlete and team, treating both random effects and residuals as draws from zero‑mean normal distributions with variances estimated by restricted maximum likelihood. To capture the fact that some players and teams may benefit more or less from rest, we then extended this model to include random slopes on the rest indicator for both grouping factors. We compared these three nested models using AIC, BIC, and likelihood‑ratio tests, and we supplemented those criteria with marginal and conditional R-squared statistics to gauge how much variance is explained by fixed effects alone versus the full hierarchical structure. We also inspected residual‑versus‑fitted plots and normal quantile-quantile plots to verify that homoscedasticity and approximate normality held in each model. Finally, we report all fixed‑effect estimates with 95 percent confidence intervals (±1.96 times the standard error), and we compute prediction intervals for player-level intercepts to quantify uncertainty in our key rest and workload effects. This combination of OLS and mixed‑effects modeling, rigorous model comparison, and thorough uncertainty quantification provides a principled framework for answering how rest and recent workload influence NBA plus/minus while accounting for athlete and team heterogeneity.

III. 2. Decision Tree

We were interested in exploring how more complex approaches, such as Gradient Boosted Decision Trees (GBDT), might benefit from incorporating load management features. Specifically, we utilized XGBoost, which implements GBDTs in R very efficiently and is easy to use. We chose this approach for its ability to model potentially non-linear relationships between predictor variables (like rest days, minutes played, rolling averages) and the target variable (plus/minus) without requiring pre-defined interaction terms. GBDTs build an ensemble of decision trees sequentially, where each new tree attempts to correct the errors made by the previous ones. This allows the model to capture complex patterns and interactions in the data that might be missed by linear models. Furthermore, XGBoost includes built-in regularization techniques to prevent overfitting. To interpret the results of this complex model and understand the contribution of each feature to the prediction for individual games, we used SHAP (SHapley Additive exPlanations). SHAP values provide a way to decompose a prediction into the contribution of each feature, offering both local (per-game) and global (overall) insights into feature importance and impact direction, thereby complementing the coefficient-based interpretation of the linear models. Model performance was evaluated using standard regression metrics like Root Mean Squared Error (RMSE), Mean Absolute Error (MAE), and R-squared on a held-out test set.

IV. Results

IV. 1. Linear Model Results

We first evaluated fit and predictive performance across our three linear models. The OLS regression explains about 4% of the variation in plus/minus (R^2 = 0.0417, residual SD = 11.18). Introducing player- and team-level random intercepts in the mixed-effects model (LMER) reduces the residual SD slightly to 10.99 and lowers AIC by roughly 1500 points. Allowing the rest-effect to vary by athlete and team (random slopes) yields a further AIC drop of ~320 and a highly significant likelihood-ratio test (ΔX^2(6)=331, $p < 0.001$), confirming that both baseline ability and the benefit of rest differ meaningfully across players and teams (see figure below). Despite these improvements, the residual SD of ~11 points highlights that a large share of game-to-game variability remains unmodeled.

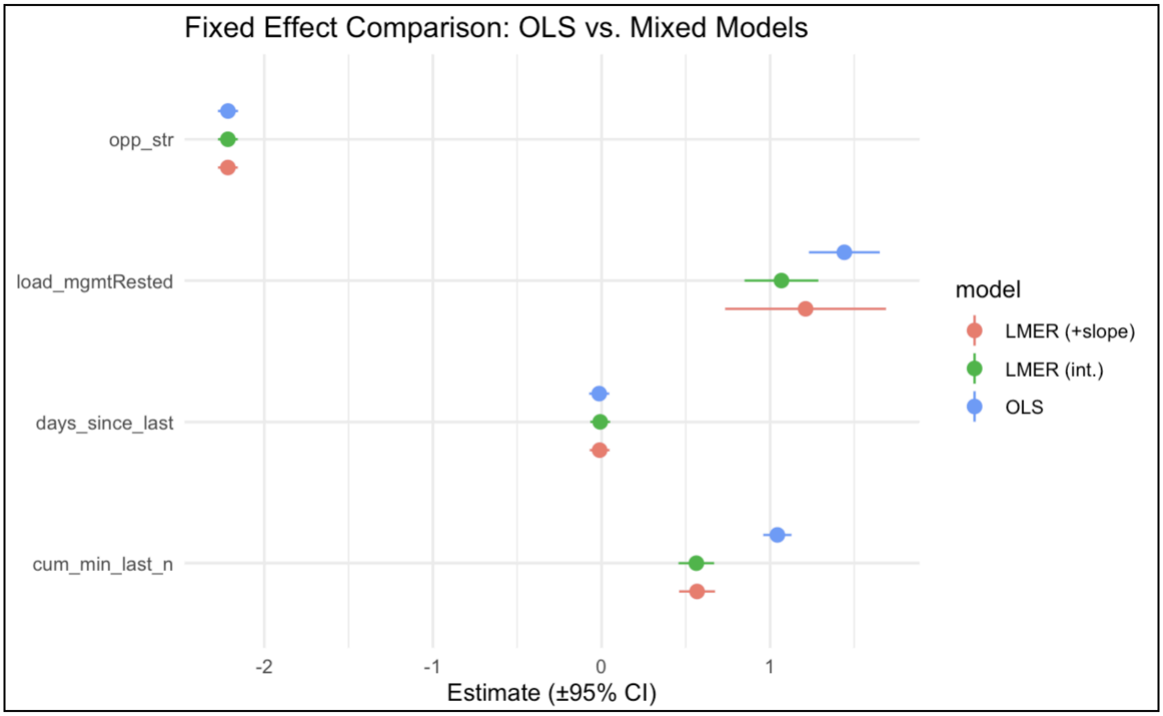

In our final random-slopes model, cumulative minutes over the past three games remains a robust positive predictor: a one-SD increase in workload corresponds to a +0.57-point gain in plus/minus (95% CI 0.46-0.68), compared with +1.04 under OLS and +0.56 in the intercept-only model. Over a typical 48-game season, this effect would accumulate to a swing of roughly 27 points—enough to flip several close outcomes and impact overall team success. The flatter slope under LMER illustrates how partial pooling tempers extreme OLS estimates (see figure below).

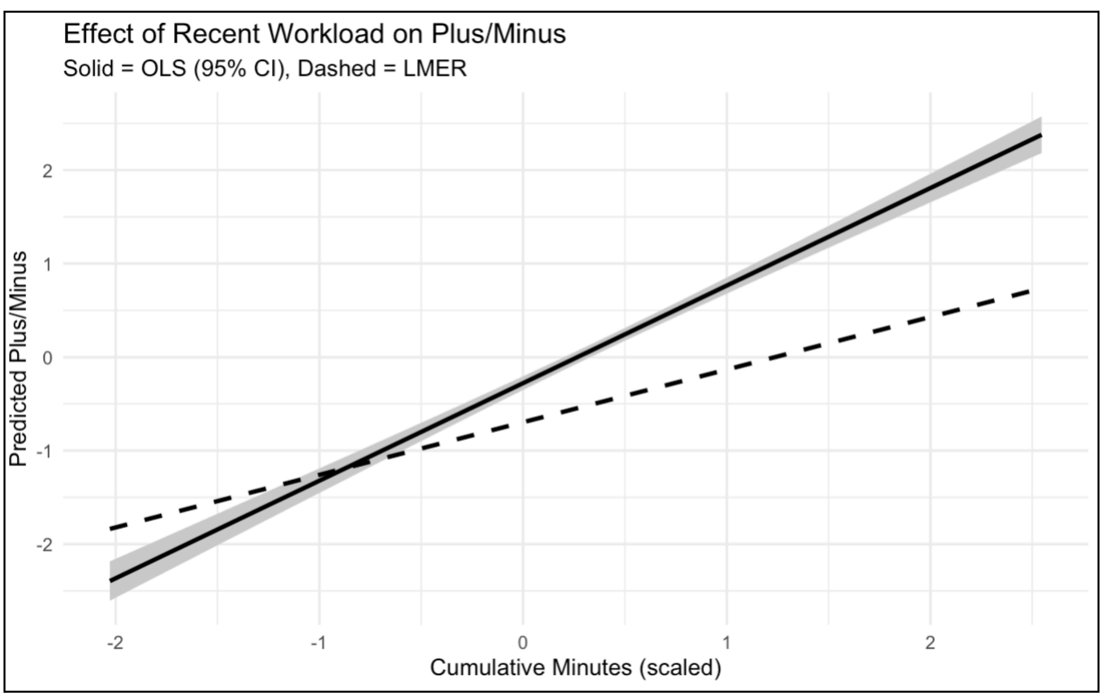

Turning to rest, the load_mgmt indicator (“Rested” vs. “No Rest”) delivers a +1.21-point boost in plus/minus (95% CI 0.73-1.69) in the random-slopes model, down from +1.44 in OLS and +1.07 in the intercept-only model. This shrinkage again shows that OLS slightly overstates the rest effect when failing to account for which player and team are resting (see figure below). A consistent +1-point lift from rest, if applied strategically, could translate into improved rotation decisions and marginal win gains in tight contests.

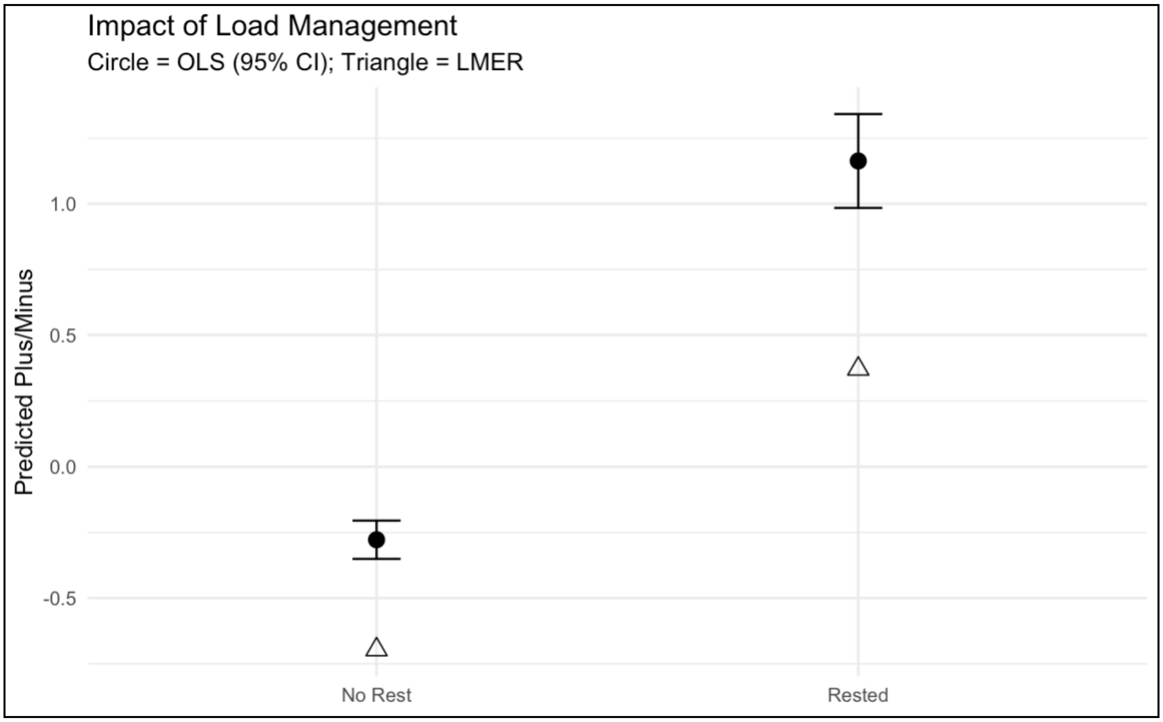

Finally, variance‑component estimates from the intercept‑only LMER reveal athlete‑level standard deviation ~1.60 points and team‑level SD ~1.18 points, compared with a residual SD of ~10.99. In practical terms, individual baselines vary by roughly ±1.6 points, teams by ±1.2, while most game‑to‑game variation remains driven by unmodeled factors. Overall, our linear and hierarchical analyses demonstrate that both moderate workloads and deliberate rest yield reliable, actionable gains in player plus/minus, even as substantial unexplained variability underscores the need to combine these insights with game-specific context and other performance factors.

IV. 2. Decision Tree Results

The XGBoost model was trained on the engineered features to predict player plus/minus. Evaluation on the test set yielded RMSE of 10.21, MAE of 7.87, and R^2 of 0.11. These metrics indicate somewhat low performance, but still non-trivial predictive power.
To understand the contribution of each feature to the model's predictions, we employed SHAP analysis. The SHAP summary plot below reveals the global feature importance.

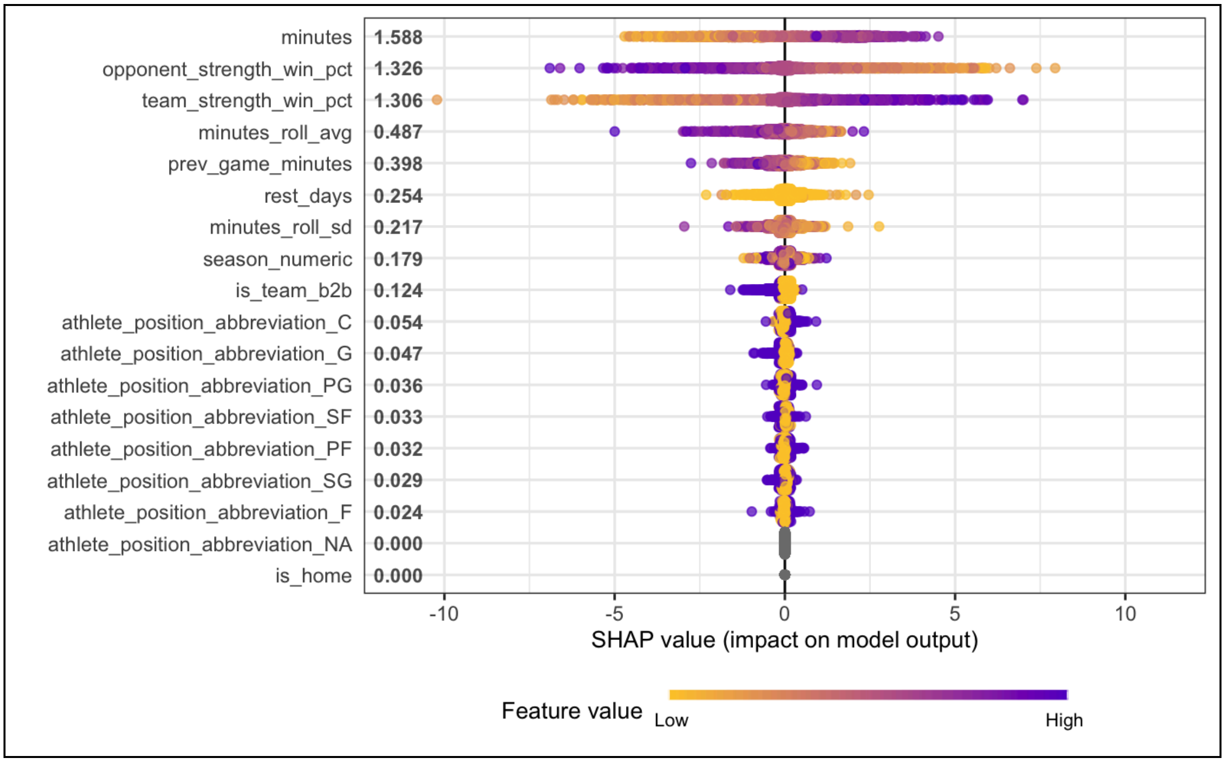

We found that minutes (current game minutes) showed the highest overall importance, followed by opponent_strength_win_pct, team_strength_win_pct, minutes_roll_avg (5 game rolling average minutes), and prev_game_minutes (previous game minutes). Notably, the raw rest_days feature itself ranked relatively low in importance.

SHAP dependence plots further illuminated the nature of these relationships. As shown in these figures, higher values for minutes_roll_avg and prev_game_minutes generally corresponded to negative SHAP values.

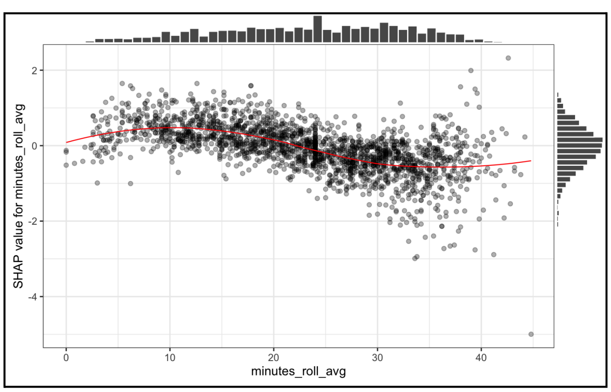

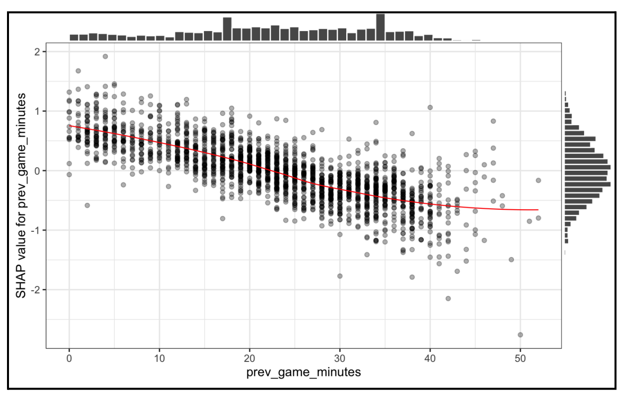

This indicates that sustained high playing time or high minutes in the immediately preceding game tended to decrease a player's predicted plus/minus, quantitatively supporting the idea that fatigue or heavy usage can negatively impact performance within this model's framework. Conversely, the SHAP values for the features representing discrete rest day counts were clustered near zero across their range, reinforcing the finding that the number of rest days had minimal predictive power in the XGBoost model compared to workload and game context variables.

V. Discussion

The results of the linear models and the decision tree allow us to examine the relationship between rest and player performance as measured by plus/minus. We observe that rest days as a predictor variable are not shown to have a significant impact on plus/minus by player, as rest days was found to not be statistically significant throughout the linear models in addition to its low feature importance in the decision tree. Instead, player statistics related to minutes had greater predictive power in relation to a player's plus/minus. Our findings suggest that NBA teams should focus on managing their player workload through considering the minutes that players play as opposed to rest days between player appearances.

The linear mixed effects model shows a positive relationship between a player's plus/minus from their cumulative minutes over their past three games as well as from their rested load management indicator. The XGBoost decision tree model suggests that high minutes in previous games can negatively impact player plus/minus, as measured through a player's three game minutes rolling average and their previous game's minutes. This difference in findings can be due to the different modeling techniques. The linear models assume a more straightforward relationship between the variables, while the decision tree may be capturing more complex, non-linear relationships in the variables. Additionally, where cumulative minutes can be correlated with star players in the linear model, the decision tree's inclusion of current minutes in its features may help control for the effect of star player usage, providing insight into how excessive minutes in recent games can negatively impact performance. This apparent contradiction highlights the complex relationship between playing time, rest, and performance that may not be fully explained in linear models.

Our approaches have several limitations that provide opportunities for additional research. One limitation in the implementation of the linear models is the focused variable selection in rest related metrics. The features used in the linear models relate primarily to player usage, with an additional variable considering the opposing team strength. Including additional player performance metrics that relate to their offensive or defensive abilities could provide further explanation for variability in plus/minus that is not currently explained. This may also enable a more-focused analysis that stratifies the player pool into buckets, allowing insight into how rest may affect performance differently for players with different roles.

Another potential limitation across the various models used to perform this analysis is the response variable of plus/minus. Player plus/minus values can include noise from teammate and opponent contributions, and it may not capture a complete image of a player's performance. Further analysis could be conducted on specific performance metrics such as true shooting percentage or player efficiency rating to assess whether rest affects specific elements of a player's abilities. In this sense, while rest days may not affect player plus/minus as a whole, it may have an impact on player abilities and can be utilized in strategic circumstances where specialized types of play are emphasized.

Our analysis provides insights into the relationship between rest and NBA player performance. Future research should expand on this work in a few directions. Firstly, stratifying players by age, position, and usage patterns could likely reveal more nuanced rest effects, as veterans or starting lineup players may benefit differently from workload management. Secondly, examining different performance metrics beyond plus/minus, such as shooting efficiency, clutch-time performances, and defensive metrics, could identify specific basketball skills that react differently to fatigue. Analysis in these areas could build on our findings and work towards developing practical load-management strategies, usable by NBA players and teams for better performances and winning outcomes.