How Do Injuries in the NFL Affect The Outcome of the Game

Andy Sun

Background:
Jenny Vrentas from the MMQB points out, "the NFL's injury surveillance data shows a slow upward trend in the total number of injuries sustained in all practices and games from 2004 to 2012." In a sport as physical as football, how much is a team affected by a loss of a player due to injury? The Packers famously won the super bowl in 2010 with 15 players on the injured reserve, but where is the breaking point?

Abstract:
In this study, we attempt to determine two main outcomes: how injuries sustained during the week of the game changes the initial predicted probability of winning, and how influential are sustained injuries relative to various other statistics used to predict team success.

To measure how injuries can change the predicted probability of winning, we look at the predicted point spread of favors teams, the outcome of the game, and the injuries sustained by players on both teams.

To measure how influential injury statistics are in terms of predicting team success, we use Football Outsiders metric of Adjusted Games Loss (AGL), which is a metric that looks at team's injury reports over a season and determines how much that team was affected adversely by injuries. We use the AGL as one of the many features in a feature vector that we then pass to a predictive model, and use feature selection to rank the features. By observing how well AGL ranks, we can determine its importance.

Data

Source	Data
Footballlocks.com	Historical betting odds and point spreads for all NFL teams during all past weeks dating from 2006 to present day
Wizardofodds.com	Historical data from 1994 – 2013 mapping point spread to predict the probability of winning.
Nfl.com	Weekly injury reports, and statistics that can be used to predict team success such as power rankings, yards gained, etc.
Footballoutsiders.com	Adjusted Games Loss metric data for the 2013 and 2012 NFL seasons

A) How injuries sustained during the week of the game changes the initial predicted probability of winning

A.1) Methodology

We can analyze how the predicted probability of the favored team winning changes based on injury by first looking at the spread of the favored team and the actual outcome of the game.

Let X = whether or not the favored team won -> 1 if favored team won; 0 otherwise.

Let p = predicted probability of the favored team winning.
Where p is calculated by converting spread using the chart provided by wizardofodds.com. See Appendix.

Let Y = Difference in Winning Probability = X – p

The value of Y is positive if the favored team wins. The smaller the value of Y, the more heavily favored the team was to win.

The value of Y is negative if the favored team loses. The more negative the value of Y, the greater the upset.

Next, we look at weekly injury reports for all of the games one week at a time and record the number of injuries on offense, and the number of injuries on defense.

We can then use this data to look at the difference in winning probability based on bins, where each bin is a different count of types of injuries each team sustains. For example, we look at the difference in winning probability only when the number of offensive injuries of the favored team is greater than the total number of injuries the defensive team sustains. We then take the averages of each bin.

By applying this to the data from the 2013 – 2014 NFL season, we have results.

A.2) Results and Discussion

Bin	Avg Diff in Winning Probability
fav_off > underdog_off	-0.117
fav_off <= underdog_off	0.594
fav_off > underdog_def	0.12
fav_off <= underdog_def	-0.079
fav_def > underdog_off	0.025
fav_def <= underdog_off	0.0195
fav_def > underdog_def	0.1197
fav_def <= underdog_def	-0.186
fav_off > fav_def	-0.1346
fav_off <= fav_def	0.05
fav_total > underdog_total	0.0676
fav_total <= underdog_total	-0.075

Ultimately, the results are contradictory. For instance, if we look at results where the average difference in winning probability is negative, meaning the favored team lost, some results intuitively make sense, while other equally statistically valid results say the direct opposite

The results show that one of the most negative average difference in winning probabilities (most negative meaning the favored team loses in an upset) occurs when the offense of the favored team has more injuries than the opposing offense. This intuitively says that the offense of the favored team may be hindered by injuries and is unable to put points on the board.

Another one of the most negative average difference in winning probability occurs when the favored team overall has less injuries than the underdog team, and lost.

These two results contradict each other and make it very difficult to come to a definitive conclusion.

If we look at the smallest positive average differences in winning probability, we can see the same contradiction.

A limitation to this model is the small sample size of inputs due to the high difficulty in gathering data. In the future, we may be able to get more accurate results if we start recording injury statistics even more commonly and cleanly than we do now. This will lead to both more data, and less preprocessing.

B) How influential are sustained injuries relative to various other statistics used to predict team success.

B.1) Methodology

In sharp contrast to the above model where it was difficult to mine data so not much data was used, this second approach uses machine learning as a "brute force" approach to look at lots of data at a time to determine the influence of injury.

We use a binary classification tree to classify whether or not a team will make the playoffs, where making the playoffs equates to team success. We build this model by training on feature vectors that consist of the Adjusted Game Loss rating for each team, followed by overall team power rating, overall offensive performance by team such as rushing yards and receiving yards, overall defensive performance, and then by starter performance statistics specific to each position.

Then we use feature selection to extract the most influential features that played a key role in helping the predictive model classify. A binary classification tree was chosen as the model since it handles mult colinearity well, meaning that the model iis not as heavily hindered by features that are not completely independent.

In order to form the nodes of the decision tree, we use Gini's Diversity Index to measure how pure each node is. The more pure the node, the more influential it is.

B.2) Results and Discussion

We can rank each of the features by Gini's Diversity Index, and Adjusted Game Loss was so non important that it was selected out by the decision tree, meaning that it was not a feature used in the decision process.

Non-surprisingly, the most important features were features to quarterback and running back performance, as well as defensive sacks.

In future studies, rather than having a majority of traditional NFL statistics thrown in the feature vector, we would use a set of features that make intuitive sense for measuring injuries. For example, in this model, we try to predict success, so we pass in features pertaining to success and try to determine if AGL was important to team success.

Another way to look at the model would be to try to predict the AGL, using features that we think would be important.

References

"Betting the NFL." - *Wizard of Odds*. N.p., n.d. Web. 12 Dec. 2014.

"Football Outsiders." *Football Outsiders Everything*. N.p., n.d. Web. 11 Dec. 2014.

"The NFL Injury Report | NFL.com." *The NFL Injury Report*. N.p., n.d. Web. 11 Dec. 2014.

"NFL Odds: Las Vegas NFL Football Odds from Week 1. Pro Football Betting Odds." *Week 1 NFL Odds: Las Vegas NFL Football Odds from Week 1. Pro Football Betting Odds.* N.p., n.d. Web. 12 Dec. 2014.