Quantifying Jay Wright's Greatness

Recently named the college men's basketball coach of the decade, Villanova coach Jay Wright has gained much media attention in the past several years. Leading the Wildcats to many victories including four conference tournament championships, seven conference regular season titles, and two NCAA national championships, Coach Wright has duly earned himself several coaching awards. Nevertheless, these awards are typically based on voting, a process which is inevitably subjective. In order to completely eliminate personal biases from these awards, a quantitative method of comparison for college men's basketball coaches is necessary, but would Jay Wright still have won coach of the decade with a quantitative analysis, rather than a voting process? This project analyzes basic stats such as win-loss percentage and Elo ratings as well as more advanced stats such as pre-game and in-game win probabilities to quantitatively answer the question 'Is Jay Wright the best coach in college basketball?'

Introduction

With the recent success of the Villanova men's basketball program, media and sports fans have turned much of their attention to the Villanova men's basketball team and, more specifically, to the Wildcats' head coach, Jay Wright. Recently named the NCAA men's basketball coach of the decade, Coach Wright is undoubtedly an accomplished coach and with two national championships within three years, he has certainly earned himself this attention. As a community rooted in pride and unity, Villanovans are proud of their men's basketball coach and often boast that they have "the best" coach out there, but how is that measured? This project aims to answer the question, 'Is Jay Wright the best coach in college basketball?' using statistics and data science to measure how Coach Wright stacks up against the great men's college basketball coaches throughout history. This begs the question, however, of how does one quantify what it means to be a great coach?

Methods

For this analysis, NCAA men's basketball data from http://www.sportsreference.com is used; this data includes statistics from the early 20th century until present. Several datasets for this project were scraped from the site directly, and additional datasets were downloaded from Kaggle (https://www.kaggle.com/c/mens-machine-learning-competition-2018). This data downloaded from Kaggle, however, was also scraped from http://www.sportsreference.com and was used to eliminate several steps of scraping in this project. Statistics for the 2019-2020 season have been omitted for this project since post season tournaments were never completed and, thus, several data points are missing. For any seasons with more than one coach listed as the head coach of a school, data will be counted for both coaches since this data cannot be split into each coach's individual contribution and omitting these years disregards crucial data of a coach's starting/ending statistics.

In this analysis, 2,057 unique coach-school dyads will be analyzed. A "coach-school dyad," by this definition, is a distinct combination of coach and school, i.e. Jay Wright - Villanova and Jay Wright - Hofstra are treated as two separate dyads. In some sections, several of these 2,057 dyads are excluded from the analysis if that data is not available for those coach/school combinations. Any coach with less than five years of experience is also omitted from this analysis as their data is less reliable; this is because most coaches tend to decrease in success in their first two to three years and then, if they are going to succeed as a coach, start to show this in the data around five years in. Every season in this analysis is represented by its latter year for simplicity in graphs and charts; for instance, the 2018-2019 season is represented as the year 2019.

All methods discussed in this paper have been applied to all coaches, however, since displaying 2,057 graphs of data is not only unnecessary but also nearly impossible, only some coaches have been selected for visualizations. These coaches include: Jay Wright, Rollie Massimino, Mike Krzyzewski, John Wooden, Bob Knight, Jim Boeheim, Charlie Schmaus, and Jerry Loyd. The inclusion of Jay Wright in these graphics is obvious, and as his mentor, Rollie Massimino is a fair choice to compare and contrast in this analysis. Mike Krzyzewski and Jim Boeheim are included as examples of well-known present day basketball coaches, while John Wooden and Bob Knight are included as famous, great basketball coaches of the past. Finally, Charlie Schmaus and Jerry Loyd are included in this analysis as examples of poor coaches, Schmaus ending his six year career at Virginia-Military-Institute with merely one win out of twenty six games played, and Loyd bringing his team down 64 percentage points from his first coaching year to his last. Schmaus and Loyd are included here to demonstrate how poor coaching stats would look on a graph as compared to great coaching stats.

Questions of Interest

The questions that will be answered in this paper include:

1. Was there an increase in team success throughout the coach's career, indicative that the coach had an impact on the team and that the team was not merely great to begin with?
2. What is the relationship between performance and familiarity with the opposing team for great coaches? A good coach would prepare for the specific team they're facing to the best of their ability, but would be able to do this better with more familiar (i.e. in-conference) teams.
3. How do coaches typically perform compared to their expected probability of winning a game? Does each coach perform better when they have a high probability of winning a game or a low probability?
4. How much does each coach increase their Elo rating throughout the season? This measure will be helpful in distinguishing any up and coming coaches that may not start with a good team, but grow their team throughout the year.
5. How quickly does a coach typically lock in a win for the game? Once a coach knows he has a high probability to win, he is likely to put in his subs to give his starters a break; the quicker a coach reaches this winning threshold, the better prepared his team is for their matchup.

Increase in Team Success

The win-loss statistic is a measure of the number of wins for a given team divided by the total number of games that team has played. Therefore, the higher the percentage, the more success a team has. The overall win-loss percentage statistic is one of the simplest ways to measure team success for any given season, thus, this analysis will begin with analyzing this. A common plot that will be referenced in this analysis is a plot of win-loss percentage over time for each coach-school dyad, therefore, it is important to be familiar with the plot below.

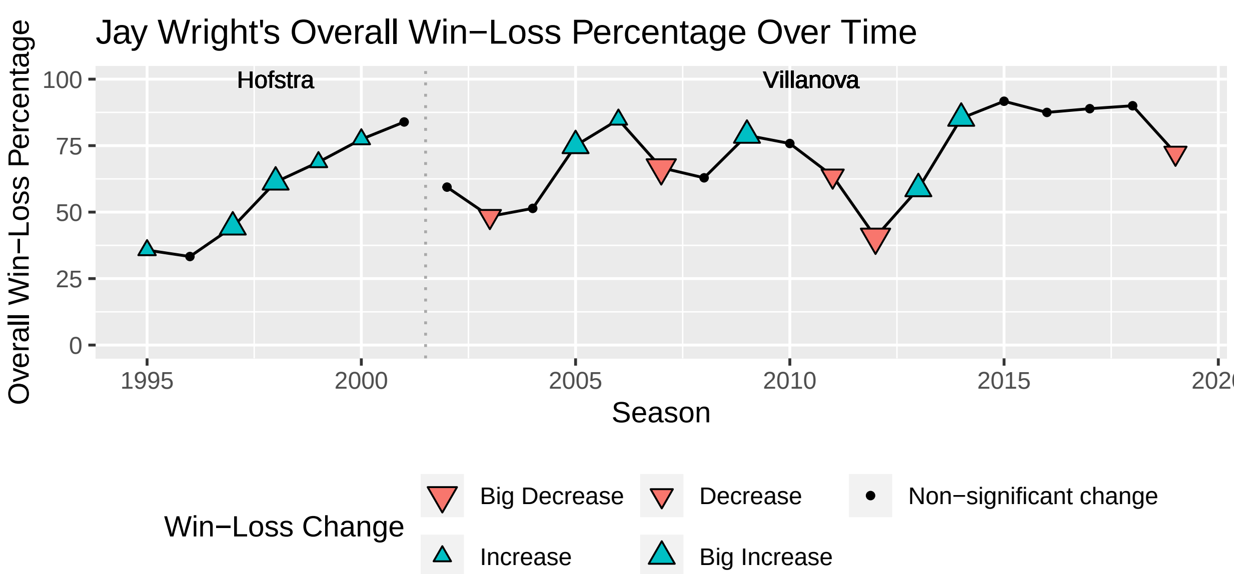

In this graph, Jay Wright's overall win-loss percentages for each season he coached at both Hofstra and Villanova are plotted on the vertical axis, with the corresponding season plotted on the horizontal axis. Also included in this graph are symbols to indicate the change in overall win-loss percentage from season to season. A pink downward-facing triangle indicates a decrease of 10 percentage points or more in overall win-loss percentage from the previous year while a large pink downward-facing triangle indicates a decrease of 20 percentage points or more in overall win-loss percentage from the previous season. Conversely, a small or large blue upward-facing triangle indicates an increase of 10 or 20 percentage points or more in overall win-loss percentage from the previous season, respectively. A black dot indicates a non-significant change in overall win-loss percentage from the previous season of less than 10 percentage points. It is interesting to note that in this plot, Villanova's overall win-loss percentage seems to dip around 2012 and suddenly increase again, likely related to the team's acquisition of many now well-known players who have made it to the NBA such as Ryan Arcidiacono, Daniel Ochefu, and Darrun Hilliard. What is also interesting and rather important to note is that while Jay Wright is more known for his coaching career at Villanova, his performance at Hofstra was rather impressive. This will come into play later in the analysis. Similar plots for all coaches analyzed in this paper are included below with a gray dotted line indicating a change in school for that coach.

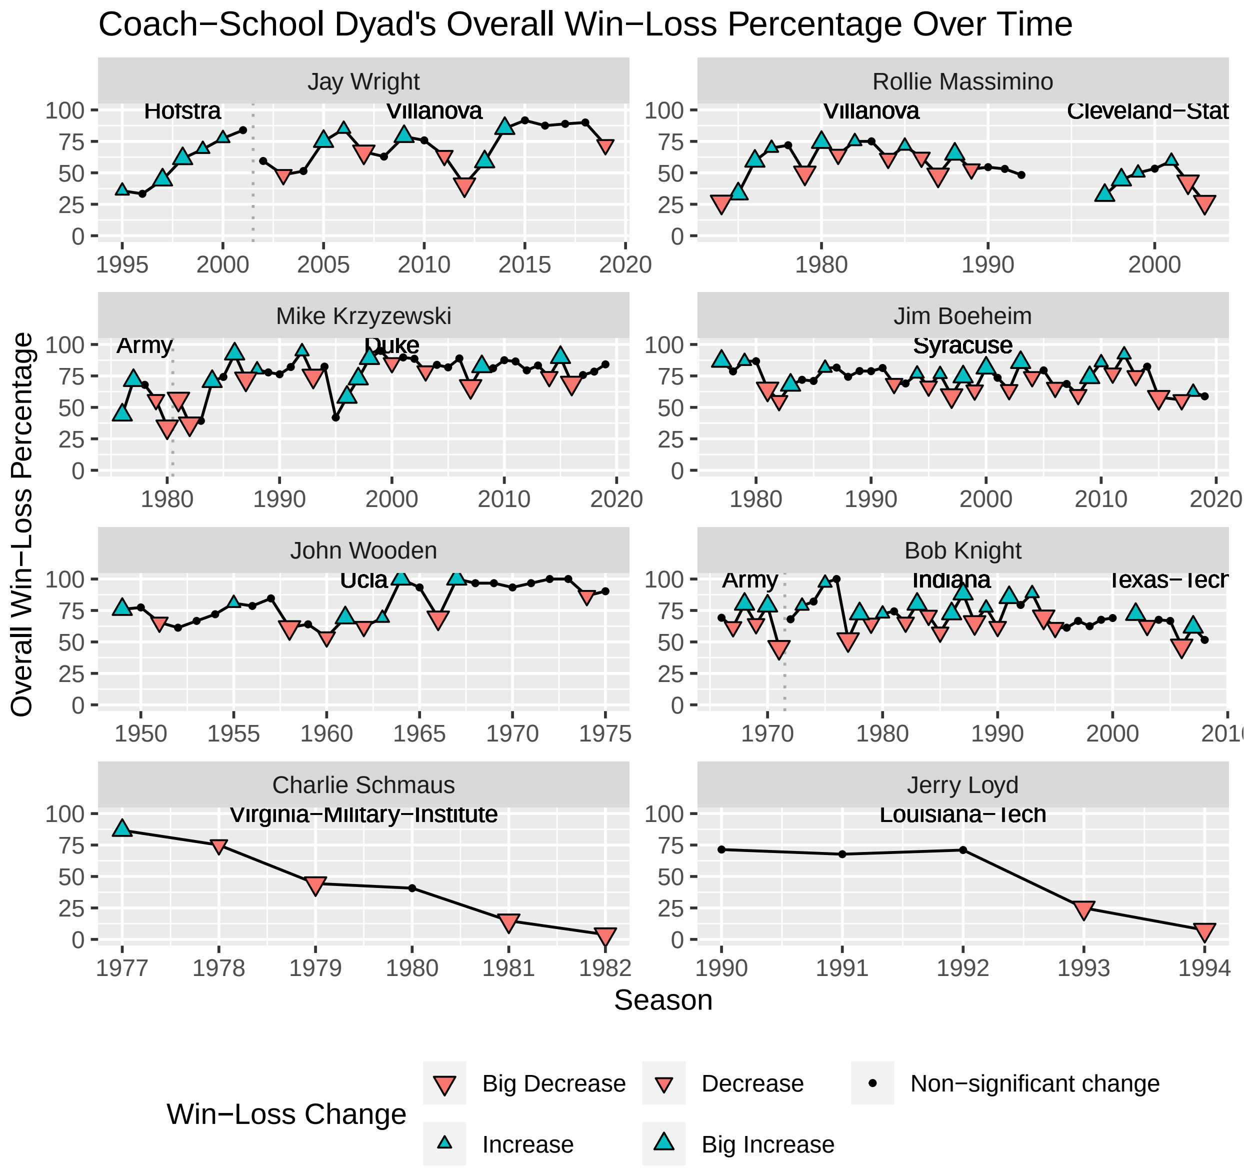

It is important to note in this plot that Rollie Massimino's graph shows an odd gap from the years 1992 to 1996 even though he did not take a break in coaching; this is when it is important to remember that coach-school dyads that lasted less than five years are omitted in this analysis. Coach Massimino, in fact, coached at six different schools, but only two of these schools are included here because of this requirement; this is the case for many coaches in this analysis. These plots will be analyzed further, but this graph is included for reference.

Percent Increase in Overall Win-Loss Percentage Throughout Career

Needless to say, a great coach would show growth throughout his or her career. Thus, recall the familiar plot of win-loss percentages, and now consider the difference in overall win-loss percentage from any given coach-school dyad's first year to final year calculated by subtracting the y-value of the first season at each school from the y-value of the last season included on this plot. This value for Coach Wright is represented by the vertical dotted line on the plot below. If this difference is then divided by the coach's first season overall win-loss percentage and multiplied by 100, the computed value represents the overall win-loss percentage percent increase for each coach-school dyad in this analysis. For instance, for Jay Wright at Villanova, this value is

$$\frac{72.2 - 59.4}{59.4} * 100 = 21.549\%$$

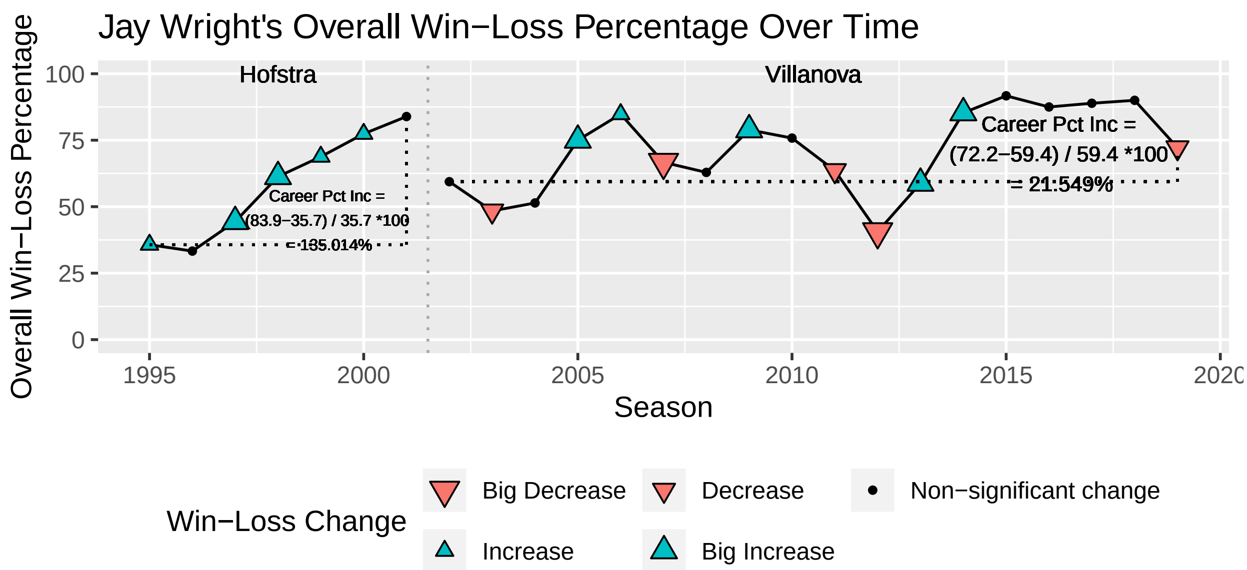

This value can then be calculated for each coach-school dyad and graphed as shown below with overall win-loss percentage percent increase on the horizontal axis and coach-school dyad on the vertical axis.

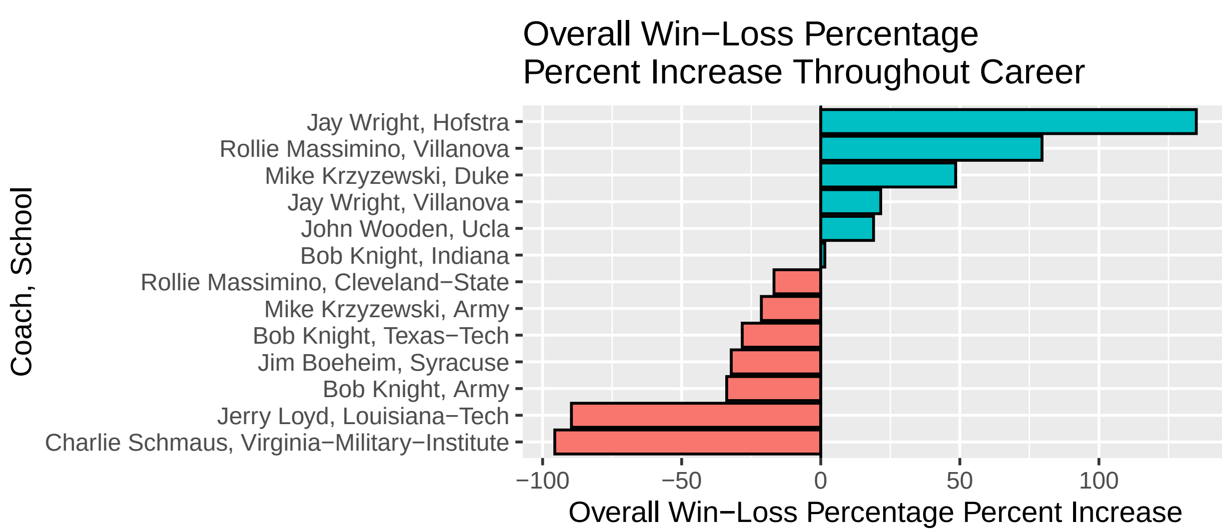

In this graph, we easily see overall win-loss percentage percent increase for every coach-school dyad and the bars are colored to indicate an increase or decrease in team success throughout the coach's career at that school. Unfortunately, this graph inaccurately depicts the data because teams with lower original overall win-loss percentages leave more room for improvement for new coaches than already-successful teams do. For instance, in Jay Wright's first season at Hofstra, the Pride men's basketball team had an overall win-loss percentage of only 35.7%, thus, Jay Wright had much room for improvement. At Indiana, however, the men's basketball team finished with an overall win-loss percentage of 67.2% in Bob Knight's first season, almost twice that of Hofstra in Coach Wright's starting season, leaving much less room for growth. Therefore, this value must be standardized to be accurately compared.

Standardizing this value can be accomplished with the following formula:

$$Standardized\,WL\,\%\,Inc = \begin{cases} \frac{Final\,WL\,\% - Initial\,WL\,\%}{100 - Initial\,WL\,\%} * 100 & when\ Final\,WL\,\% > Initial\,WL\,\% \\ \frac{Final\,WL\,\% - Initial\,WL\,\%}{Initial\,WL\,\% - 0} * 100 & when\ Final\,WL\,\% \leq Initial\,WL\,\% \end{cases}$$

This formula essentially calculates the potential reached towards increasing the team's overall win-loss percentage to the maximum win-loss percentage, 100%, for coaches whose overall win-loss percentage increased over time and, conversely, calculates the potential reached towards decreasing the team's overall win-loss percentage to the minimum win-loss percentage, 0%, for coaches whose overall win-loss percentage decreased over time. If this number is calculated for all coaches and this bar graph is replotted with these new, standardized values in the same order, the graph below, which better depicts the percent increase of each coach-school dyad's overall win loss percentage, is created.

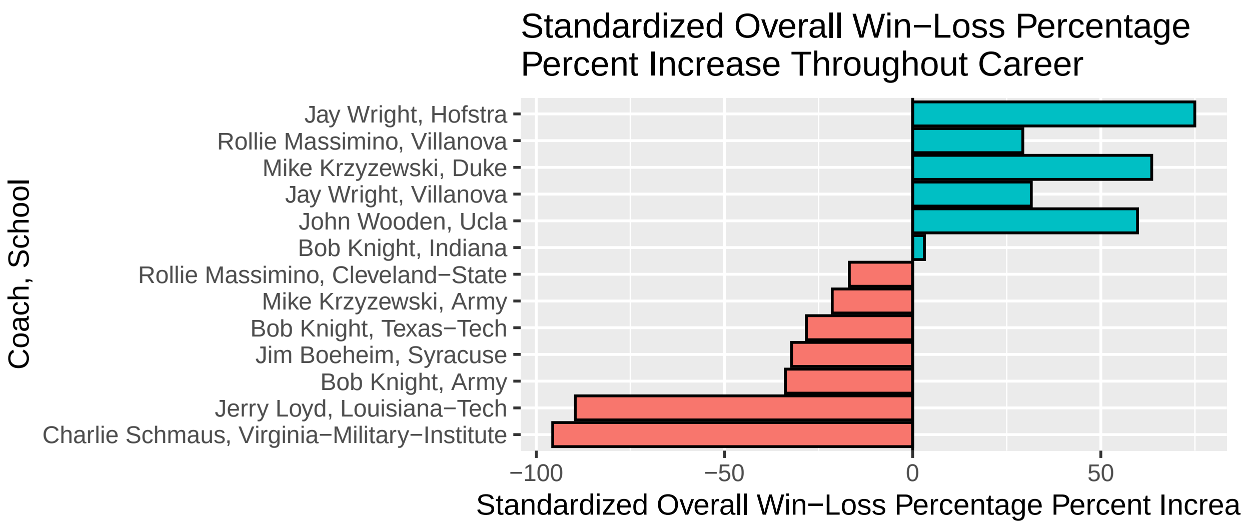

This bar graph is more informative than the previous graph; here we see that while Coach Massimino increased his overall win-loss percentage at Villanova by a greater amount than Coach Krzyzewski did for his respective overall win-loss percentage at Duke, indicated by the order of the bars, Coach Krzyzewski has reached more of his potential to increase during his time at Duke, indicated by the size of the bars. Since one downside to this visualization is that only a handful of coaches can be effectively analyzed at once with this graph, we plot these same values on a density plot and include dotted lines to indicate the coaches compared in this analysis for reference.

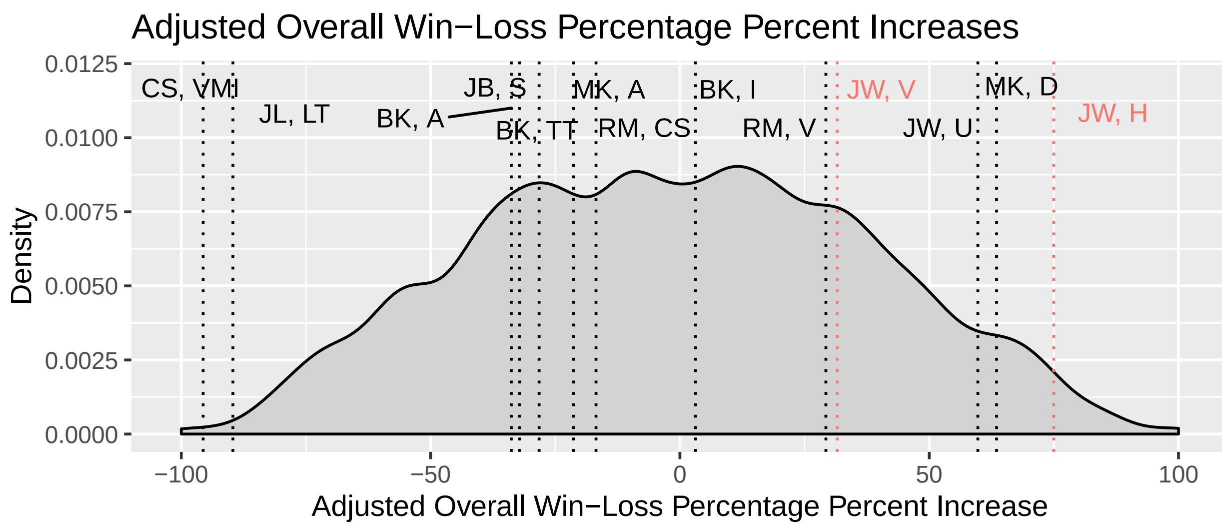

With this graph we can easily see the distribution of the adjusted overall win-loss percentage percent increase, which is roughly normal. Note that these lines are in the same order as the bars in the previous bar graph because the same data is plotted in these two plots.

Several different statistics will be analyzed in this paper; since combining different statistics would be rather difficult, these values will be converted to percentiles to easily compare all coach-school dyads simultaneously. A percentile can be calculated from this potential reached value by finding the mean and standard deviation of this standardized overall win-loss percentage percent increase, then using these values to calculate z-scores for every coach-school dyad. Afterwards, these z-scores can be used to calculate the percentile of every coach-school dyad in this analysis. The results of this calculation are shown in the graph below where the percentiles for the coaches compared here are indicated with dotted lines labeled with the initials of the coach and school. Jay Wright is represented by the colored dotted lines.

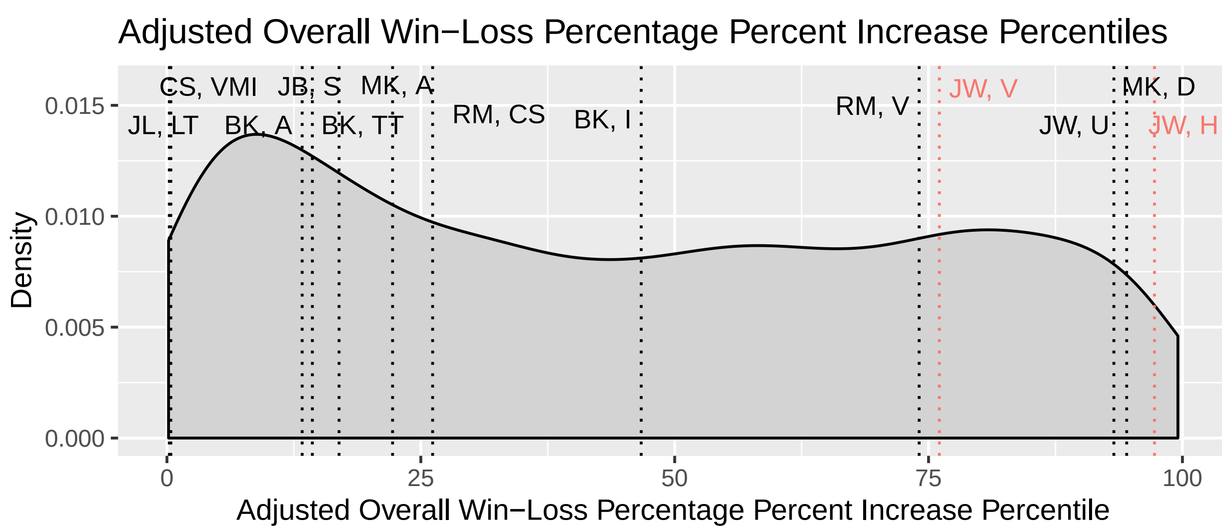

We see in this plot that Jay Wright is in approximately the 97th percentile for his success at Hofstra and the 76th percentile for his success at Villanova according to this measurement of standardized overall win-loss percent increase. This means that Jay Wright reached more of his increase potential than 97 percent and 76 percent of his colleagues, respectively. Note that the data is not normal; the plot peaks around the 15th percentile, approximately. We see from this plot that, according to this analysis, even renowned coaches such as Bob Knight and Jim Boeheim have lower percentiles than would be expected.

Slope of Overall Win-Loss Percentage Throughout Career

Again, consider the graph below which plots Jay Wright's win loss percentage at both Villanova and Hofstra. The dotted line represents a linear model of overall win-loss percentage and season, also known as the slope. The slope of the linear model of this plot reflects how rapidly the coach was able to improve his or her team's success; a higher slope indicates quicker growth.

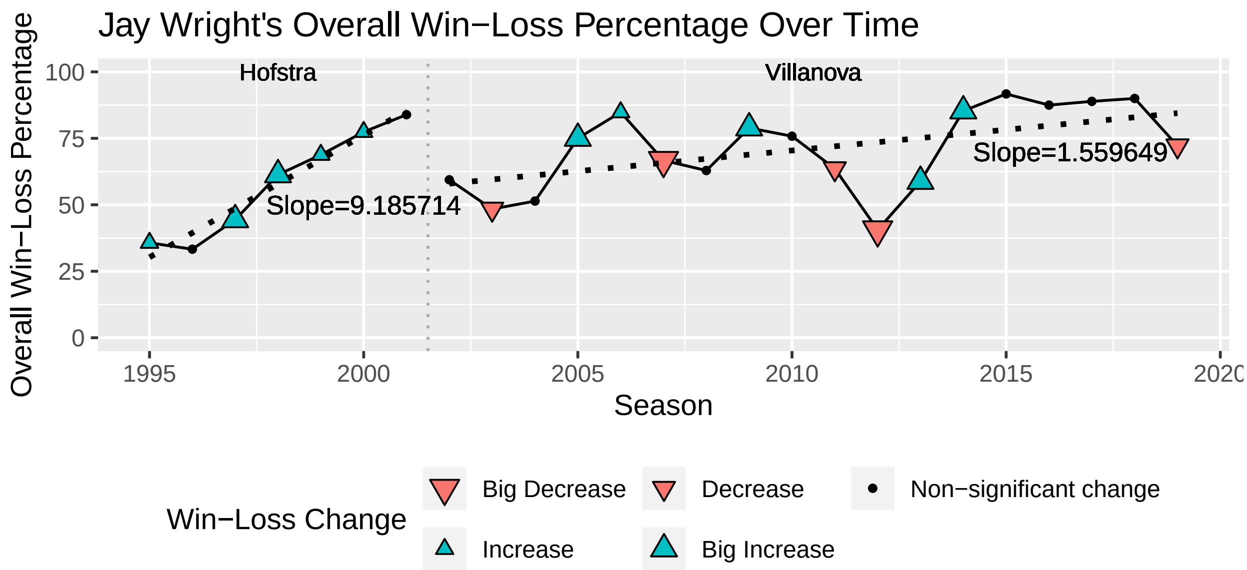

We see here that Jay Wright had a slope of 1.559649 for his performance at Villanova and a slope of 9.185714 for his performance at Hofstra. Therefore, it would appear that Jay Wright was a better coach at Hofstra than he is at Villanova. That does not make too much sense, however. An important note about these slopes can be found in a plot of the overall win-loss percentage slope versus coaching years, as shown below.

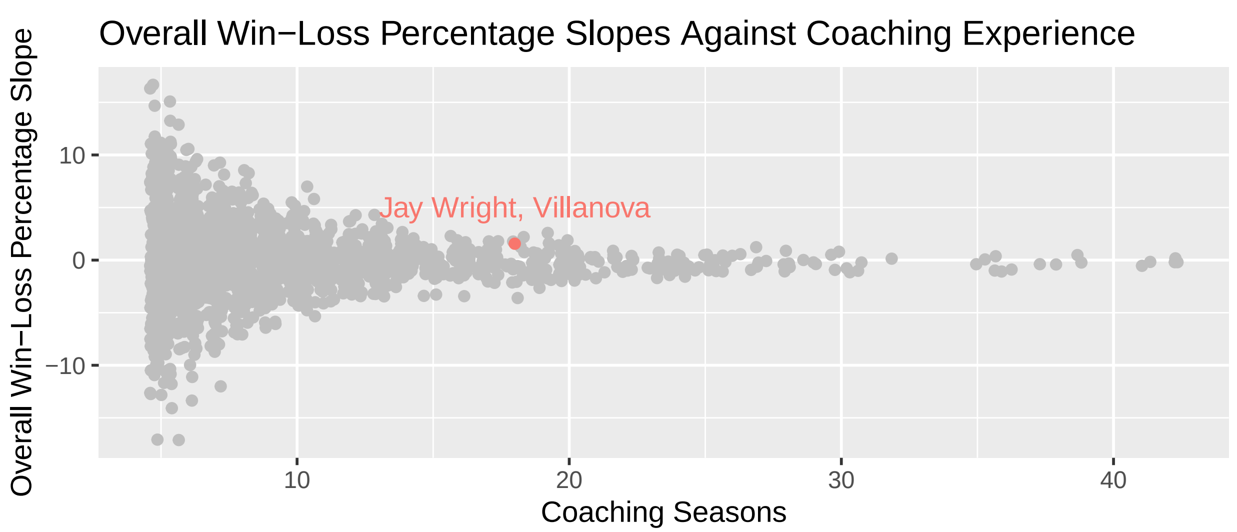

In this plot, a strong relationship between overall win-loss percentage slope and coaching years is clearly evident in the shape of this graph, almost resembling a trumpet. It is clear here that the more years a coach has been at a single school, the closer the slope of his or her overall win-loss percentage plot will be to zero. This plot is essentially showing a regression towards the mean, zero. Therefore, Jay Wright's lower percentile for his time at Villanova than at Hofstra is to be expected since Coach Wright has coached at Villanova approximately 2.5 times longer than he did at Hofstra. This is important to remember because it tells us that some of the most well-known coaches may have an apparently low percentile by this standard because of their longer coaching careers.

A potential solution to this problem would be to group coaches by coaching years and calculate percentiles within these groupings. Doing this would prevent any short-term coaches from outshining longer-term, arguably better coaches because the slopes of these coaches would not be compared to one another. By doing this, Jay Wright's performance at Hofstra would not be compared to his performance at Villanova because of the difference in years spent at each school. To see what effect this solution would have on these calculated percentiles, coach-school dyads were grouped into one of three categories based on how long each coach was the head coach of his school's men's basketball team: 5-9 years, 10-14 years, and 15+ years. Then, the percentile for overall win-loss percentage slope was recalculated within these groups and is displayed in the graph below.

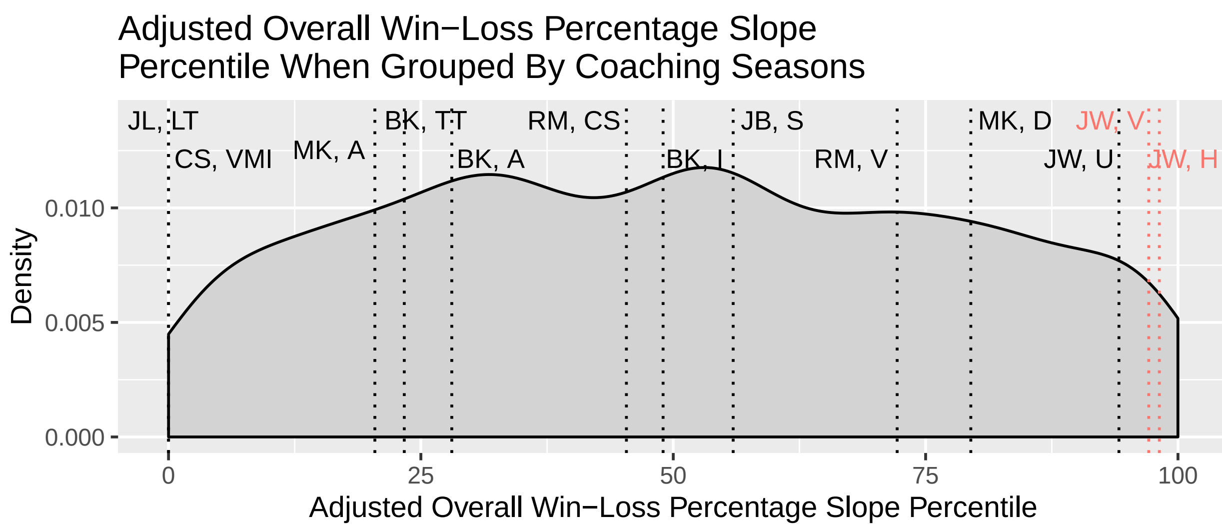

Still, Jay Wright is in the upper percentiles for his performance at both Villanova and Hofstra, but we see the advantage of grouping by coaching years. The percentiles for all selected coaches have shifted because this tendency for overall win-loss percentage slopes to approach zero as years increase is accounted for, and we see that this data is more uniform. We also notice that most of these "great" coaches have benefitted from this grouping; John Wooden, Mike Krzyzewski and Rollie Massimino all shifted from approximately the 50th percentile to the upper percentiles, where their percentiles would be expected to be. Jay Wright's percentile for Villanova also increased quite a bit compared to before; he is now in the 97th percentile. This makes a lot more sense than before. Coach Wright did not get worse when he came to Villanova; he merely stayed longer.

Preparation for In-Conference Play

Similar to the overall win-loss percentage statistic, the in-conference win-loss percentage statistic measures a team's performance against teams in the same conference. These teams tend to be closer in skill to each other and are fairly familiar with one another from years of playing each another. Thus, this statistic is just as important to analyze, if not more important, since game preparedness certainly reflects coaching skill. Similar methods from the previous analysis will be used to analyze this statistic since these measurements are extremely similar; in fact, the data used for these analyses are merely a subset of the data previously used. It is important to note that conference win-loss percentage tends to be higher than overall win-loss percentage for most coaches.

In this graph, which will be used for both sections of the in-conference analysis, season is plotted on the horizontal axis and conference win-loss percentage is plotted on the vertical axis. The same symbols from before are included to indicate the change in conference win-loss percentage from season to season. A sample plot for conference win-loss percentage versus season for Jay Wright at Villanova is shown below, including the dotted lines indicating conference win-loss percentage percent increase and slope.

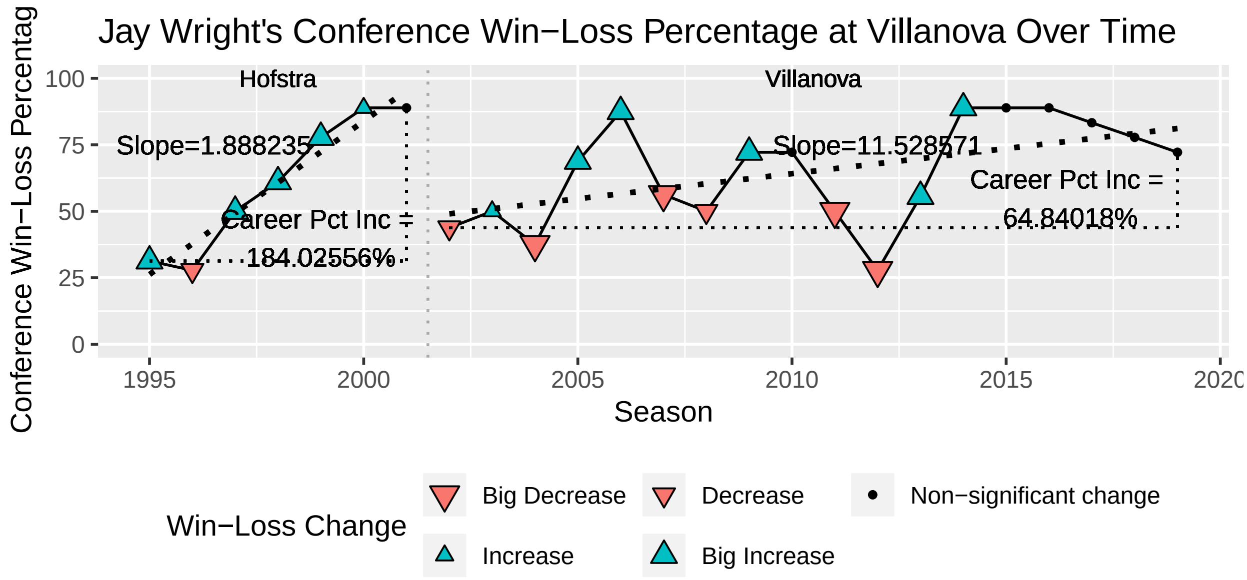

Percent Increase in Conference Win-Loss Percentage from Beginning to End of Career

Again we calculate the win-loss percentage percent increase, this time for only conference play. This is accomplished by dividing the difference in conference win-loss percentage for the final season minus the conference win-loss percentage for the first season by the first season conference win-loss percentage and multiplying by 100. This calculation can be repeated for each coach-school dyad in this analysis and, similar to before, these numbers should be standardized according to the previously stated formula. These values can then be used to derive percentiles and plotted to yield the graph below. To eliminate the redundancy of repeated plots, only the final percentile plot is included here. Note that this graph does not contain information for all coaches in this analysis because many years of data are missing for in-conference play prior to 1980; the coaches for which any of this data is missing are omitted from this graph to ensure these numbers are complete and accurate.

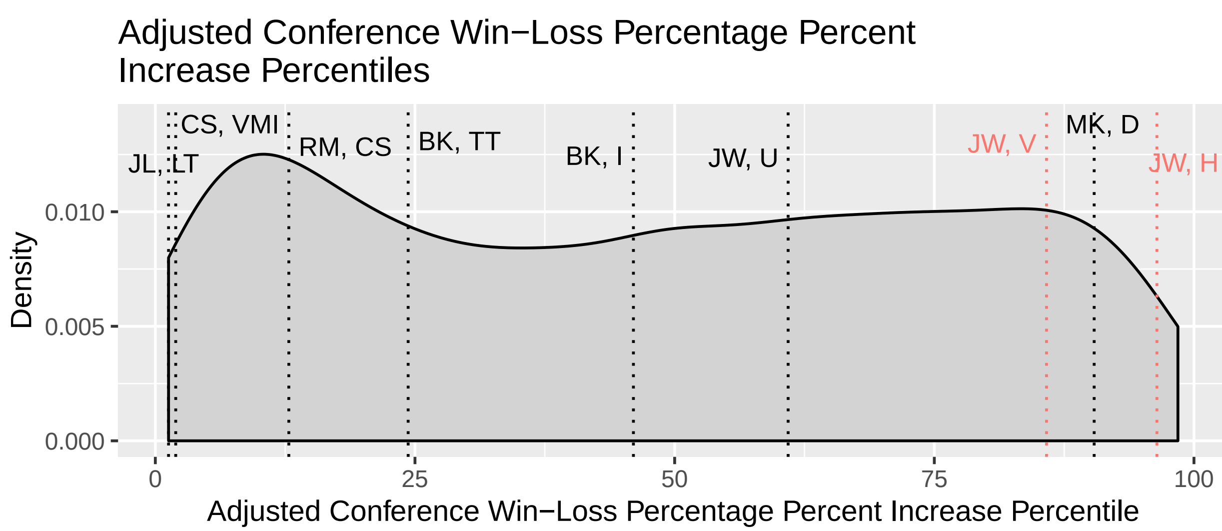

As can be seen on this plot, according to this analysis, Jay Wright is in approximately the 96th percentile for his success at Hofstra and the 85th percentile for his success at Villanova. These values are similar to the overall win-loss percentage percent increase percentiles described above, as would be expected.

Slope of Conference Win-Loss Percentage Throughout Career

Once more, consider the slope of the line that is created when conference win-loss percentage is plotted against season, indicated by a dotted line on the plot at the beginning of this section. The line's slope measures how rapidly each coach was able to improve his or her team's success against in-conference teams. The slope of every coach-school dyad is calculated and grouped by number of seasons coached at that school. Then, percentiles are calculated from and these values are shown on the graph below.

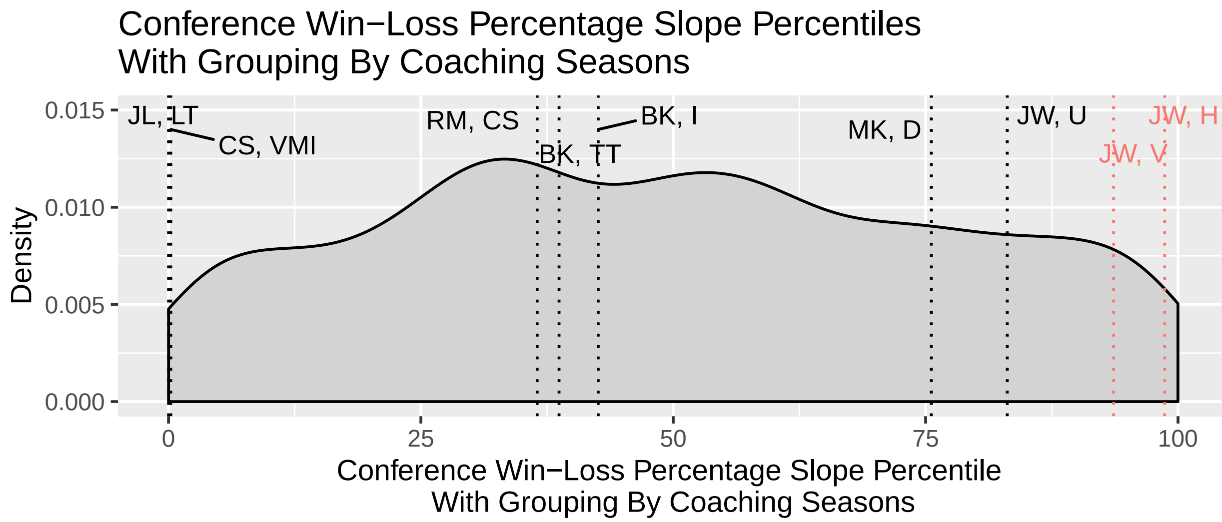

We see in this plot that Jay Wright falls in the 94th percentile for his conference win-loss percentage slope when grouped by coaching years and compared to his colleagues.

Expected Game Performance vs Actual Game Performance

For every game in college basketball, the probability that a given team wins can be calculated before the game even begins using what is known as an Elo rating. This probability is crucial to analyzing expected performance vs actual performance for every coach-school dyad, but first Elo ratings must be explained. Essentially, Elo ratings are zero-sum measures that can be used in any form of game where two opponents play each other head-to-head: basketball, soccer, chess, etc. When an Elo rating is first created, every team starts with the same score of 1500, the mean Elo rating. Then, as teams play one another, the winning team gains points based on the formula

$$Elo = Elo_0 + k\left(1 - \frac{1}{1 + 10^{\frac{Elo_0 - Elo_{Opp}}{400}}}\right)$$

and the losing team loses points based on the similar formula (where a negative number means a loss of that number of points)

$$Elo = Elo_0 + k\left(0 - \frac{1}{1 + 10^{\frac{Elo_0 - Elo_{Opp}}{400}}}\right)$$

In these formulas, k is known as the k-factor which varies based on how quickly the Elo reacts to game changes for a given sport. For the NBA, this value is accepted to be 20 and, thus, k=20 will used for these calculations.

There are a couple important notes about Elo ratings. First, this calculation accounts for a home-court advantage, or the fact that the home team tends to have an easier time throughout the game due to cheering fans and familiar facilities. For basketball, the accepted home-court advantage value is a 100 point increase in the Elo used to calculate the denominator in the formula given above. This essentially accounts for approximately 3 extra points a team would gain by playing their game at home. Also, it is important to note that these values are calculated after every game and change throughout a season. Therefore, at the end of the season, Elo ratings must be adjusted for the next season to account for the expected loss and gain of players according to the formula below.

$$Elo = 0.75(Elo_0) + 0.25(1505)$$

This formula takes $\frac{3}{4}$ of the previous season's Elo and adds it to $\frac{1}{4}$ of 1505, an Elo rating just above the mean Elo. This brings every team slightly closer to average yet maintains the rankings of teams compared to one another based on their Elos. Finally, to calculate the probability of a team winning a game, the denominator of the first Elo equation is used, namely

$$P(Win) = \frac{1}{1 + 10^{\frac{Elo_0 - Elo_{Opp}}{400}}}$$

These formulas were taken from https://fivethirtyeight.com/, which provides an in-depth explanation of Elo rating calculations for NBA teams. A Kaggle dataset including every college basketball game since 1985 was used to calculate an Elo rating and win probability for each team and every game. These values will facilitate the answer to the question of how a coach performs compared to their expected probability of winning a game.

Weighted Excess Wins

Consider the sum of all of Jay Wright's games as Villanova's head coach. If the probability of winning each game is calculated and then these results are split into ten categories of win probability (0%-10%, 10%-20%, 20%-30%, etc.) the plot below is created, where the percent chance Villanova wins the game is shown on the horizontal axis and the actual percent of games Villanova won in that category is shown on the vertical axis. The dashed lines are included for convenience and indicate the midpoint of that bar. For instance, in the category 50%-60%, the midpoint line lies at 55%, halfway between 50% and 60%. This bar indicates what the expected height of each bar is. Bars are colored to indicate if they reach/exceed the midpoint line, or if they fall short. In this plot, multiples of ten are represented in the lower category (i.e. 10% would fall into the 0%-10%)

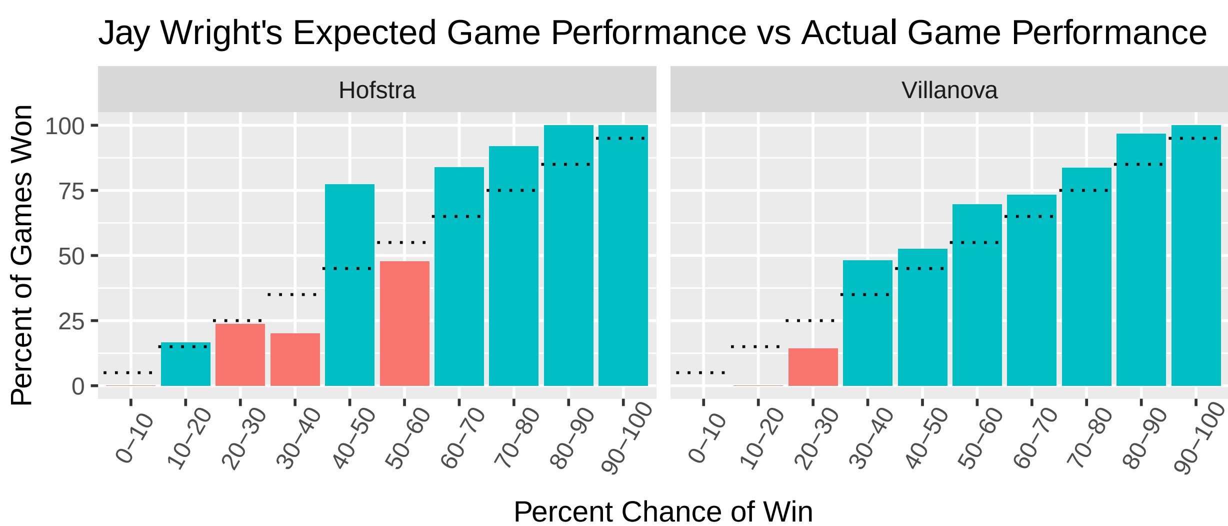

There are a few key points to be noted here. First, there are no bars in three sections of this plot; it is important to realize that this visualization does not indicate whether no games were played in this category, or games were played in this category but none of these games were won. This situation on this plot is very ambiguous, in fact, both cases are true in this plot alone. Next, notice the bar for Hofstra with a probability of winning between 40% and 50%; this bar is well above the level it is expected to fall at. This tells us that

Jay Wright's team at Hofstra performed very well when they were considered to be the underdog in a fairly even matchup. Finally, notice that most of these bars are blue and, thus, are above the midpoint value. If a bar is above this midpoint value, a coach's team performed better than they were expected to for that category. This analysis is concerned with how much each coach is above this value. If this plot is repeated for the coaches included in this analysis, the plot below is created.

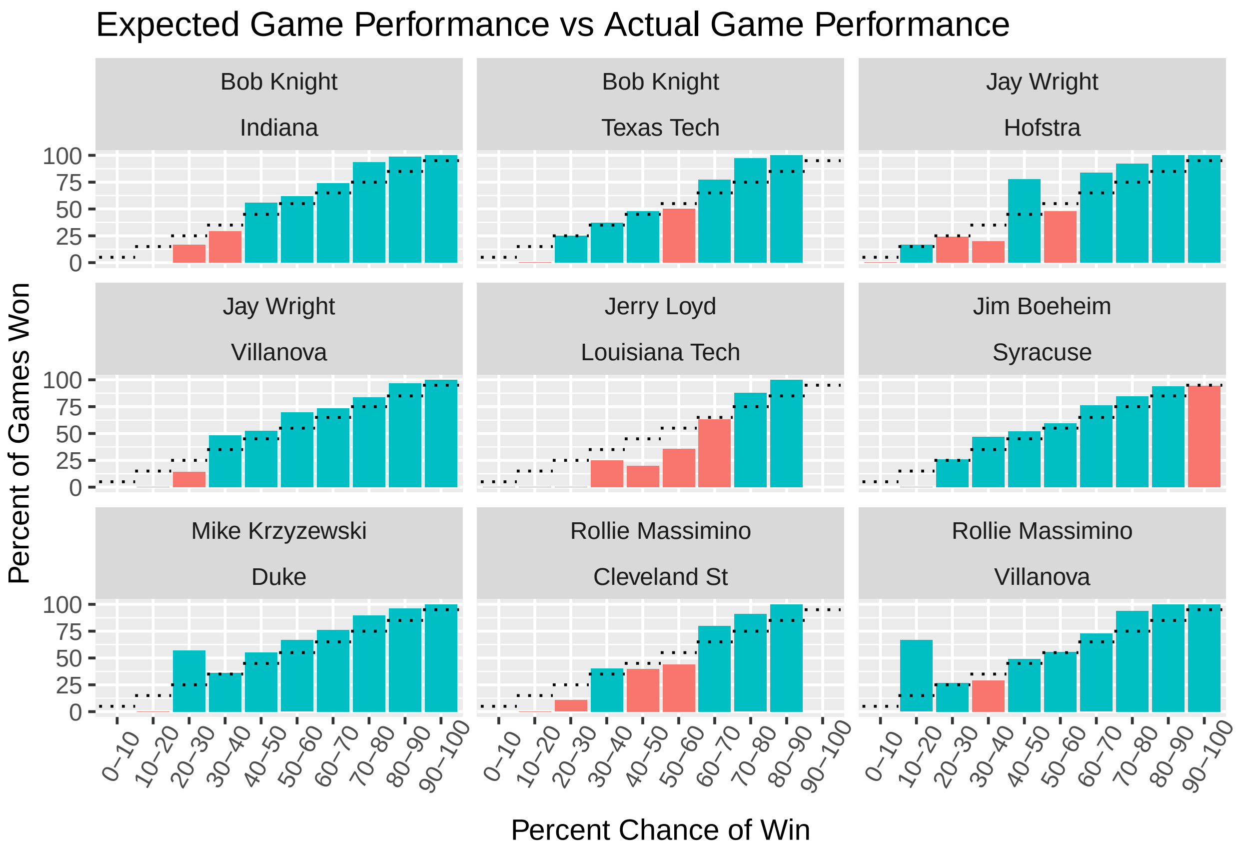

Very similar patterns to those discussed in the Jay Wright plot above are noticed in this plot. Again, we are concerned with how much each coach is above or below this midpoint value, how many excess wins a coach has. Thus, a statistic referred to here as weighted excess wins is calculated based on the formula below.

$$weighted\ excess\ wins = \sum_{i=1}^{10} \frac{(x_i - p_i)g_i}{g}$$

In this formula, $x_1 - p_i$, the mean actual win percentage minus the predicted win percentage also known as the excess wins, is calculated for each bar, and limits of one and ten are used to add up the values for all ten bars. The variable g indicates the number of games played and this value is used to weight the result. It is crucial to weight the result here because, for example, if a team won only one game but also only played one game in any given section of this chart, they would have a percent of games won equal to 100%, making it appear that the team performs very well in this category while this may not be an accurate description. Weighting the excess wins with g accounts for this issue. After calculating these values, the units of weighted excess wins are hard to understand, however, these values can be converted into percentiles to create the plot below. With these values we can easily compare coaches to one another.

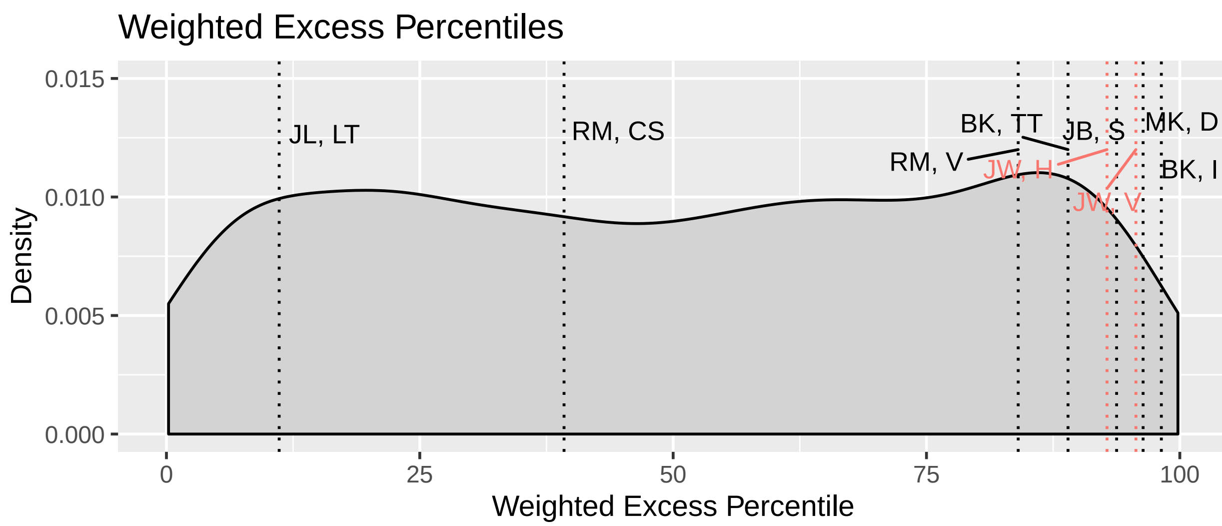

We see that Jay Wright falls in the 95.7th and 92.8th percentiles for his performance at both Villanova and Hofstra, respectively. These values are impressive for any coach but particularly for a coach of a team like Villanova where the bar has already been set pretty high.

Elo Ratings

The Elo rating, as described in the section above, is essentially a measure of how good a team is compared to its opponents on any given day. Since Elo ratings take into account opponent Elo scores as well as whether the game was won or lost, the amount an Elo increases after a win depends on the opponent's skill level; thus a win against a more skilled opponent will earn a team a greater increase in Elo rating than a win against a similarly ranked team. Because of this, the Elo rating is also beneficial to use to assess a team's growth. Most measures used in this paper thus far have resulted in similar percentiles, however, this analysis will spotlight not only consistently great coaches, but also coaches that may not have the best team, but show the most growth. The goal of including this measure in this report is to account for team skill at the beginning of the season and not falsely attribute high performing teams with great coaching ability when, perhaps, this great performance is a result of the players' skills, and not coaching skill. A great coach team must still grow, no matter how high their team's Elo is at the beginning of the season.

Elo Rating Increase

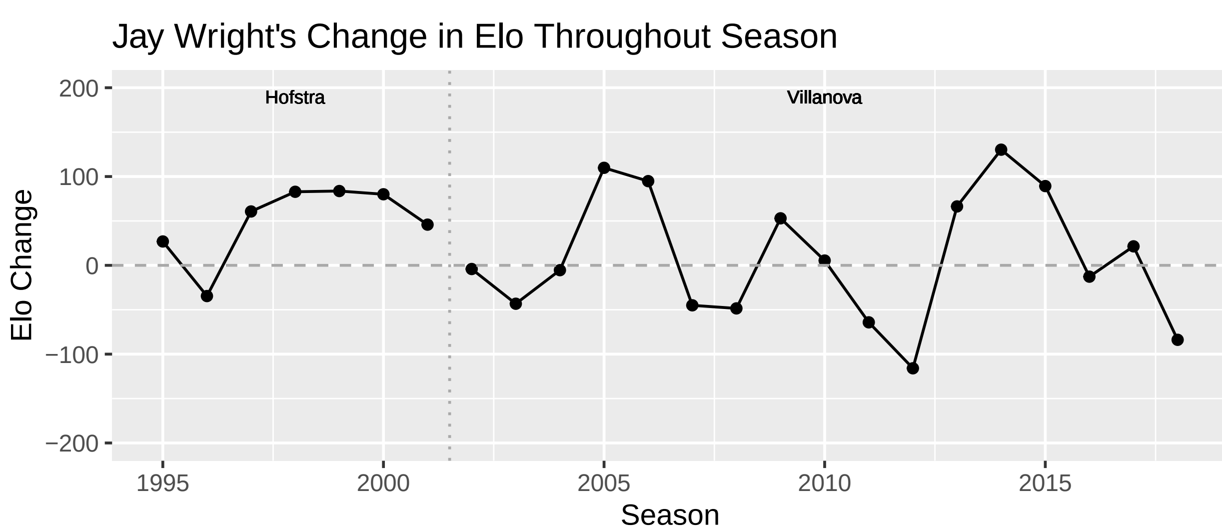

Consider the previous plot, where season is plotted on the horizontal axis and Elo change is plotted on the vertical axis for Jay Wright at Hofstra and at Villanova; a dashed line is included at zero for convenience. Any point above zero indicates an Elo rating increase throughout that season and any point below zero indicates an Elo rating decrease throughout that season. One important note is that this plot does not account for how high an Elo is for that season – this is the precise goal of this measurement. For instance, from this graph it may appear that Jay Wright was better in 2005 than in 2016, however, in 2005, Jay Wright's team increased their Elo rating from 1556 to 1666 while in 2016 Wright's team decreased their Elo rating from 1832 to 1819. An increase in 110 Elo points is not unimpressive, however, an Elo rating of 1832 is clearly better than 1556, and explains how the wildcats could win the National Championship that year. Thus, it is important to remember that this plot does not account for how high an Elo rating is, but how much it changes.

The same variables can be similarly plotted and this graph can be expanded to include all coaches in this analysis as shown below. As mentioned above, this plot does not indicate a team's skill level, but its improvement throughout the season.

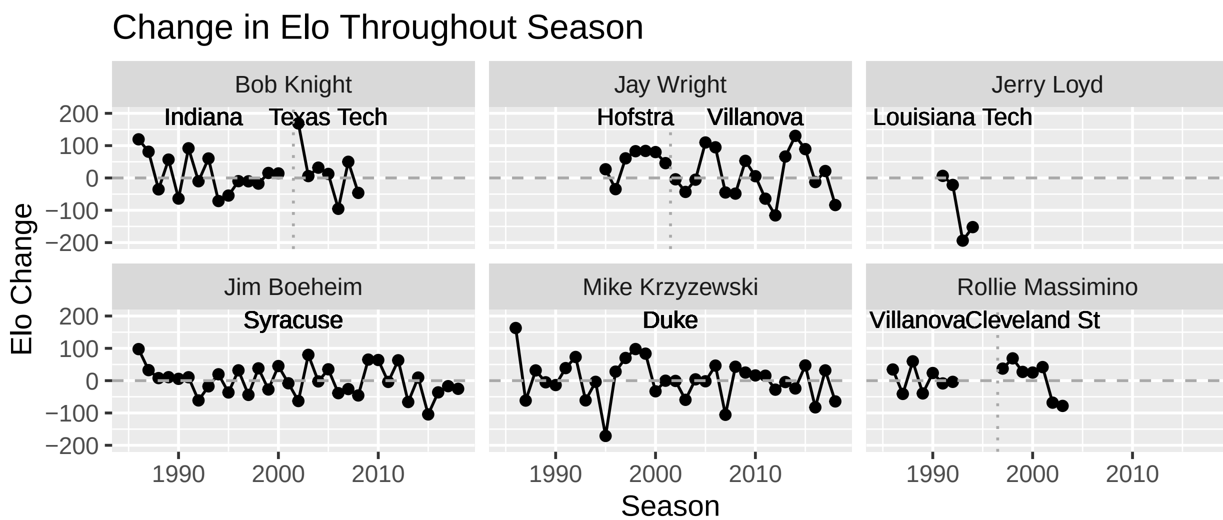

In these plots, one point indicates the change in Elo for each coach-school dyad and each season. When the change in Elo rating from season to season is averaged for each coach-school dyad, one value of mean Elo change is returned for each coach-school dyad and these values can then be converted to percentiles. The results are shown in the plot below.

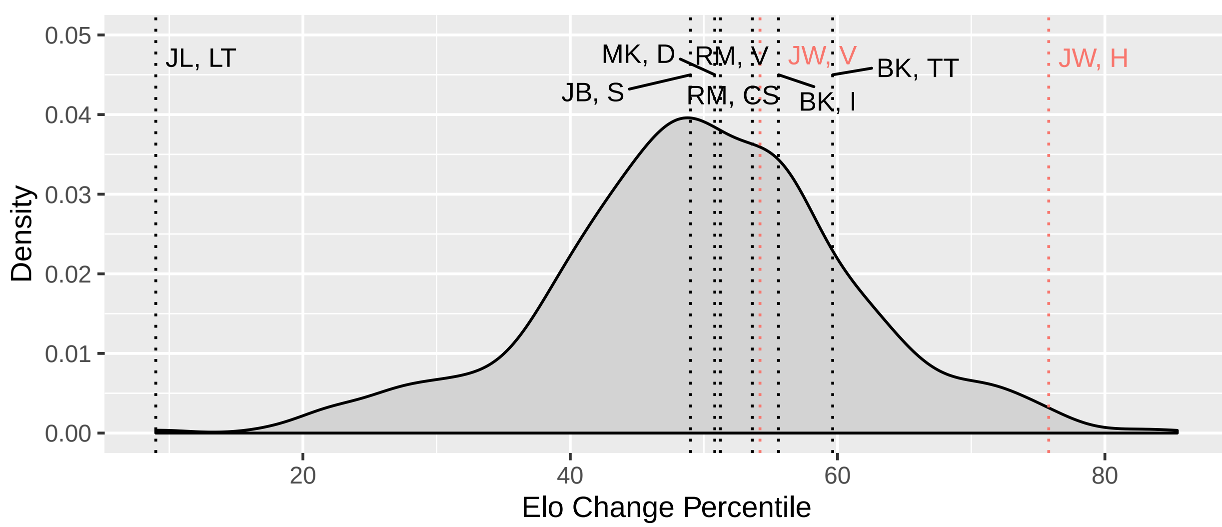

Minutes left							
Lead	35	30	25	20	15	10	5
0	500	500	500	500	500	500	500
1	514	520	526	534	539	547	569
2	529	541	553	568	578	593	636
3	543	561	579	602	616	637	698
4	557	581	604	634	653	679	753
5	572	601	630	666	688	719	801
6	586	620	654	696	721	755	842
7	600	639	677	724	751	788	876
8	613	658	700	751	780	818	903
9	627	676	722	775	806	844	925
10	640	694	743	798	829	867	942
11	653	711	762	820	850	888	956
12	666	728	781	839	869	905	966

Figure 1: In-game Win Probabilities for Even Matchup - Kenpom.com

Jay Wright at Villanova has a lower percentile here than what has been seen thus far in this report. This is likely because this percentile is from a statistic measuring growth, not skill. Jay Wright falls in the 75.8th percentile for his years at Hofstra and the 54.2th percentile for his years at Villanova according to this measurement.

Win Speed

One major problem with most basketball game statistics is that sometimes what would be considered a blowout game by spectators does not appear to have been an easy win on paper. This happens when, with a decent amount of time left, the winning team's coach recognizes that his team will win this game and chooses to replace his starters with his subs to give his starters a break. Analyzing how quickly a coach gets to this point in the game will help account for this issue and allow us to consider a "win speed" statistic indicating how quickly a coach is able to lock in a win.

Kenpom In-Game Probabilities

On his website https://kenpom.com, Ken Pomeroy, a well-known basketball statistician, describes how he used data from approximately 700 games to calculate the in-game win probability for each team in a 50/50 matchup game. This means that, for any game with opponents who are approximately equal in ability, the probability that a given team wins at any minute can be calculated based on the time remaining and the team's lead. Essentially, Pomeroy sampled 700 games by recording the lead of the winning team every five minutes and recording who won the game. After collecting this, Pomeroy smoothed the data and applied a regression to derive a table of values for the probability of winning a game at every five minute mark. The probabilities Pomeroy found and published on his website are shown in Figure 1 with each value equal to 1000 times the probability (i.e. a value of 500 indicates a 0.5 chance of winning); these are the values that will be used in this analysis. Nevertheless, there is one major flaw in using this sampling to determine in-game win probabilities which Pomeroy discusses later in his article: this method assumes every game is a 50/50 matchup, despite the fact that this is often not the case. There are multiple methods Pomeroy uses to overcome for this flaw, and this analysis has been designed based off these methods.

Before adjusting for uneven matchups, the values for any minute that is not a multiple of five had to be interpolated since Pomeroy only recorded probabilities every five minutes. Since these values were roughly linear, they were simply estimated based off the surrounding multiples of five minutes using a linear interpolation. This returned a probability of winning for every minute and every possible lead for 50/50 matchups. This can then be used in combination with the probability a team wins that is calculated before the game to give a more accurate in-game win probability for uneven matchups.

The first way Pomeroy adjusts his calculation does not account for any difference in team ability, but, instead, accounts for how hard each player on a team will play based on the time remaining. In his article, Pomeroy treats time as nonlinear; instead of using the time remaining in the game, Pomeroy uses the square root of the fraction of time left based on the assumption that players will try harder when less time is remaining than they will in the first few minutes. For example, halfway into a game, when 20 minutes are left, 28 minutes are considered to be left because $\sqrt{\frac{1}{2}} * 40 \approx 28$.

Log5

The second way Pomeroy accounts for disparities in team abilities is using what is known as the log5, or a formula invented by Bill James to compare teams to the average, 0.500. James wrote about this formula, "I call it the log5 system since it is, essentially, a logarithmic system which is based upon a weighted comparison of each team to a .500 team" (Kushner). This calculation is exactly what is needed here because the probabilities included in our in-game probability table are all those of 0.500 teams. The formula for a log5 calculation is included below and can be found in the article "Kenpom, Pythag, and Expected Scoring Margin: A Reader's Question."

$$P_{A,B} = \frac{P_A - P_A P_B}{P_A + P_B - 2P_A P_B}$$

$$P_A = initial\ estimate\ of\ favorite \qquad P_B = even\ strength\ estimate\ of\ opposing\ team$$

In this formula, P_A is the initial probability that the favored team will win, that is, the probability calculated using Elo scores and found before the game is even played. P_B is the estimate that the opposing team will win at a selected moment, assuming the game is an even matchup. Since we have previously found all these values, we can use this formula to calculate the log5. Finally, this log5 value is used to calculate the probability that either team wins at any given moment of the game with a linear calculation. For this value, the percent chance of winning is determined to be equal to the product of percent change of winning at that minute and the fraction of the game played plus the product of the initial percent chance of winning and the fraction of the game remaining. This calculation is necessary to account for the fact that there is still time to be played in the game. This is all calculated using the data downloaded from Kaggle and the table taken from https://kenpom.com. It is important to note that since Pomeroy only recorded data for a maximum lead of 25, any games where a team had a lead greater than 25 will be truncated at 25. This means that, for instance, a team that reaches a lead of 30 points at a given minute will be calculated as if it has a 25 point lead at that minute.

Threshold for Guaranteed Win

In order to find how quickly a team secures a win, a threshold must be selected as a percentage at which the percent chance of winning the game is considered high enough to practically guarantee a win. Essentially, this threshold would say that if a coach's team reaches a win probability higher than whatever percentage, they are extremely likely to win the game. Since there is no agreed upon value for this threshold, the data was used to select an appropriate threshold. To do this, the highest percent chance of winning during the game for each opponent in every game was selected. This value was then used to determine if the team should have won the game based on a tested probability threshold. Values tested include 75%, 80%, 85% and many values in between and the expected outcome was either "W" or "L" based on if this highest win probability was higher than/equal to or lower than the threshold, respectively. A confusion matrix was created for these various thresholds and the percent of false positives and false negatives was analyzed, using trial and error to determine the best threshold. As an example, consider a game where one team is favored to win going in to the game with an 85% chance of winning the game at tip-off; now consider that the underdog actually won the game. Probability wise, the favored team should have won because they had a higher probability;

they had the expected win. Nevertheless, the favored team did not win, they had an actual loss. This is an example of a false positive. The goal was to minimize both false positives and false negatives, and since when the percent of false positives decreased the corresponding percent of false negatives increased, the sum of these two percentages was taken to find the total percent error. This total percent error was minimized at a threshold of 81%. For simplicity purposes, the threshold to secure a win was set at 80%; the confusion matrix for a threshold of 80% can be found below. We see here that approximately 4.7% of predicted wins resulted in losses and 6.2% of predicted losses resulted in wins.

		Predicted Loss	Predicted Win
1	Actual Loss	95.334140	4.66586
2	Actual Win	6.248037	93.75196

Time Left in Game When Probability of Winning Passes 80%

Consider the three plots below, each of which represents a single game played in 2017. The horizontal axes on these plots display minutes left in the game and the vertical axes display the probability that the team, either Villanova or Duke, will win that game. A blue smooth curve is included to display the general trend of the team's win probability throughout this game and a dashed line is included at the chosen threshold of 80%. An additional dashed line is included as a reference point to mark a 50% chance of winning. In this analysis, a coach is considered to have secured a win when they pass the 80% threshold.

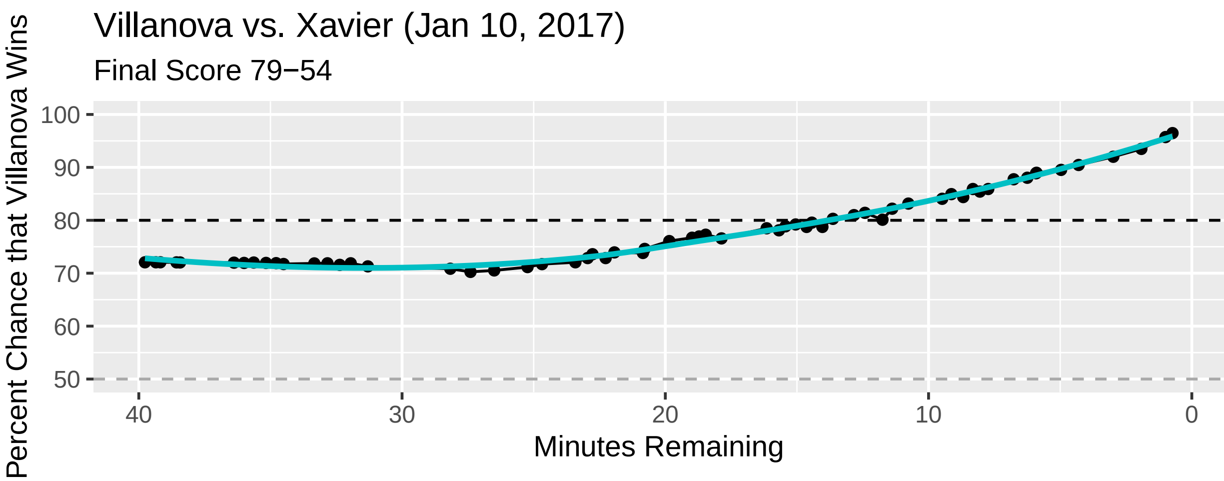

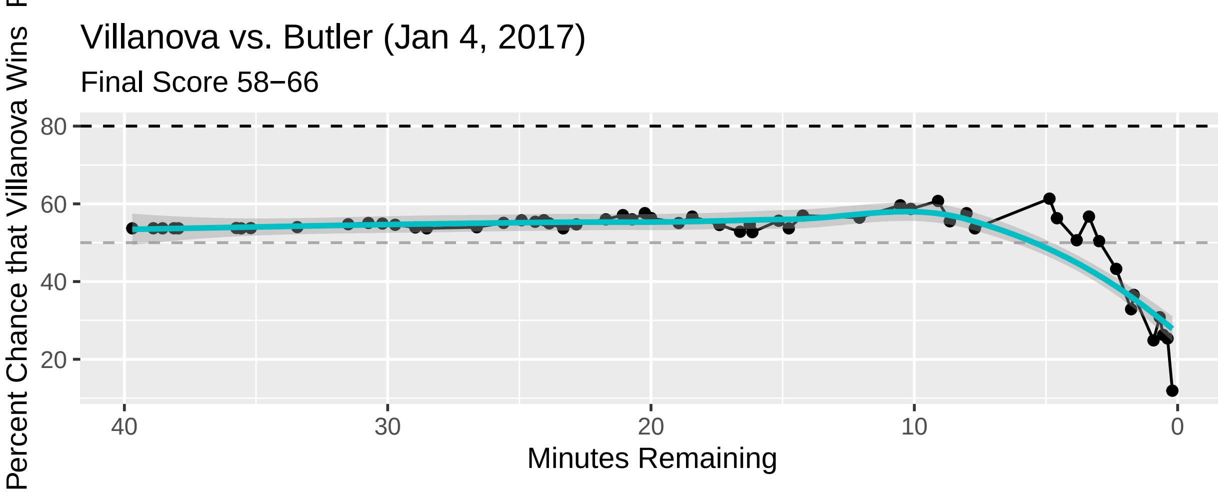

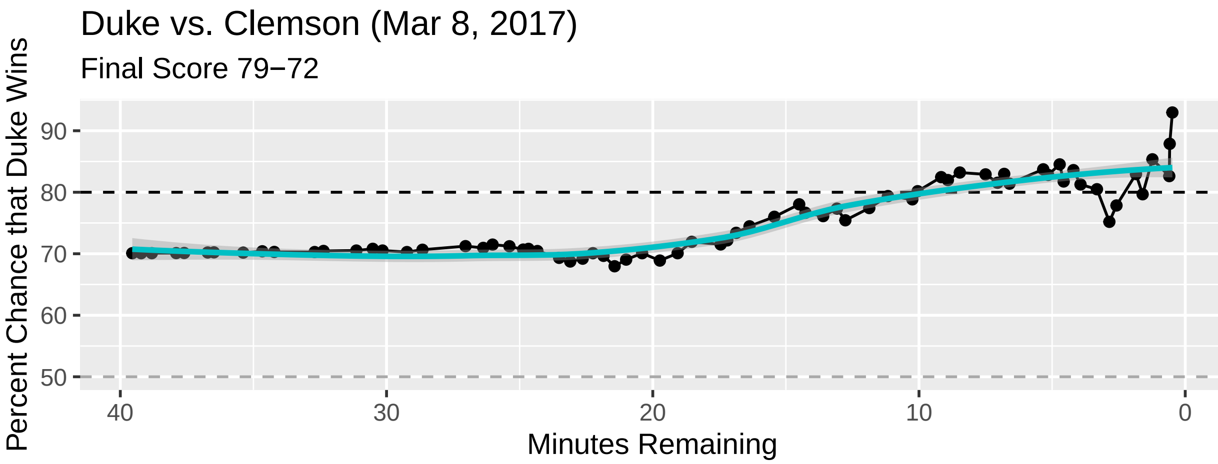

These specific games were selected as instances of the three different cases that occur. The first plot depicts a game where the Wildcats under Jay Wright achieved a "guaranteed" win with 15 minutes still remaining in the game; that's almost half the game, an impressive feat for sure. Consider now the second plot, which depicts a Villanova loss against Butler; here Villanova's chance of winning never exceeded 65%. In the final plot, which depicts a game against Duke and Clemson in March 2017, we see that Duke begins with a 70% chance of winning, exceeds the 80% threshold with approximately 10 minutes left in the game, and later dips below this threshold, exceeding the threshold again briefly before the game finishes with a score of 79-72. It is important to note that, in cases such as this, the time a coach will be considered to have secured the win will be the final time their percent chance of winning exceeds 80%. If the time a win is secured is derived for every game and every coach, percentiles can be based on how quickly a coach reaches this threshold to produce the following graph.

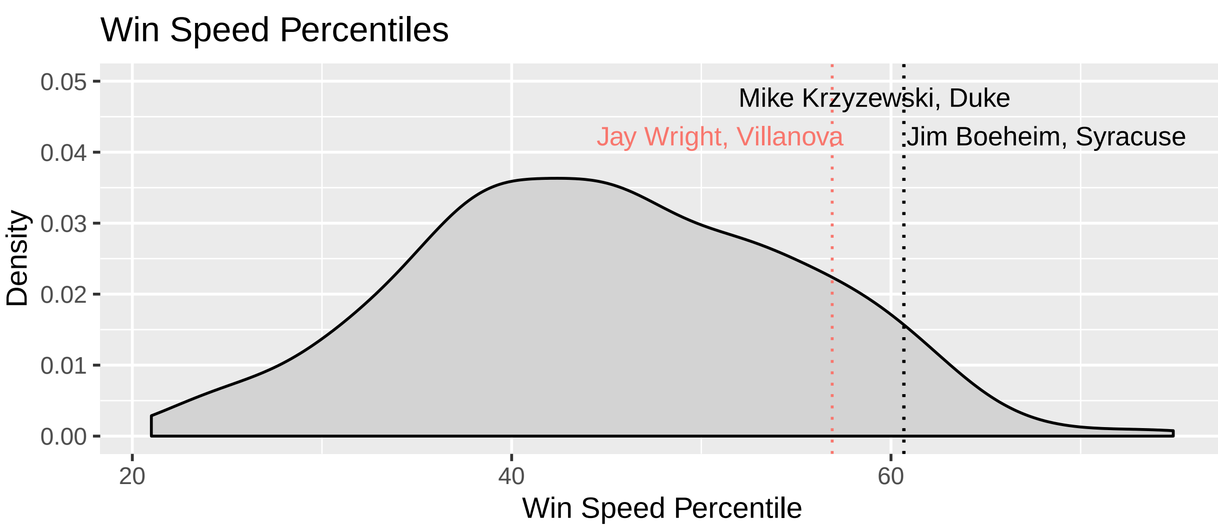

As can be seen in the plot above, Jay Wright lies in the 53rd percentile. This means that Coach Wright reached this 80% threshold in games faster than 53 percent of his colleagues when he coached at Villanova. Note that Jay Wright's performance at Hofstra is, again, not included in this visualization because this data is only available since 2010, when Jay Wright was already coaching at Villanova.

Conclusions

This analysis presently measures coaches according to their overall win-loss percentage percent increase, overall win-loss percentage slope, conference win-loss percentage percent increase, conference win-loss percentage slope, weighted excess wins, seasonal Elo change, and win speed. These numbers were calculated and converted to percentiles and the results are shown altogether on one plot below. While repeating the same methods for two different datasets, namely overall win-loss percentage and conference win-loss percentage, may seem redundant, this method of analysis was deemed the best way to examine the win-loss percentage; to analyze such similar statistics two different ways would be inconsistent. Thus, the same methods were performed on the overall win-loss percentage and the conference win-loss percentage. Some of the percentiles ultimately calculated are extremely close, nevertheless, these similarities strengthen the analysis in how they affirm one another.

If all calculated percentiles are plotted on the same graph (excluding overall/conference win-loss percentage without grouping because this measurement was deemed misleading earlier in this analysis), we get the following graph.

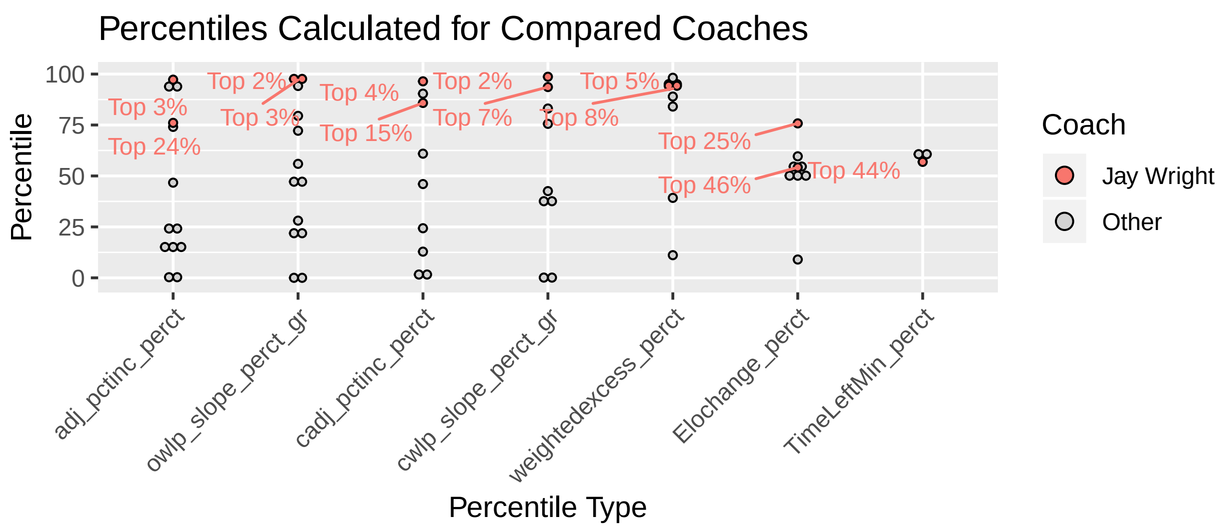

As can be seen here, compared to other coaches, Jay Wright is clearly one of the greater coaches in college basketball. Coach Wright falls in both the 97th and 76th percentile for his percent increase in win-loss percentage and, similarly, lies in the 96th and 85th percentile for his increase in in-conference win-loss percentage. For his slope of his teams' win-loss percentage, Coach Wright falls in the 97th and 98th percentile and for the slope of his teams' in-conference win-loss percentage, Jay Wright falls in the 92nd and 98th percentile. For weighted excess, the measure of how much a coach exceeds his teams' expected wins, Coach Wright falls in the 95th and 92nd percentile. For Elo change, Jay Wright's is classified in the 75th and 54th percentiles for his performance at Hofstra and Villanova. Finally, for his win speed, Jay Wright lies in the 56th percentile.

For his performance at both Villanova and Hofstra, Jay Wright ranks above the 90th percentile 8 times out of 13 and is in the top 75% of coaches in all but two measurements. On top of this, Coach Wright falls in the top 5% of coaches in 6 of these analyses, an incredible achievement. Nevertheless, it would not make sense to simply average all of these percentiles. The analyses used in this project are only a small subset of possible measurements of coaching ability and, thus, cannot be combined to give an accurate, single-value representation of coaching ability. It is worthy of noting, however, that, of the coaches included in this analysis, Jay Wright has the highest percentage of percentiles in the top 10% of coaches. Jay Wright ranks in the top 10% of coaches for 62% of the analyses used in this project. This means that, though Jay Wright has more percentile values for some measures used here due to the fact that he has coached at not one but two schools, he also has a higher number of high percentiles than those coaches. The values for each coach can be found below.

	Coach	# of Top 10 Pctl.	# of Measures	% of Top 10 Pctl.
1	Jay Wright	8	13	61.538462
2	John Wooden	2	4	50.000000
3	Mike Krzyzewski	3	9	33.333333
4	Jim Boeheim	1	5	20.000000
5	Bob Knight	1	14	7.142857
6	Charlie Schmaus	0	4	0.000000
7	Jerry Loyd	0	6	0.000000
8	Rollie Massimino	0	10	0.000000

While this project does not even begin to explore every way coaching ability can be measured, Jay Wright's performance in this analysis is outstanding. Compared to the coaches included here, Jay Wright would rank #1 by these measures. Thus, it is certainly safe to say that Coach Wright is a great basketball coach, particularly according to his teams' win-loss percentages. If the attention the media has given Jay Wright in recent years and the success of his team were not enough to qualify Coach Wright as a great basketball coach, the numbers here certainly are.

Further Investigations

There are multiple ways to analyze coaching ability, only a small fraction of which could be included in this project. Ideas for further research include, but are not limited to:

1. What relationship exists between the variation in player abilities and coaching success? Does a greater gap (indicating a "star" player(s)) affect coaching in any way (it may be easier to coach a team with greater chemistry), or vice-versa?
2. Do better coaches have higher player retention rates (excluding NBA recruits since this would be indicative of success on either the coach's/player's end)? In addition to this, do better coaches acquire better players?
3. Is there a difference in performance against teams with better defensive/offensive skills? Does the coach prepare for that?
4. How have the individual players on a coach's team improved? A good coach would not only help the team grow but individual players should show improvement as well.
5. When a team reaches the 80% threshold to guarantee a win, does their defensive pressure increase?
6. What statistic is most correlated to the determined coach of the year?
7. Do good mentors typically produce better coaches? Is there an relationship between how good a coach is and how good the coaches who taught him were?

Works Cited

"College Basketball Statistics and History: College Basketball at SportsReference.com." *SportsReference*. https://www.sports-reference.com/cbb/.

"Kenpom, Pythag, and Expected Scoring Margin: A Reader's Question." *SBNation*. 27 Feb 2008. https://www.rockytoptalk.com/2008/2/27/165520/578.

Kushner, James. "The Best Teams in Baseball History." *Baseball Prospectus* 28 Jul 1998. https://web.archive.org/web/20110118052847/http://www.baseballprospectus.com/article.php?articleid=169.

Massey, Kenneth and Jeff Sonas. *Google Cloud & NCAA® ML Competition 2018-Men's*. 2 Apr 2018. Web. https://www.kaggle.com/c/mens-machine-learning-competition-2018/data.

Pomeroy, Ken. "In-game win probabilities." *Kenpom*. 3 Apr 2010. https://kenpom.com/blog/ingame-win-probabilities/.

Silver, Nate, and Reuben Fischer-Baum. "How We Calculate NBA Elo Ratings." *FiveThirtyEight* 21 May 2015. https://fivethirtyeight.com/features/how-we-calculate-nba-elo-ratings/.