

2013

The Complete Plus-Minus: A Case Study of The Columbus Blue Jackets

Nathan Spagnola
University of South Carolina - Columbia

Follow this and additional works at: <https://scholarcommons.sc.edu/etd>



Part of the [Statistics and Probability Commons](#)

Recommended Citation

Spagnola, N. (2013). *The Complete Plus-Minus: A Case Study of The Columbus Blue Jackets*. (Master's thesis). Retrieved from <https://scholarcommons.sc.edu/etd/2708>

This Open Access Thesis is brought to you by Scholar Commons. It has been accepted for inclusion in Theses and Dissertations by an authorized administrator of Scholar Commons. For more information, please contact digres@mailbox.sc.edu.

THE COMPLETE PLUS-MINUS:
A CASE STUDY OF THE COLUMBUS BLUE JACKETS

by

Nathan Spagnola

Bachelor of Arts
Kenyon College, 2010

Submitted in Partial Fulfillment of the Requirements

For the Degree of Master of Science in

Statistics

College of Arts and Sciences

University of South Carolina

2013

Accepted by:

John Grego, Major Professor

Leslie Hendrix, Committee Member

David Hitchcock, Committee Member

Don Edwards, Committee Member

Lacy Ford, Vice Provost and Dean of Graduate Studies

© Copyright by Nathan Spagnola, 2013
All Rights Reserved.

DEDICATION

This thesis is dedicated to Dr. John Grego and David Busby. Both have been instrumental in creating this finished project. I was very excited when Dr. Grego said he would be my advisor for this process. Looking at the datasets together and trying to figure out how to solve various issues throughout the analysis really was a lot of fun. It was through David's master coding that we were able to compile a dataset for this analysis. We have since been able to create a program that analyzes NHL teams and have hopes of marketing this program to professional teams. Without Dr. Grego and David, this would not have been possible.

ABSTRACT

A new hockey statistic termed the Complete Plus-Minus (CPM) was created to calculate the abilities of hockey players in the National Hockey League (NHL). This new statistic was used to analyze the Columbus Blue Jackets for the 2011-2012 season. The CPM for the Blue Jackets was created using two logistic regressions that modeled a goal being scored for and against the Blue Jackets. Whether a goal was scored for or against the team were the responses, while events on the ice were the predictors in the model. It was found that the team's poor performance was due to a weak defense and severe underperformances by key players.

TABLE OF CONTENTS

DEDICATION	iii
ABSTRACT	iv
LIST OF FIGURES	vi
CHAPTER 1: INTRODUCTORY HOCKEY 101	7
CHAPTER 2: SPORTS STATISTICS	9
2.1 HOCKEY SABERMETRICS	10
2.2 THE PLUS-MINUS AND ADJUSTED PLUS-MINUS STATISTICS	11
2.3 NEEDED IMPROVEMENTS IN HOCKEY ANALYSES	12
CHAPTER 3: METHODOLOGY	13
3.1 AIMS OF THIS STUDY	13
3.2 METHODS	14
3.3 STATISTICAL ANALYSIS.....	16
CHAPTER 4: RESULTS AND DISCUSSION	21
4.1 RICK NASH: THE FACE OF THE FRANCHISE	21
4.2 THE STATE OF THE TEAM	25
4.3 SUGGESTED MOVES.....	29
4.4 ANALYSIS OF ACTUAL TRANSACTIONS.....	30
CHAPTER 5: CONCLUSIONS AND FURTHER STUDIES	34
REFERENCES/BIBLIOGRAPHY/WORKS CITED.....	36

LIST OF FIGURES

Figure 4.1 Rick Nash CPM Over 40 Games.....	22
Figure 4.2 Rick Nash Offensive Ability Over 40 Games	23
Figure 4.3 Rick Nash Defensive Ability Over 40 Games.....	24
Figure 4.4 Blue Jacket Average CPM Over 40 Games	26
Figure 4.5 Blue Jackets Average Offensive and Defensive Production Over 40 Games ..	27
Figure 4.6 Maple Leafs Offensive and Defensive Production Over 40 Games.....	28
Figure 4.7 CPM for Players No Longer With the Blue Jackets.....	32

CHAPTER 1

INTRODUCTORY HOCKEY 101

To understand this analysis, it is not important to comprehend all the nuances of hockey. It is, however, important to understand key aspects of the game. This section will provide the necessary information so that even a hockey novice will be able to understand the following report. To begin this introductory course in hockey, we will start with the basics.

Each team has 5 men on the ice at one time, not including the goalies. The 5 skaters consist of 3 offensive players (2 wings and a center) and 2 defensive players. Any of these players can score a goal at any time. When a player is on the ice, that time on the ice is called a shift. Typically a shift lasts about 1 minute and a player can expect to have anywhere from 10 to 25 shifts a game. Better players are given longer and more numerous shifts than lesser players.

A single game consists of 3 periods of 20 minutes each. If a game is tied after regulation, the game goes into sudden death overtime where each team plays down a man during a 5 minute period. If at the end of the 5 minutes neither team has scored, the game goes into a shootout. The two teams take turns sending a skater out onto the ice one at a time to try and score against the opposing goalie. After each team has had a total of 3 skaters attempt to score a goal, the team with the most goals wins. If the game is still tied at this point, then the teams continue in a sudden death shootout until a winner is determined.

There are also penalties that may be called by the referees during the game. If a penalty is called, the guilty player must leave the ice for a determined amount of time. Minor penalties are 2 minutes long and double minors are 4 minutes. The team of the penalized player must play down a man for this period of time. Once the penalty time has expired, the player may return to the ice and the game proceeds 5-on-5. Some penalties, such as penalties for fighting, result in a 5 minute penalty being given to each player involved in the fight, but the teams continue at full strength. A 10 minute penalty may also be given for misconduct and while the guilty player cannot be on the ice for 10 minutes, both teams continue to play 5-on-5.

CHAPTER 2

SPORTS STATISTICS

The world of sports has dramatically changed over the past few years. Statistics are being implemented more and more as a better way to determine player ability and value. Even sports such as baseball, which has always had a focus on statistics, have adopted new measures to calculate how good players are.

The most famous example is Moneyball, a book written by Michael Lewis which was later made into the movie starring Brad Pitt. This book follows the true story of Billy Beane, the general manager of the Oakland A's. Billy Beane was limited by his owner's unwillingness to spend the same amount of money on players that other teams in the league were spending. This of course meant that the Oakland A's had a distinct disadvantage playing teams such as the New York Yankees who had 3 times the payroll as the A's. Billy Beane knew he had to do something differently. That something was the use of sabermetrics.

For years baseball had been a sport about batting average and homeruns. These were the pillars of baseball statistics, the first things anyone looked at on the back of a baseball card or when a player came to the plate. Beane realized, however, that these were poor statistics for measuring the impact a player had on a game. Instead, Beane started looking at new statistics such as on-base percentage and slugging percentage. Using these sabermetrics, Beane compiled a team of undervalued players. These players

were able to amount to something greater than their sum and in 2002 they made history by breaking the American League record for consecutive wins and won the division title.

2.1 HOCKEY SABERMETRICS

The advent of sabermetrics, and the success of teams such as the Oakland A's that use these neo-statistics, has opened the flood gates in all sports. Basketball is perhaps the sport that uses sabermetrics the most next to baseball, and other sports are beginning to follow suit. This has created an arms race of sorts for statisticians to create the next best statistic.

The usual statistics for hockey are admittedly extremely poor for measuring a player's ability (Gramacy et al., 2013). The number of shots a player takes, the number of goals scored, and the number of hits a player has made simply don't measure ability accurately. New statistics have started to become more popular such as Corsi numbers, but there are still questions as to the accuracy of such indexes (Vollman, 2010).

Corsi numbers (equation 2.1) are based on statistics readily compiled for each

Equation 2.1:

$$\text{Corsi} = (\text{Shots on Target For} + \text{Missed Shots For} + \text{Blocked Shots Against}) - (\text{Shots on Target Against} + \text{Missed Shots Against} + \text{Blocked Shots For})$$

player. As we can see, this is not so much a statistical model as it is a simple index. Corsi numbers assume that the weights are equal for shots on target, missed shots for, and blocked shots against. Clearly having a shot on target, i.e. a shot on net, is more beneficial than a shot that was missed. Yet Corsi forces these two events to have equal coefficients. This index also allows no room for team variation. For example, a defensive

team would be expected to have very different coefficients than an offensive team. While the Corsi number is a good attempt to improve hockey sabermetrics, it fails to take a truly statistical approach.

2.2 THE PLUS-MINUS AND ADJUSTED PLUS-MINUS STATISTICS

One of the better common hockey statistics is called the plus-minus. This is another simple statistic: a player gets a plus 1 if a goal is scored for their team when they are on the ice and a minus 1 if a goal is scored against the team when that player is on the ice. This statistic was created to measure player effect better than the usual statistics. The reasoning behind this statistic is that while a player may not actually score the goal, they may still have an important impact on the play. A player might have obscured the goalie's view of the puck, or perhaps the defense swarmed to that player leaving another player open to score the goal. The plus-minus is supposed to measure this effect, but still fails to do so with sufficient accuracy and precision (Gramacy et al., 2013).

Hockey sabermetrics have taken the plus-minus a step further and created the adjusted plus-minus (APM). Statisticians such as Brian Macdonald (2010) have taken steps to improve the basic plus-minus statistic by using weighted least squares regression to create his version of the APM. The issues with these analyses mainly lie in extremely large confidence intervals. Macdonald used fixed effects in his model to estimate player contribution and looked at all players in the league. He has failed, however, to attain confidence intervals for player effects that are sufficiently narrow. He has later taken aims, such as ridge regression, to alleviate this issue, but has yet to succeed in narrowing these intervals a substantial amount (Macdonald, 2011).

2.3 NEEDED IMPROVEMENTS IN HOCKEY ANALYSES

There are many issues with the approaches other statisticians have taken. Most of these are the same issue Macdonald has faced in that confidence intervals for player effects are simply too wide. A recent paper by Gramacy et al. (2013) used regularized logistic regression to take a Bayesian approach to narrowing intervals. This analysis, however, used a monstrous 4 years of data to achieve this goal.

Another common weakness with other published papers is that they tend to focus solely on proven players and perennial All-Stars. Any model, no matter how poor, should be able to properly credit these elite athletes as good players. This is not something that is hard to do. What is more difficult is to separate the average players from each other, and what is more difficult yet is to find undervalued players. This study aimed to do just that.

CHAPTER 3

METHODOLOGY

3.1 AIMS OF THIS STUDY

The goal of this study was to change the way the common adjusted plus-minus statistic is used. All other papers focus on analyzing every player in the league and only discuss how that study's version of the APM correctly labels elite players as elite. While this may be interesting from the standpoint of a fan, this approach is not that interesting or helpful from the standpoint of a team.

Contrary to previous analyses, this analysis focused solely on the Columbus Blue Jackets. The Blue Jackets have always been one of the worst teams in the league and have only made the playoffs once during their 13 year existence. (They were swept in the first round of those playoffs by the Detroit Red Wings.) Needless to say, the Blue Jackets have a reputation of blown trades and wasted draft picks, so it may be argued that the real issue the franchise faces is incorrectly evaluating players.

This study took the novel approach in analyzing the team in the form of the following question: "If we were heading into the trade deadline of the 2011-2012 season and had been using the models later described in this paper, what changes would we have suggested the team make, and were the moves that the team actually made beneficial?"

3.2 METHODS

There are certain statistics that are tallied during a game. For our analysis, we tallied statistics for each individual player during every shift when both teams were at full strength (playing 5-on-5 during regulation) for the first 40 games of the season. This means that we excluded all statistics that occurred during a power play (when a penalty had been called) or events that occurred during an overtime period. Eliminating time when a power play occurred resulted in about 10 minutes per game being excluded. The statistics that we used for the analysis were the following:

Goal for (GF): a goal is scored for the team

Goal against (GA): a goal is scored against the team by the opponents

Shot on goal (SOG): the player in question makes an attempt at scoring a goal and the goalie is forced to make a save

Shot on goal by other player (SOGOP): another player makes a shot on goal

Shot on goal against (SOGA): the opponents have a shot on goal

Missed shot for (MSF): the player in question fails to score a goal and the goalie does not need to make a save. This means that the player shot the puck wide of the net.

Missed shot other player (MSOP): another player misses a shot

Missed shot against (MSA): a player on the opposing team misses a shot

Blocked shot for (BSF): the player in question has his shot blocked by another player on the opposing team who is not the goalie

Blocked shot other player (BSOP): another player has his shot blocked by another player on the opposing team who is not the goalie

Blocked shot against (BSA): the player in question blocks a shot made by the opposition

Hit for (HF): the player in question makes physical contact with a player that controls the puck, or just after that player had possession of the puck (this is a slightly subjective statistic as it is up to the stat-keeper to determine whether a hit was actually made)

Hit against (HA): the player in question was hit by an opposing player

Turnover Created (TC): the player in question stole the puck

Turnover against (TA): the player in question lost the puck

There are of course other statistics that are tallied during a game. The statistics that are listed here were chosen because the player in question had some effect on that event. For example, we did not include whether a faceoff was won or lost. A faceoff is when the referee drops the puck between two centers to begin play and each center tries to gain possession of the puck. A defensive player therefore does not have an effect on a faceoff.

We did include SOGOP, however, because a player may have passed the puck to the player who took the shot and therefore does have an effect on that play. This method for event selection is not the same rationale used in other papers where essentially every aspect of the game is included. Whether a faceoff was won and even where a faceoff occurs were considered for model selection in other studies (Schuckers et al., 2011).

While there may be an argument for including the result of a faceoff, we do not feel that the location and other factors are appropriate predictors because they are out of the

control of the players. We want to calculate player effect, so we look solely at events affected by players.

Creating a dataset that tallied the counts for the events for every shift for each player was no small feat. We had to use three separate datasets in order to generate the final database that was used for the analysis. Using both ESPN.com and NHL.com, we created spreadsheets for the time intervals for each of the players' shifts, the time intervals for each power play, and the time that each event occurred. These three files had to be created for each individual game. Using SAS, a program was written that used the information from these three original files and tallied the counts for each of the events previously listed. The events were counted for every shift of every player of every game.

3.3 STATISTICAL ANALYSIS

Using the final database, we used two logistic regressions that modeled whether a goal was scored for or against the Blue Jackets using the count for each event per shift as predictors (equations 3.1 and 3.2). Backwards selection was used to eliminate

Equation 3.1:

$$\log(\text{odds for GF}) = \beta_0 + \beta_1\text{BSA} + \beta_2\text{BSF} + \beta_3\text{BSOP} + \beta_4\text{HA} + \beta_5\text{HF} + \beta_6\text{MSF} + \beta_7\text{MSOP} + \beta_8\text{MSA} + \beta_9\text{SOG} + \beta_{10}\text{SOGOP} + \beta_{11}\text{SOGA} + \beta_{12}\text{TA} + \beta_{13}\text{TC}$$

Equation 3.2:

$$\log(\text{odds for GA}) = \delta_0 + \delta_1\text{BSA} + \delta_2\text{BSF} + \delta_3\text{BSOP} + \delta_4\text{HA} + \delta_5\text{HF} + \delta_6\text{MSF} + \delta_7\text{MSOP} + \delta_8\text{MSA} + \delta_9\text{SOG} + \delta_{10}\text{SOGOP} + \delta_{11}\text{SOGA} + \delta_{12}\text{TA} + \delta_{13}\text{TC}$$

insignificant events using a significance level of 0.1. This analysis was performed at 5 game increments for the first 40 games of the season. As the season progressed, the dataset was updated with data from the new games and the analyses were re-run. It was

possible, therefore, for the predictors in the model, and the coefficients for those predictors, to vary during the season. While BSA, HF, and SOG may have been significant predictors for the GF logistic regression for game 10, game 15 may have had BSA, HF, SOGOP, and TC as significant instead.

While other papers, such as Macdonald (2010) have included players in the model as fixed effects (equation 3.3), we chose to eliminate fixed effects from the model and

Equation 3.3:

$$y = \beta_0 + \beta_1 X_1 + \dots + \beta_J X_J + \delta_1 D_1 + \dots + \delta_J D_J + \gamma_1 G_1 + \dots + \gamma_K G_K + \varepsilon$$

where: J = number of skaters in the league
 K = number of goalies in the league
 y = goals per 60 minutes during an observation
 $X_j = \begin{cases} 1, & \text{skater } j \text{ is on offense during the observation} \\ 0, & \text{skater } j \text{ is not playing or not on offense during the observation} \end{cases}$
 $D_j = \begin{cases} 1, & \text{skater } j \text{ is on defense during the observation} \\ 0, & \text{skater } j \text{ is not playing or not on defense during the observation} \end{cases}$
 $G_k = \begin{cases} 1, & \text{goalie } k \text{ is on defense during the observation} \\ 0, & \text{goalie } k \text{ is not playing or not on defense during the observation} \end{cases}$

only use events as predictors. We believe that it is the usage of fixed effects to estimate player ability that causes the lack of precision in other studies because we become limited by the sample sizes for individual players. If we wished to analyze players throughout a season, we could not take this approach due to further decreased sample sizes.

What we did instead was take the average count for the significant events for each player and plug those means into the two logistic regression models. For example, let's assume that the GF logistic regression model at game 20 was:

$$\log(\text{odds}) = 0.3(\text{MSA}) - 0.1(\text{SOGA}) + 0.2(\text{SOG})$$

To estimate the probability that a goal was scored for the team when Rick Nash was on the ice, we would use the average MSA, SOGA, and SOG counts for each of Nash's

shifts. Assuming his averages were 0.1, 0.9, and 0.6 for MSA, SOGA, and SOG respectively, we would have the following estimate:

$$\log(\text{odds for GF}) = 0.3(0.1) - 0.1(0.9) + 0.2(0.6) = 0.06$$

Our estimate for Rick Nash at game 20 is therefore that there is a 51% chance of having a goal scored for the Blue Jackets each time Nash has a shift. This is of course a ridiculous estimate and is only just an example.

Plugging in the mean counts for each player may seem a bit ad hoc, or that we are indirectly measuring player effects. By choosing to only use event counts rather than fixed effects for players, we increased the sample size used for the regression and decreased the confidence interval widths. We might only have a few thousand observations per player at the end of a season using a fixed effects model. By removing fixed effects and solely using event counts as predictors, we increased the observations to over 20,000 in that same time period.

Unfortunately, there were issues encountered with the goodness of fit test for the logistic regression models. At certain points in the season, we found that there was a sufficient fit for the two models. After a few more games had been played and we ran the analysis again, however, the p-values became significant. The goodness of fit would change back and forth as more games were played and there did not appear to be any pattern or clear reason behind this.

Attempts to alleviate this issue were taken using various methods such as changing the link functions and using the lasso method, but none of these approaches were successful. The lasso method used a similar reasoning as the methods used in Macdonald (2011) and Gramacy et al. (2013). Macdonald used ridge regression as a

shrinkage model while Gramacy et al. used a Bayesian approach and set a prior distribution centered at zero on team partial effects and player partial effects. They imposed these priors to “guard from overfit and provide stable estimates of individual player effects”.

Macdonald was not successful in his ridge regression approach and has yet to suggest another solution that might further shrink his confidence intervals. Gramacy et al., on the other hand, were successful but as we have already pointed out, that study used 4 years of data. It was hoped that through using the lasso we would be able to gain the sufficient goodness of fit statistics. Considering that the two previous studies were unable to do this, it is not surprising that we too failed in this attempt.

The following protocol was used for estimating player effects as the season progressed in order to fix the lack of fit: if the p-value for the goodness of fit test was insignificant, the most current models were used to calculate the estimates using the most recent means for each player. If the goodness of fit tests were significant, then the most recent previous models that did have a sufficiently good fit were used instead. To attain estimates, we input the updated means for each player into the previous equations.

An example might make this protocol a bit more clear. Suppose that we are once again trying to estimate player abilities for week 20. If we ran a regression for GF and got a good fit, we simply plugged in the mean counts for the significant events to attain estimates for each player just as we did in the previous Nash example. If the fit was poor, however, we would go back to the most recent regression model that did have a good fit. For example, while using all the data up to game 20 might not result in a good fit, using all the data up to game 18 might. We would therefore use the regression models from

week 18, but plug in the updated mean counts from week 20. That way although we were using a dated regression model, we were still using updated means.

The estimates for each player from the two models reflected the probabilities that a goal was scored for or against the Blue Jackets when that specific player was on the ice. These two probabilities effectively measured the offensive and defensive abilities of the players. Using the estimates for each player, we took the probability that a goal was scored for the Blue Jackets and subtracted the probability that a goal was scored against the Blue Jackets (equation 3.4). If the resulting probability was positive, there was a

Equation 3.4:

$$\text{Complete Plus-Minus (CPM)} = \text{Probability of GF} - \text{Probability of GA}$$

greater chance a goal would be scored for the team when that player was on the ice; if the final probability was negative, there was a greater chance a goal would be scored against the team. This final statistic was termed the Complete Plus-Minus (CPM).

CHAPTER 4

RESULTS AND DISCUSSION

4.1 RICK NASH: THE FACE OF THE FRANCHISE

To begin the analysis of the Blue Jackets, we will first look at the face of the franchise, Rick Nash. Rick Nash is a 5 time All-Star, and plays on the Canadian National team. While he does have an elite status, he is sometimes criticized for disappearing during Columbus games. Looking at Nash's CPM over the first 40 games, we see that at no point does Nash have a positive CPM (figure 4.1). This means that he consistently hurts the team. Moreover, there is no trend to suggest that he is at least improving or moving towards a positive CPM at a significant rate.

Breaking the CPM into its offensive and defensive components, it is evident that Nash's offense does fluctuate during the first 40 games, but overall his offensive play is fairly consistent (figure 4.2). His offensive ability has an average probability of a goal being scored for the Blue Jackets of about 2%. His defensive play, on the other hand, is extremely troubling. While his defense is about as stable as his offensive production, the probability a goal is scored against the Blue Jackets is consistently higher (figure 4.3). His average probability for giving up a goal was about 3.5%. This is way too high a rate and he fails to at least cancel out his poor defensive play with offensive production. This is the tale for the entire team as well.




Figure 4.1 The CPM for Rick Nash over the first 40 games of the season. Nash's CPM is always negative indicating that Nash consistently hurts the team. There is a slight upward trend but nothing great enough to suggest that Nash would be able to attain a positive CPM by the end of the season.

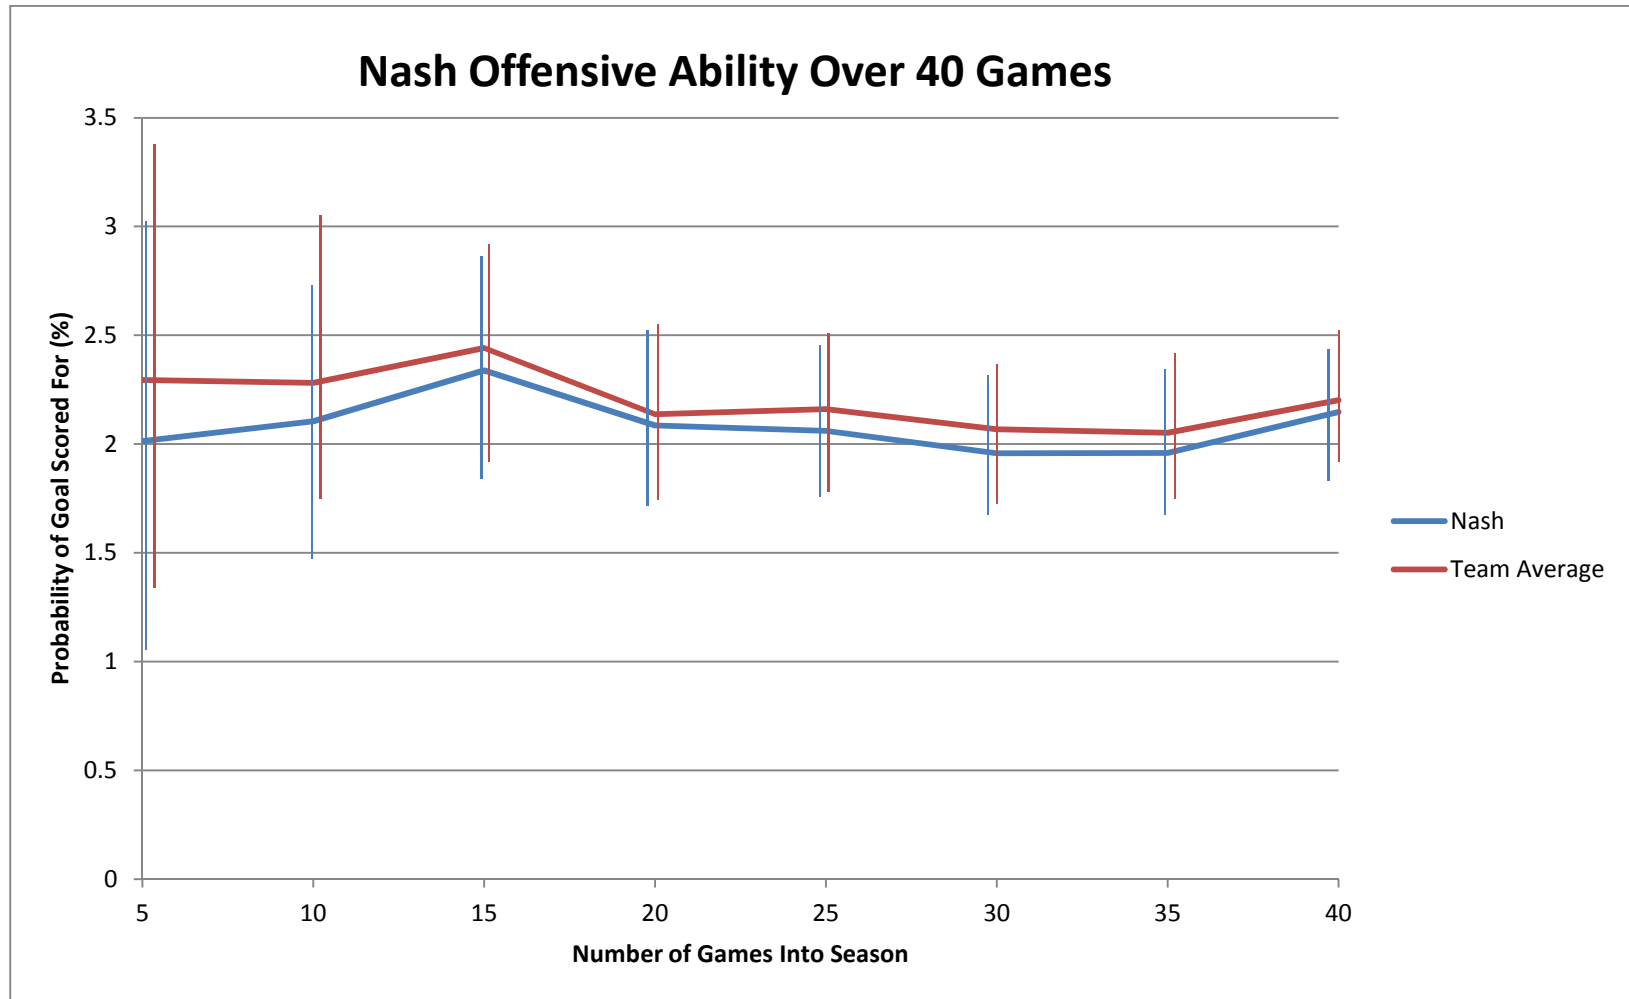


Figure 4.2 Rick Nash's offensive ability is consistent over the first 40 games of the season averaging around a 2% chance that a goal will be scored each shift when Nash is on the ice.

Nash Defensive Ability Over 40 Games

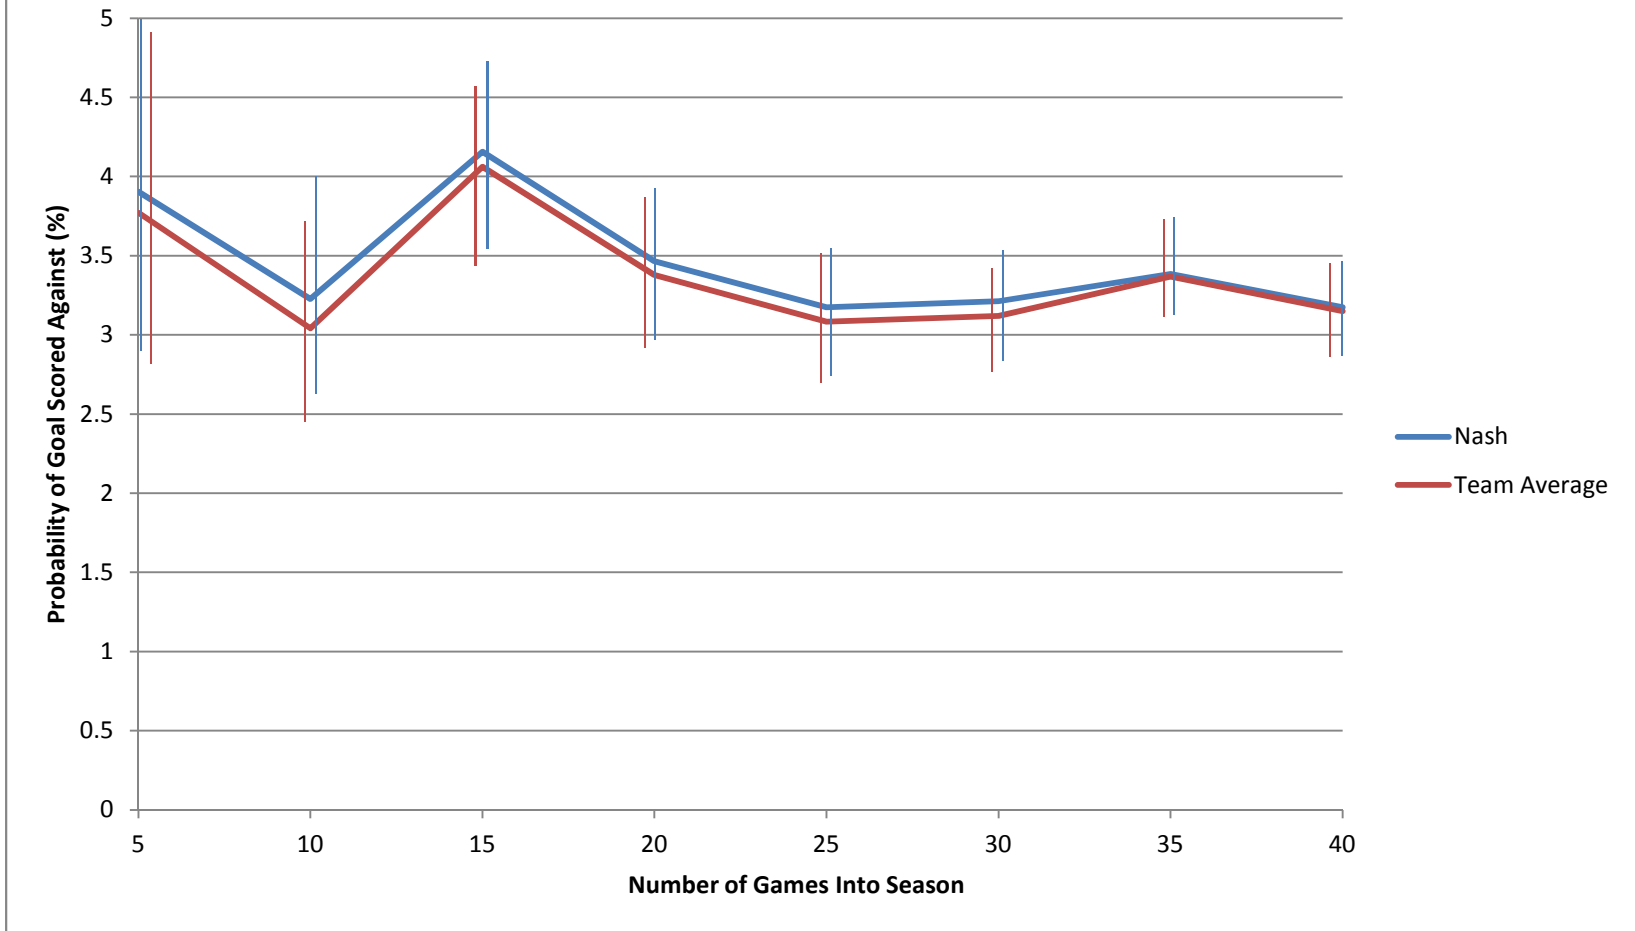


Figure 4.3 Rick Nash’s defensive ability averaged around 3.5%. This indicates that there is a 3.5% probability that a goal will be scored against the Blue Jackets when Nash is on the ice. This is 1.5% higher, on average, than his offensive ability. Nash’s poor CPM may be attributed to his poor defensive play.

4.2 THE STATE OF THE TEAM

If we look at the average CPM for the team, we see why the Blue Jackets played so poorly. At no point is the average CPM positive and there is great inconsistency (figure 4.4). This means that the Blue Jackets as a team always had a higher probability of giving up a goal than scoring a goal. This is more evident when looking at the offensive and defensive production overlaid on the same graph. Both the offensive and defensive probabilities remain approximately the same for the first half of the season, and the probability that a goal is scored by the opposition is around a full percentage point greater than the probability the Blue Jackets score (figure 4.5). We can compare this to the Maple Leafs, another team who did not make the playoffs.

Unlike the Blue Jackets, the Maple Leafs managed to either have their offensive production above or at least close to their defensive production over the first 40 games (figure 4.6). What this graph tells us is that while the offense consistently scores goals, the defense is the reason the Maple Leafs lose games as it is the defenses that worsens as the season progresses. In the case of the Blue Jackets, it is the defense that causes them to lose games as well. The offense needs to improve for the Blue Jackets, but the majority of the blame for the first half of the season falls squarely on the shoulders of the defensive play.

For both teams, these figures tell the story for the rest of the season. The Maple Leafs continued to play at a 0.500 win percentage, coinciding with our CPM after 40 games, and finished the season 35-37-10. The Blue Jackets also continued their trajectory

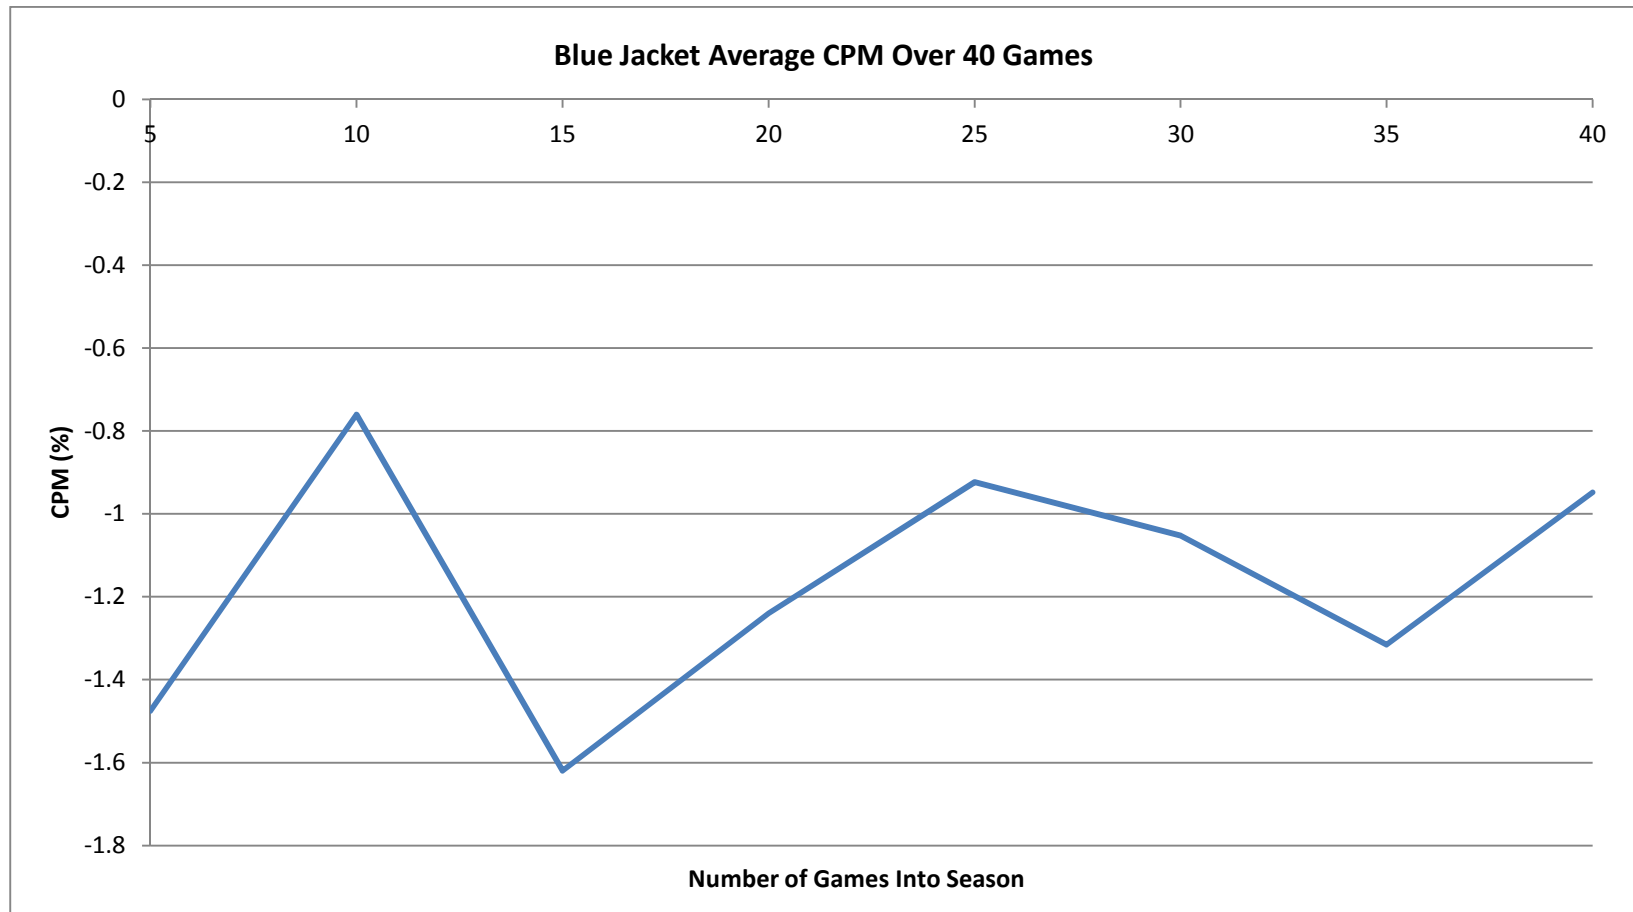


Figure 4.4 The Columbus Blue Jackets had a CPM that was always negative during the first 40 games of the season. While there are evident fluctuations, there does not appear to be an overall trend suggesting a change in the average CPM.

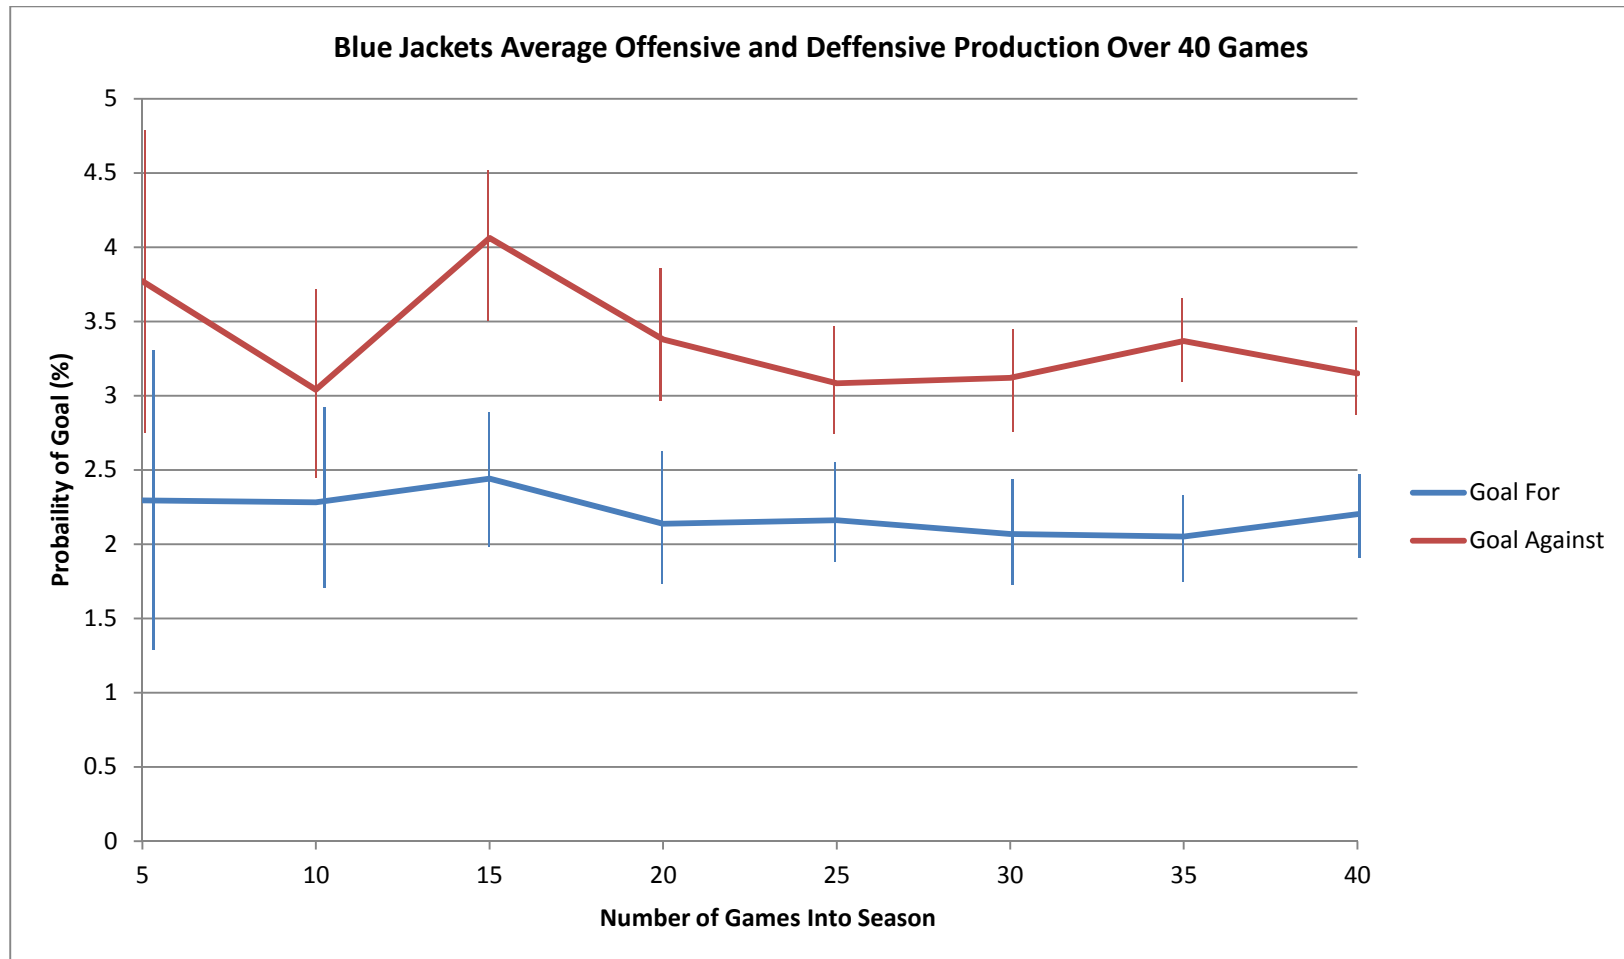


Figure 4.5 The offensive and defensive production for the Blue Jackets is plotted over the first 40 games. At no point does the probability that a goal is scored for the Blue Jackets become greater than the probability that a goal is scored against the Blue Jackets. Both probabilities are also constant suggesting that the current poor performance for the team may be expected to continue.

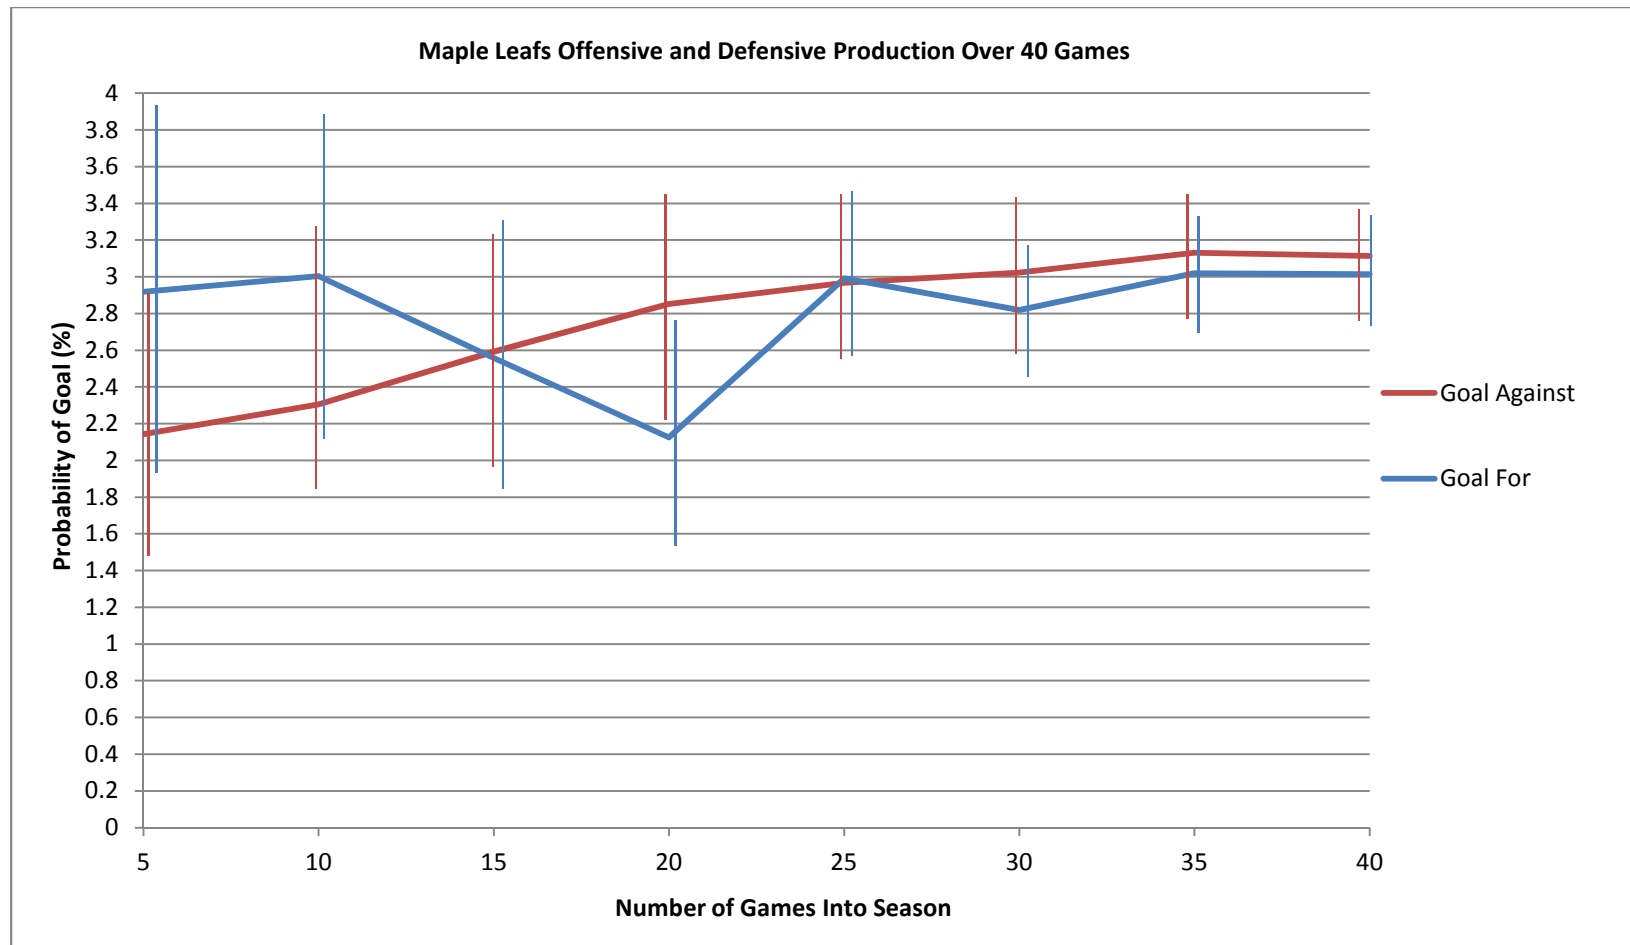


Figure 4.6 The offensive and defensive production for the Maple Leafs over the first 40 games of the season is plotted. The Maple Leafs start the season well with the probability of a goal being scored for the team being greater than the probability a goal is scored against the team. The defensive production worsens as the season progresses and once the team reached the 30 game mark, the two probabilities were equal.

towards mediocrity and finished last in the league with a record of 29-46-7. It may be argued that the fluctuation at the start of the season for both teams' CPM values is due to a small sample size. This is most likely the case because the trends for each team appear to stabilize as more games are taken into account (figures 4.5 and 4.6). If we had projected how each team would finish using the data at the 40 game point, we would have correctly projected that the Maple Leafs would finish around the 0.500 mark and the Blue Jackets would have far more losses than wins.

4.3 SUGGESTED MOVES

The biggest adjustment that could easily be made is to give Matt Calvert more playing time on a top line. Calvert is a prospect who has had some good playing time at the NHL level but plays sparingly for the most part. For the first 40 games of the season, Calvert has been one of the best players on the team according to the CPM. As of game 40, he had an average CPM of -0.834. While that number is still negative, it is the largest CPM on the team. Calvert does suffer from a small sample size, but with a team that clearly had no chance of making the playoffs, it certainly couldn't hurt to give a young player more playing time to prove his worth. We project such a move would only benefit the Blue Jackets.

Another move would be to trade both Rick Nash and Derick Brassard. Rick Nash was traded once the season ended (we analyze this trade in the next section), but the Blue Jackets may have been able to get more value for him if they had made a trade during the season instead. Brassard was a first round draft pick for the Blue Jackets and was expected to be a staple of the team for years to come, but Brassard ranked 19th on the team in CPM over the first 40 games. He hurts the team and while the Blue Jackets may

not bring in a huge haul for him, there should be interest from other teams in a former first rounder.

Another easy adjustment Columbus could make is to change the line combinations. The following lines would maximize the goal differential by giving the players that have the largest CPM the most playing time. (CPM values are given next to the names of the players.)

Line 1: Matt Calvert (-0.834), Antoine Vermette (-0.991), Cam Atkinson (-1.1)

Line 2: Ryan Russell (-1.11), Ryan Johansen (-1.072), Vinny Prospal (-1.19)

Line 3: Sammy Pahlsson (-1.14), Jeff Carter (-1.239), R.J. Umberger (-1.219)

Line 4: Derick Brassard (-1.294), Derrick Mackenzie (-1.353), Derrick Dorsett (-1.31)

Defense 1: Nikita Nikitin (-0.965), Fedor Tyutin (-1.02)

Defense 2: David Savard (-1.149), John Moore (-1.247)

Defense 3: Mark Methot (-1.293), James Wisniewski (-1.331)

Notice that Nash is not even listed in the suggested lineup. This is because his CPM is so low every player listed on offense is a better player according to the CPM. This again gives reason to why he should be traded. By getting multiple players in return, it would also make more line combinations possible to further maximize the goal differential.

4.4 ANALYSIS OF ACTUAL TRANSACTIONS

There were numerous moves made by the Blue Jackets to try and right the ship. They got rid of Grant Clitsome, Jeff Carter, Antoine Vermette, Rick Nash, and changed their head coach. The question is, were these the right moves? Taking a look at these four

players and comparing their CPM values to the team average CPM over the first 40 games should give a better idea as to whether these moves improved the team.

Clitsome was put on waivers just after the 30 game mark. As we can see, he was almost exactly average throughout the first half of the season (figure 4.7). The same can be said about Carter (figure 4.7). For both these players and for Carter in particular, expectations were that they would be above average and in fact leading the team in the right direction. As the CPM shows, neither helped the team and both failed to even be above average players for the Blue Jackets.

We can see that Vermette was above the team average and in fact ranked 4th on the team for CPM (figure 4.7). It can be argued that he should be expected to rank 4th given the amount of money he was being paid. It became clear, however, that the team was moving towards rebuilding as they neared the break. If a team is looking to rebuild, they must give up a valuable asset in order to stockpile draft picks and prospects. Vermette was traded for a 2012 second-round pick, a conditional 2013 fifth-round pick and a goalie. The Blue Jackets did get good value for Vermette and it was the right move towards building a better team in the future.

Then of course we have perhaps the biggest trade in Blue Jackets history. There was a lot of controversy concerning whether the GM, Scott Howson, was waiting too long or looking for too much to trade Rick Nash. He had 30 goals and 29 assists at the end of the season. One would think that he must help the team at least a bit, but the CPM says otherwise. Not only did Nash never achieve a positive CPM, he never even got above the team average (figure 4.7). If you asked anyone last year who the top 3 players were on the Blue Jackets every person would have Nash on that list. That was not reality.

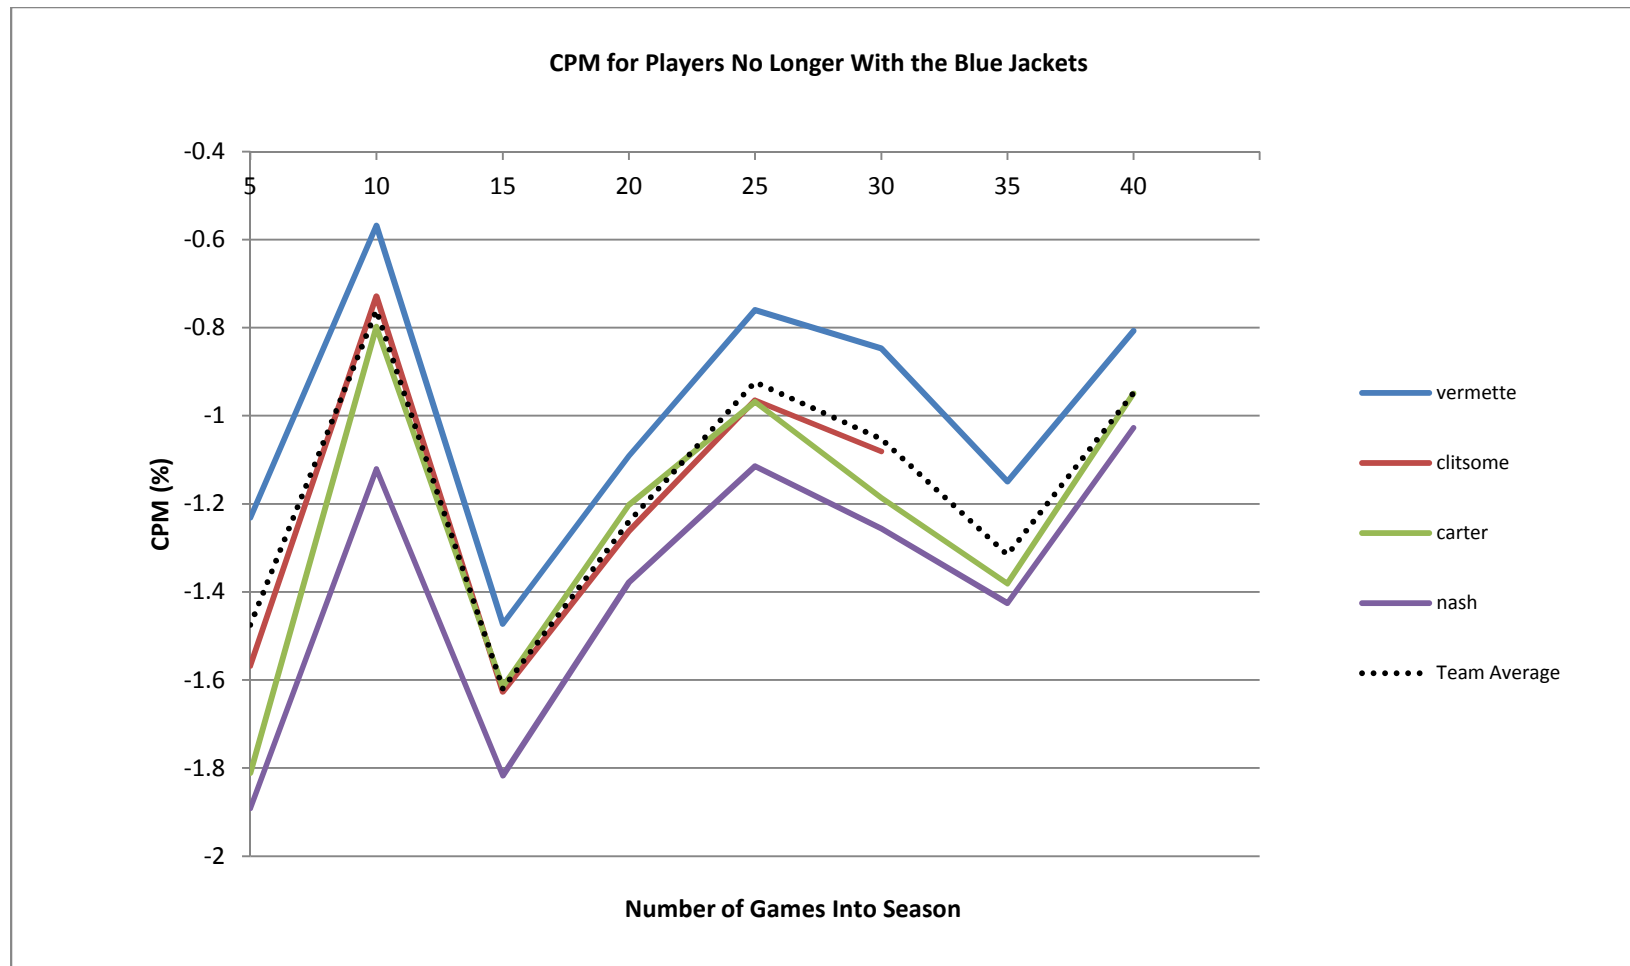


Figure 4.7 The CPM values for Antoine Vermette, Grant Clitsome, Jeff Carter, Rich Nash, and the team average are shown for the first 40 games. Vermette is well above the average while Nash is far below it. Clitsome and Carter are almost perfectly average throughout the first half of the season.

By trading Nash, the Blue Jackets shockingly improved the team. To get younger players that the team can use right away and build around was a best case scenario and it was only a question of who those players would be. The two moves of trading Carter and Nash were the best moves of the season while the decision on Clitsome was a questionable one.

CHAPTER 5

CONCLUSION AND FURTHER STUDIES

The Columbus Blue Jackets are anything but synonymous with success. Their attempts to make trades or draft players that will change the team's fortunes have mostly been met with failure. Something must be done differently and that something is the use of hockey sabermetrics.

Using the CPM to analyze the Columbus Blue Jackets, we were able to conclude that the poor record of the team during the first half of the 2011-2012 season was mostly due to a defense that gave up too many goals and players underachieving. Those players, such as Rick Nash, mostly struggled to equal their offensive production with their defensive ability. We also found that certain moves made by the Blue Jackets bettered the team; moves like letting Grant Clitsome leave on waivers, however, were mistakes.

During the current 2012-2013 season, the Blue Jackets have so far shown improvement. The players that were acquired for Rick Nash are contributing to the team. It is counterintuitive that shipping off a perennial All-Star would improve a team, but just as the CPM suggested, that is exactly what has happened. The team is not yet of a playoff caliber, so there are still many more moves that must be made in order to reach that point.

Future studies should focus on correcting the goodness of fit tests so that the two regression models fit sufficiently well throughout the entire season. We may want to include effects for opposing goalies and opposing teams. It remains in question whether this would greatly improve the model, however, since the Blue Jackets play most of the

teams the same amount. It would also be important to study other teams in the NHL to see if the same analysis works universally. We could also begin to analyze trade scenarios using data from other teams. For example, if we had data for the players that came in the Rick Nash trade, we could predict their CPM for the Blue Jackets and project if they would contribute to the team even before the trade was made.

Using the CPM, it might also be useful to take a longitudinal approach to predict how players and the team will play later in the season. This could prove extremely useful as we could predict whether the team in its current state has a chance of making the playoffs. It would also be possible to determine what players on other teams would help make the chance of reaching the playoffs greater if they were acquired through trades. These improvements would naturally require generalizations of the model explained here.

The CPM is a statistic that is team specific and is able to measure how well players fit with that team. We can rank the players on a team which should help with making minor roster changes. Dropping the lower level players and slowly increasing the team average CPM is the way to form a playoff caliber team. We can also project how well players on other teams will fit with the Blue Jackets which allows us to determine what players are the best to make trades for. The CPM is not a tool for fans to rank players in the NHL; it is a tool for a general manager. Just as Billy Beane revolutionized baseball with his sabermetrics, we hope that this too can help change the way hockey looks at players. All it takes is a general manager who isn't afraid to think outside the box and someone who can think of a better title than "Money puck".

REFERENCES

Gramacy, R. B., Jensen, S. T., and Taddy, M. (2013). “Estimating Player Contribution in Hockey with Regularized Logistic Regression.” Tech. rep., arXiv:1209.5026.

Macdonald, B. (2010). “A Regression-based Adjusted Plus-Minus Statistic for NHL Players.” Tech. rep., arXiv: 1006.4310.

Macdonald, B. (2011). “An Improved Adjusted Plus-Minus Statistic for NHL Players.” MIT Sloan Sports Analytics Conference, 2011.

Schuckers, M. E., Lock, D. F., Wells, C., Knickerbocker, C. J., and Lock, R. H. (2010). “National Hockey League Skater Ratings Based upon All On-Ice Events: An Adjusted Minus/Plus (AMPP) Approach.” Tech. rep., St. Lawrence University.

Vollman, R. (2010). “Howe and Why: Ten Ways to Measure Defensive Contributions.” *Hockey Prospectus*, March 4, 2010.

Wheeler, G. (2010). “The Lasso Logistic Regression Model: Modifications to aid causality assessment for Adverse Events Following Immunization.” London School of Hygiene & Tropical Medicine, 2010.