National Hockey League Skater Ratings Based upon All On-Ice Events: An Adjusted Minus/Plus Probability (AMPP) Approach

Michael E. Schuckers[1,2], Dennis F. Lock[3], Chris Wells[1,2], C. J. Knickerbocker[1,4], Robin H. Lock[1]

[1]*St. Lawrence University*

[2]*Statistical Sports Consulting, LLC*

[3]*Iowa State University*

[4]*Sensis Corporation*

Abstract

In this paper we present a new methodology for evaluating and comparing National Hockey League non-goalie performance based upon each play for which a given player on the ice. This methodology is based upon a least squares regression that adjusts for other players that are on the ice for each play. The value we assign to each play is the difference between the expected outcome for that play and the observed outcome. The specific outcome we study is an indicator of whether or not a goal was scored in 10 seconds following the play. The predictors in our model are indicators of which players are on the ice for a given play. Our regression adjusts for all players on the ice for a given set of plays. Each player is then ranked based upon their average rate at which they contribute to the response. We apply this methodology to each play of the 2006-7 National Hockey League season and present results for the top ten forwards and defensemen based upon these ratings.

Keywords: regression, play-by-play, moneyball, time on the ice, player ratings

I. Introduction

The goal of this paper is to develop a single methodology for ranking and comparing National Hockey League (NHL) skaters based upon all the plays for which they were on the ice. Currently, the typical measures for comparing NHL players are plus/minus or points scored. Plus/minus is the number of even strength or shorthanded goals scored by his team while a player is on the ice minus the number of similar goals scored by the other team when that player is on the ice. This measure does not account for the complete range of times when a player is on the ice or reflect positive or negative contributions that did not result in a goal being scored.

Points are the total number of assists and goals attributed to each player. This does not take into consideration, for example, the other players on the ice, defensive play, power play performance or the total number of penalties committed.

Mason and Foster (2007) suggest that one reason for the lack of a single metric to measure player performance is a lack of quality data. Below we present an analysis that makes use of the NHL's Play-by-play files (National Hockey League, 2007) which describes every play that the NHL recorded for the 2006-7 season and which are freely available for download. Our regression model extends the adjusted plus-minus work described by Ilardi (2007) and Rosenbaum (2004) for basketball by adjusting the response for the expected outcome of a given play. We have as our observational unit each play that is recorded by the NHL rather than doing our analysis over different changes in personnel as Ilardi and Rosenbaum do.

There are two main results for this paper. First, we develop methodology for assigning value to each play recorded by the NHL and, second, we combine those values to establish ratings for each player. Our paper is organized in the following manner. Section II describes the model we are proposing and the data we use to evaluate skaters or non-goaltenders based upon this model. We apply this model to data from the 2006-7 NHL season in Section III and give ratings for the top forwards and defensemen. We discuss these results as well as ramifications and limitations of our approach in Section IV.

II: Methodology and Description of the Data

A. Description of the Data

We analyzed every play from the 2006-7 NHL regular season by combining information from the play-by-play (PBP) files and time on the ice (TOI) files for all 1200 games during that season. For examples of these files, see National Hockey League (2007a) and (National Hockey League, 2007b). These files were downloaded from nhl.com. A play, as we define it, is an event that the NHL records in the PBP files from each game. Events are the possible values that a play can take and are listed in Table 1. For this analysis, we focused on the following events: Blocked Shot, Faceoff, Giveaway, Hit, Missed Shot, Penalty, Shot, Takeaway and Goal because these plays reflect 'on ice during game action' events. Thus, Goalie Pulled and Stoppage were not considered as events since they did not indicate a 'during game' event (Stoppage) or because they were accounted for by information in the TOI files (Goalie Pulled).). Faceoff Won (L) is the event that a team wins a faceoff in Zone L. There are three possible zones: Defensive (Def), Neutral (Neu) and Offensive (Off). Additionally, because shootout performance is inherently different from other events we excluded events that occurred during a shootout.

Table 1: List of events in NHL PBP Files, 2006-7

Blocked Shot	Face-Off (L)	Giveaway	Goal
Goalie Pulled	Hit	Missed Shot	Penalty
Shot	Stoppage	Takeaway	

Having extracted the plays from all regular season games, we then processed these files so that each play was accompanied by the names of each of the players on the ice (from among all n=1053 possible players) for that play as well as the time that each play occurred. Each player's team was also ascertained and which of the teams playing was at home. This information was recorded for w=359,222 plays. Note that among these 1053 players were goalies who are classified here as non-skaters. From the data above, we created a 359,322 x 1,053 player matrix, X, where each of the 359,322 rows represents each play of the season, and each of the 1,053 columns represents each individual player. We note here that players who appeared with more than one team are listed as different players.

B. The Model

The model that we are proposing here is a least squares model similar to the adjusted plus/minus proposed by Rosenbaum (2004) and Ilardi (2007) for basketball. These authors use as the response for their model the change in score, Y, when each group of players is on the field of play. This model can be summarized as follows:

$$Y_i = \theta_0 + \sum \theta^{(H)}{}_j - \sum \theta^{(A)}{}_m \quad (1)$$

where Y_i is the margin of difference in the score for the i^{th} unit of time in a game where no substitutions were made, θ_0 is included in the model for those that play less than so small number of minutes, $\theta^{(H)}{}_j$ is the rating for the j^{th} player for the home team and $\theta^{(A)}{}_m$ is the rating for the m^{th}

player for the away team (Rosenbaum, 2004). We propose an extension of the Rosenbaum/Ilardi model . Given the different average scoring rates for the two sports -- basketball and hockey -- it is perhaps not surprising that we use a different model. The value that we will assign to a given play is based upon the probability that either team will score in a finite interval following play *i*. Further, we subtract the expected or average outcome, $E[Y_i]$, of play *i* from the observed outcome, Y_i, of a play *i* to give each play a value relative to the expected outcome from that play. Below we will discuss assigning value to the i^{th} play. Thus, we reward players for outperforming the expected value for plays in which they were involved. In a way this standardizes the player ratings by referencing our ratings based upon each play rather than on a set of players who appear in games a fraction of the time. We follow Rosenbaum and Ilardi in rating players by accounting for all other players with whom they appear on the ice. Thus, our ratings indicate how well a player performs relative to what we would expect for the events that occurred while they were on the ice and adjusts for the other players on the ice with them. This general approach is explicitly given by Equation (2).

$$Y_i - E[Y_i] = \sum \theta^{(H)}{}_j - \sum \theta^{(A)}{}_m . \qquad (2)$$

For a given combination of players credit is given only for consistently outperforming the average performance for a given play. Thus, the player who performs at an average rate adjusting for the other plays on the ice with them will receive a rating of zero. To that end, we use as our outcome for play *i*, Y_i, a difference of two indicator functions; whether or not the home team scores for a given period of time minus whether or not the away team scores in that same time interval. Then $E[Y_i]$ is a difference in cumulative probability between the home team scoring and the away team scoring over that interval.

As mentioned above the scoring rate for the NHL (5.89 goals per game) is much less than for the National Basketball Association (197.48 points per game) for the respective 2006-7 regular seasons. As a consequence we chose to focus on each event that occurs and the likelihood that it leads to a goal for a given team within k seconds. Then expanding on Equation (2) our model becomes

$$(I_H(i,k) - I_A(i,k)) - (P_{GH}(i,k) - P_{GA}(i,k)) = \sum \theta^{(H)} j - \sum \theta^{(A)} m \qquad (3)$$

where $I_H(i,k)$ and $I_A(i,k)$ are indicator variables that take the value 1 if a goal was scored and a 0 otherwise in the *k* seconds following play *i* by the home (H) and away (A) team, respectively, and $P_{GH}(i,k)$ and $P_{GA}(i,k)$ are the cumulative probabilities that play *i* will lead to a goal for the home team and for the away team in the next *k* seconds, respectively.

To estimate the ratings, θ's, for each player we write the above equation as a system of linear equations

$$Y_i - E[Y_i] = \theta_1 X_{i,1} + \theta_2 X_{i,2} + \ldots + \theta_n X_{i,n},$$

$X_{i,j}$ is the indicator of whether the j^{th} player was on the ice for the i^{th} play such that

$$X_{i,j} = \begin{cases} 1 & \text{if } j^{th} \text{ player is on the ice for the home team during play i,} \\ 0 & \text{if } j^{th} \text{ player is not on the ice for play i,} \\ -1 & \text{if } j^{th} \text{ player is on the ice for the away team during play i.} \end{cases}$$

θ_j is the average contribution of the j^{th} player to the outcome of the i^{th} play adjusting for the

Figure 1

D(t)-D(t-1) for t seconds following an event

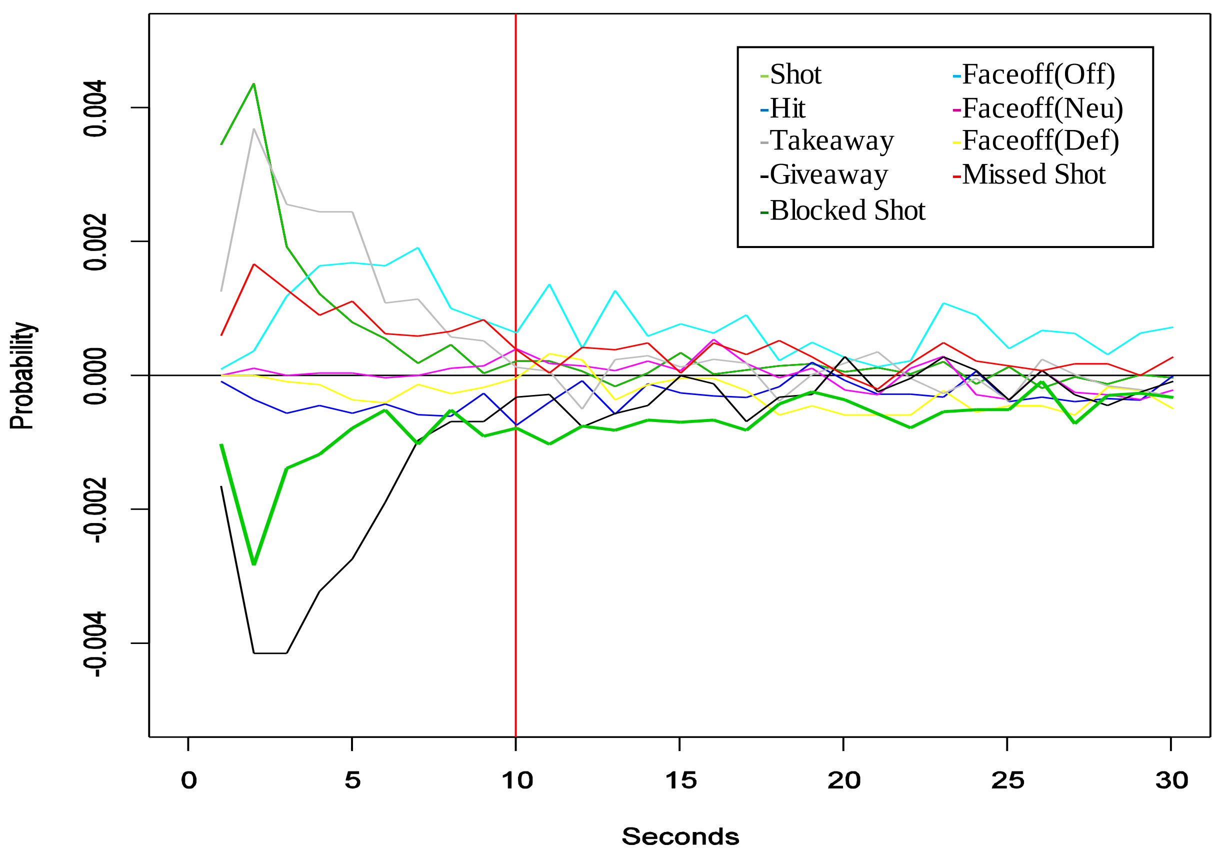

contributions of the other players on the ice for the i^{th} play. By having the X_{ij}'s in the model we can identify each player's individual average contribution to Y_i because we have accounted for the contribution of the other players who are on the ice. In order to solve for these coefficients, the θ_j's, we set this system of linear equations as a regression equation

$$\Delta = X\theta \tag{4}$$

where Δ is a vector $(Y_1-E[Y_1],\ldots,Y_w-E[Y_w])^T$ of probability values, X is a w by n matrix comprised of the values $X_{i,j}$ and $\theta = (\theta_1,\ldots,\theta_n)^T$. To estimate θ, we use least squares (Mendenhall & Sincich, 2003).

The vector Δ represents the standardized value of each play with, $\Delta_i = I_{GH}(i,k) - I_{GA}(i,k) - (P_{GH}(i,k)-P_{GA}(i,k))$. In order to determine the appropriate value for *k*, we examined the change in $P_{GH}(i,k) - P_{GA}(i,k)$ for a range of *k*'s. For example in order to find the value of a takeaway we examine all 17,634 takeaways from the 2006-7 season and find the proportion of plays that the team making the takeaway scores and the team that was the victim of the takeaway scores in exactly *t* seconds. Thus, we investigated the change in $D(t)=P_{GH}(i,t) - P_{GA}(i,t)$ for each event for t=1,…,60 seconds. A graph of D(t)-D(t-1) against *t* (t=1,…,60) is given in Figure 2 for each event. Based upon this graph we chose k=10 since the change in probability for these events seems to stabilize to zero at approximately 10 seconds. That is, it appears that there will be little change in the cumulative difference in probabilities, $P_{GH}(i,t) - P_{GA}(i,t)$, after 10 seconds for each event. We note that these conclusions are similar to those found in Thomas (2006) where the slopes of cumulative probabilities due to different events were relatively constant after approximately 10 seconds. The only play with which we choose a different value of *k* is penalties, where the value of k is the length of the penalty.

Having established k=10, we now summarize the values assigned to each event. Table 2 contains summaries of the values assigned to different events. There we denote that team committing the event as S and the opponent as O. Thus if the home team takes a shot and misses, the value assigned to that play is 0.0094, but if it is the away team (A) that misses a shot, the value assigned to that play is -0.0094. This is done so that values are assigned appropriately to players on the home and away teams according to Equations (2) and (3).

Table 2: Values of each play based on P_{GS}-P_{GO}

Play	$P_{GS(i,10)}$	$P_{GO(i,10)}$	$E(Y_i)$
Blocked Shot	0.0040	0.0164	-0.0124
Faceoff Won (Def)	0.0011	0.0027	-0.0017
Faceoff Won (Neu)	0.0013	0.0006	0.0007
Faceoff Won (Off)	0.0119	0.0009	0.0110
Giveaway	0.0021	0.0247	-0.0226
Hit	0.0039	0.0085	-0.0046
Missed Shot	0.0125	0.0029	0.0094
Penalty	0.0225	0.1759	*-0.1534
Shot	0.0171	0.0024	*0.1226
Takeaway	0.0205	0.0029	0.0176

* *Note that E(Yi) for shots, penalties and blocked shots is based upon a different calculation than the other events.*

One final adjustment was made to the value of a shot, since a shot has a probability of scoring on its own. So the value of a shot is as follows,

$$E(Y_i) = \frac{\#(Goals)}{\#(Shots)+\#(Goals)} + P_{GH}(i,10) - P_{GA}(i,10).$$

when play i is a shot.

III. Results

Estimates of ratings for the j[th] player, $\hat{\theta}_j$, were obtained for all 1053 players via a least squares solutions to Equation (4). $\hat{\theta}_j$ is the estimated change due to player j in the difference in team scoring probability for each 10 second interval that player j was on the ice relative to an average player adjusting for the other players on the ice. Because of the nature of the response we refer to these ratings as the adjusted minus/plus probability (AMPP) for each player. Tables 3 and 4 give the top ten estimated ratings for forwards (centers or wings) and defensemen, respectively, who played more than 1000 minutes (TOI>1000) in the 2006-7 regular season.

Because of the large number of plays that occur while goalies are on the ice for which they are not involved we have not included goalies here. See Desjardins (2009) for a proposed methodology for assessing goalie performance by adjusting save percentage for difficulty of shots faced. Each table of skaters below also gives players overall rank as well as the traditional +/- for each players as well as the player's team for the 2006-7 season. We note that as expected there is little relation between a player's traditional +/- and their AMPP. Thus, Thomas Vanek listed as the top forward had an AMPP of 2.47% meaning he increased the average difference in the probability of a goal for his team by 2.47% over what the average player did for every 10 seconds he was on the ice accounting . Similarly Chris Pronger the highest rated defensemen had an AMPP of 1.40% and so he increased this difference by 1.40% for every 10 seconds he was on the ice over an average player adjusting for other players on the ice with Pronger. As would be expected there was a wide range of values for all 1053 players. All of the top players had positive values but many players had negative values indicating that they decreased the probability of their team scoring a goal. The range of values for skaters who were on the ice for more than 1000 minutes was from 0.0247 to –0.0235. The standard deviation for these ratings among skaters is 0.0083.

To assess the stability and reliability of our ratings, we studied the players who performed for multiple teams during the 2006-07 NHL season. If our method does indeed adjust for other players on the ice, then we should expect that these ratings are similar. One of the advantages of our adjusted approach is to permit such a comparison of the different ratings, $\hat{\theta}_j$'s, that players possess when switching teams. The mean absolute difference between AMPP's for traded players with more than 1000 minutes during the 2006-7 season is 0.0070. Table 5 lists a sample of players traded during the 2006-7 NHL season as well as their estimated AMPP ratings for both teams. Since the mean absolute difference is smaller than a standard deviation of the player ratings For all of these players there is little difference between their rating with the different teams suggesting that an individual's rating is relatively stable and independent of the other players on the ice with a given individual.

Table 3

AMPP for Top Ten Forwards

Overall Rank	Team	Player	Plus/Minus	AMPP
1	BUF	T. Vanek	+47	0.0247
2	PIT	S. Crosby	+10	0.0192
3	NJD	J. Madden	-7	0.0177
4	NJD	S. Gomez	+7	0.0176
5	BOS	S. Donovan	-13	0.0175
6	NJD	T. Zajac	+1	0.0170
7	PIT	J. Staal	+16	0.0166
8	PIT	M. Talbot	-2	0.0146
9	PHX	O. Nolan	-2	0.0143
10	ANA	S. Pahlsson	-4	0.0140

IV. Discussion

We have presented a methodology for rating NHL skaters based upon every play that was recorded in the 2006-7 NHL PBP and TOI files. For each play we assign a relative value based upon the difference between the observed outcome and the expected outcome where outcome is measured as the probability that that play will lead to a goal for the team committing the play against the probability that the opposing team will score within 10 seconds. We used 10 seconds based upon an analysis of the length of impact for each play (Figure 1) and following work by Thomas (2006). The response that we chose here is the difference in probabilities that each team scores a goal within 10 seconds for the events: Faceoff, Giveaway, Hit, Missed Shot, Takeaway and Goal. For a penalty we use a larger time interval due to the consequence of taking a penalty. For Blocked Shot and Shot we adjust to account for the chance that those events would lead to a goal. We note that four events – Faceoff (Def), Giveaway, Hit and Penalty --have negative values suggesting that these events are better for the opposing team.

The model that was used to derive our estimates of a rating for a given player is based upon a linear regression that accounts for the other players on the ice with each player for each play. This aspect of our model borrows its approach from work by Rosenbaum (2004) and Ilardi

(2007) for basketball. The resulting ratings, the AMPP's, allow direct comparison of the values of each skater.

To calculate an AMPP for each player we used data available in the 2006-7 regulare season PBP and TOI files. In subsequent seasons, the NHL has added additional information to the PBP files that may allow for model improvement. In particular that data includes the zone (Offensive, Neutral or Defensive) where a play occurs, the situation (A on B where A and B are the number of players on the ice) as well as information on each shot: distance and shot type (Wrist, Snap, etc.). Here since we lack that information we are effectively averaging across this additional information by deriving the marginal probabilities. While our comprehensive play approach is an improvement over current approaches, the primary weakness in this regard is that our method for dealing with non-even strength situations, i.e. those where the number of individuals on the ice for both teams is different, is implicit. For events that occur during these situations we distribute the value of those plays to only those on the ice at the time.

The response here, Y_i-$E[Y_i]$, is one that assigns positive value to a play for exceeding the expected outcome and negative value for failing to achieve the expected outcome. If a player is succeeding – producing goals -- at the league average rate then the total, and hence, the average value of the plays for which they are on the ice will be zero. If a player is exceeding the expected rate of success for the plays that they are on the ice then the total value assigned to those plays here will be positive. One possible alternative response that could be used is to divide our current response by its' standard deviation and, thus, create a z-score type response. This would create a normalized response. Given the binary response here it seems that this is not necessary and would alter the interpretation of the ratings to those based upon standard deviations from expected rather than those given above.

The ratings derived from the AMPP method give here are potentially a useful tool for evaluating and for appraising a players complete contribution to their team in all facets of on-ice performance. These ratings are potentially valuable for making personnel decisions and for finding players who performance is not reflected by current measures but who consistently contribute to positive outcomes for their team.

Table 4

AMPP for Top Ten Defensemen

Overall Rank	**Team(#)**	**Player**	**Plus/Minus**	**AMPP**
13	ANA	C. Pronger	+27	0.0140
15	NYR	M. Malik	+32	0.0136
18	PHI	R. Jones	-14	0.0126
21	NJD	B. Rafalski	+4	0.0123
26	PHI	J. Pitkanen	-25	0.0120
37	PHI	D. Hatcher	-24	0.0098
45	ANA	F. Beauchemin	+7	0.0134
48	BUF	D. Kalinin	+19	0.0089
50	NJD	J. Oduya	-5	0.0088
68	BOS	Z. Chara	-21	0.0080

Table 5

Traded Players AMPP Ratings Before and After

Player	**Before Team**	**After Team**	**Before Ratings**	**After Ratings**
M. Jurcina	WAS	BOS	0.0021	0.0052
P. Mara	BOS	NYR	0.0085	0.0071
L. Nagy	DAL	PHX	0.0017	0.0023
W. Primeau	CGY	BOS	0.0068	0.0069
R. Robitaille	PHI	NYI	0.0002	0.0036
R. Smyth	EDM	NYI	0.0095	0.0047
L. Vishnevsky	NSH	ATL	-0.0056	-0.0056
J. Williams	CHI	DET	-0.0022	-0.0016

References

Desjardins, G. (2009, April 24). *2008-09 5v5 Goaltender Performance*. Retrieved from Behind the Net: http://www.behindthenet.ca/blog/2009/04/2008-09-5v5-goaltender-performance.html

Ilardi, S. (2007). Adjusted Plus-Minus: An Idea Whose Time Has Come. *http://www.82games.com/ilardi1.htm* . www.82games.com.

Mason, D. S., & Foster, W. M. (2007). Putting Moneyball on ice? *International Journal of Sport Finance , 2* (4), 206-213.

Mendenhall, W., & Sincich, T. L. (2003). *A Second Course in Statistics: Regression Analysis* (6th Edition ed.). Prentice Hall.

National Hockey League. (2007a). *NHL Play-by-Play File for Game 1217*. Retrieved from NHL.com: http://www.nhl.com/scores/htmlreports/20062007/PL021217.HTM

National Hockey League. (2007b). *Time on the Ice File for Game 1217*. Retrieved from NHL.com: http://www.nhl.com/scores/htmlreports/20062007/TH021217.HTM

Rosenbaum, D. T. (2004). *Picking the Difference Makers for the All-NBA Teams*. Retrieved from 82games.com: www.82games.com/comm30.htm

Thomas, A. C. (2006). The Impact of Puck Possession and Location on Ice Hockey Strategy. *Journal of Quantitative Analysis in Sports , 2* (1).

Vincent, C. B., & Eastman, B. (2009). Defining the Style of Play in the NHL: An Application of Cluster Analysis. *Journal of Quantitative Analysis of Sport , 5* (1).