

Assessing the offensive productivity of NHL players using in-game win probabilities

Stephen Pettigrew
Harvard University,
Cambridge, MA, USA, 02138
Email: pettigrew@fas.harvard.edu

Abstract

Hockey journalists and statisticians currently lack many of the empirical tools available in other sports. In this paper I introduce a win probability metric for the NHL and use it to develop a new statistic, Added Goal Value, which evaluates player offensive productivity. The metric is the first of its kind to incorporate powerplay information and is the only NHL in-game win probability metric currently available. I show how win probabilities can enhance the narrative around an individual game and can also be used to evaluate playoff series win probabilities. I then introduce Added Goal Value which improves upon traditional offensive player statistics by accounting for game context. A player's AGV has a strong positive correlation between seasons, making it a useful statistic for predicting future offensive productivity. By accounting for the context in which goals are scored, AGV also allows for comparisons to be made between players who have identical goal-scoring rates. The work in this paper provides several advances in hockey analytics and also provides a framework for unifying current and future work on Corsi, Fenwick, and other NHL analytics.

1 Introduction

Despite progress that has been made in recent years, NHL executives, writers, and fans lack many of the analytical and empirical tools that are available for analyzing baseball [1], basketball [2], football [3], and other sports [4]. This paper introduces two new analytical tools for the NHL. The first is a metric for calculating in-game win probabilities for NHL games. The second tool, Added Goal Value, evaluates offensive contributions by a player in terms of the impact that they have on the team's probability of winning the game. Throughout the paper I demonstrate the value that these two statistics can have for journalists and bloggers, as well as team executives who are looking to improve their rosters.

The paper is organized as follows. In section 2 I derive the win probability metric. To my knowledge it is the only one currently available for calculating in-game probabilities for NHL games. It is built on a framework which can easily unify new empirical findings and new sources of data, like spatial data which is widely used in other sports.

In section 3 I demonstrate how the win probability metric can be used by somebody writing about an individual game. The metric provides empirical estimates of momentum swings that occur as a result of a goal or penalty, which may be useful in building a narrative about a game [5]. In this section I also present an extension of the win probability metric to the evaluation of playoff series win probabilities.

In section 4 I introduce the Added Goal Value metric (AGV), which improves upon standard goal scoring statistics by accounting for the value each goal adds to the team's probability of winning the game. For those interested in player evaluation, the metric provides a way to identify which players' offensive production tends to occur in high stress situations. I demonstrate how AGV can be useful for a team comparing players with similar goal-scoring productivity.

In the final section I conclude with several possible extensions of the win probability and AGV metrics.

2 A win probability metric for NHL games

Although win probability graphs have become commonplace in other sports [6, 7], the NHL currently lacks such a metric.¹ The method described in this section is thus the only metric currently available for assessing second-by-second probabilities that a team wins an NHL game. The statistic incorporates information about the current score, penalty situation, and home-ice advantage. To my knowledge, this is the first hockey win probability metric that accounts for penalty time remaining at each second of the game. The play-by-play data available from the

¹ Briefly at the end of the 2013-2014 season, Extra Skater provided win probability graphs which were based exclusively on the score differential. That site became inactive when its owner was hired by an NHL team.

NHL's website does not explicitly include this information, so I developed an algorithm that calculates the time remaining based on unexpired penalty time, goals that may erase penalty time, and nuances in the rulebook that dictate how penalties are assessed in unusual situations like line brawls. The model underlying the metric has a framework that it allows for the future incorporation of other factors that influence goal scoring.

The model for calculating the probability that the home team wins the game when there are t seconds remaining is described by the formula below.

$$P_t(w) = P_t(w | \delta_t + 1) \cdot \Lambda(\boldsymbol{\gamma}_h \cdot \mathbf{v}_t) + P_t(w | \delta_t - 1) \cdot \Lambda(\boldsymbol{\gamma}_a \cdot \mathbf{v}_t) + P_t(w | \delta_t) (1 - \Lambda(\boldsymbol{\gamma}_h \cdot \mathbf{v}_t))(1 - \Lambda(\boldsymbol{\gamma}_a \cdot \mathbf{v}_t))$$

w : the event that the home team wins the game

δ_t : the score differential (home minus away) at time t

$\boldsymbol{\gamma}_h$: a vector of goal-scoring rates (goals per second) for the home team at 5-on-4, 5-on-3, 4-on-3, 3-on-4, 3-on-5, and 4-on-5

$\boldsymbol{\gamma}_a$: the same as $\boldsymbol{\gamma}_h$ except reversed in order so that the values correspond with away team rates

\mathbf{v}_t : a vector of seconds remaining in each of the six non-even strength situations at time t

$\Lambda(\cdot)$: the probability mass function of the Poisson distribution, evaluated for $P(x = 1)^2$

The simplest version of the model occurs when both teams are playing at full strength. The functions $\Lambda(\boldsymbol{\gamma}_h \cdot \mathbf{v}_t)$ and $\Lambda(\boldsymbol{\gamma}_a \cdot \mathbf{v}_t)$ signify the probability that home or away team, respectively, scores a goal before the teams return to full strength. These probabilities equal zero when the teams are playing at full strength because \mathbf{v}_t is a vector of zeros. Thus, the model simplifies to $P_t(w) = P_t(w | \delta_t)$ during 5-on-5 scenarios (or 4-on-4 scenarios in the case of regular season overtime periods).




Figure 1: Empirical probability that the team with the lead wins game, conditional on score difference, δ

I estimate $P_t(w | \delta_t)$ based on data extracted from the summaries of all regular season games from 05/06 to 12/13 ($n > 9000$) [8]. For each possible score differential and at each second of regulation time, I calculated the percent of time that the team with the lead ultimately won the game.³ When the score is tied ($\delta = 0$) I calculated the probability that the home team wins the game. Home teams win about 54% of games, but have almost exactly even odds of winning in overtime.⁴ I then applied a weak beta prior distribution to minimize the variance generated by

² I have simplified the model here to exclude the possibility that multiple powerplay or shorthanded goals occur during the course of a penalty. For minor penalties, a PPG will cause the penalty time to expire so this simplification is justified. Multiple powerplay goals during a major penalty, or multiple shorthanded goals are very rare so the simplification has very little impact on the result of the model.

³ I collapsed the small number of cases in which the score differential, δ , was 6 or greater into the $\delta = 5$ category.

⁴ Future versions of the metric will use a different function for when the score is tied, based on pre-game odds of winning. The downside of this approach is that it makes it virtually impossible to compare statistics derived from the metric, like AGV, across teams.

the tiny number of observations used to estimate the probabilities early in the game, particularly for the large goal differentials. These empirical estimates of $P_t(w | \delta_t)$ are shown in Figure 1.

In order to smooth out the curves in Figure 1, I regress the empirical win probabilities for each goal differential on a third order polynomial for the time remaining in the game. Using the coefficients from each of these regressions, I calculated the expected win probabilities at each time in regulation, conditional on the score. These values provide the inputs used for $P_t(w | \delta_t)$ in the model. Figure 2 shows these smoothed estimates.




Figure 2: Smoothed probability that the team with the lead wins game, conditional on score difference, δ

I use the same dataset to estimate the inputs for the γ vectors. I calculate the total amount of time played at each possible non-even strength situation (5-on-4, 5-on-3, 4-on-3), as well as the number of powerplay and shorthanded goals. Dividing the time by the number of goals provides estimates for the PPG and SHG rates at each penalty situation. These values, on the left of Table 1, make up the γ_h and γ_a vectors.⁵ I then assume that goals are generated from a Poisson distribution, and calculate the probability of a PPG and SHG during some amount of penalty time using the Poisson PMF [9]. When two minutes of 5-on-4 are played, the probability of a powerplay goal is about 17%. These probabilities match up well with the fact that the average powerplay percentage in the NHL is approximately 20%.

Table 1: Powerplay and shorthanded goal rates and probabilities

Situation	PPG per second	SHG per seconds	P(PPG in 2 minutes)	P(SHG in 2 minutes)
5-on-4	0.00172	0.000239	16.79%	2.79%
5-on-3	0.005245	0.000101	33.54%	1.20%
4-on-3	0.002499	0.000172	22.22%	2.02%

An advantage to this model of win probabilities is that it can unify existing and new knowledge or data together under a common framework. Future versions of the metric will incorporate the findings of research on the impact of zone starts or turnovers on goal-scoring rates. These findings could enter the win probability metric through the $\Lambda(\cdot)$ function. Similarly, if spatial data about the location of players and the puck becomes available, we could manipulate the rate parameter inside of $\Lambda(\cdot)$ to reflect how the probability of a goal being scored changes with the flow of the game.

3 Building a narrative from the win probability metric

⁵ Because a powerplay goal censors the data, these estimates will slightly underestimate the true goal-scoring rates during powerplays. This would be more of a problem if the rates were high, but because they are so low the degree of bias induced is likely to be tiny.

Just as journalists and bloggers use win probabilities to describe the impact of a homerun in baseball or a 4th down conversion in football [10], hockey observers can use this metric to empirically assess the importance of a goal or the costliness of a penalty taken. Figure 3 illustrates how the metric works in an especially back-and-forth game early in the 2014/2015 season. Although Philadelphia ultimately won the game in overtime, Dallas had the upper-hand for most of the third period. Just before the Flyers made the score 4-3 with 11:28 remaining in the period, Philadelphia had just a 5.5% chance of winning. Later in the period when the teams traded three goals in the span of two minutes, the Flyer's probability of winning swung from 10.0% to 48% back down to 9%, then back up to 48%. Both times Dallas took a penalty at the end of regulation and in overtime, the Flyer's chances of winning increased by 9 percentage points. These statistics could be invaluable to a reporter who wants to quantify the back-and-forth that occurred during a particularly exciting game.




Figure 3: In-game win probability for a recent game between Philadelphia and Dallas

Another application of the metric for writers is the assessment of the probability that a team will win a playoff series. To do this, I calculate the in-game win probabilities for each game and use them to estimate the series win probability based on the following formula:

$$P(\text{win series}) = P(\text{win series} \mid \text{win current game})P(\text{win current game}) + P(\text{win series} \mid \text{lose current game})(1 - P(\text{win current game}))$$

The conditional probabilities on the right hand side of the equation are estimated based on the results of every best-of-7 NHL playoff series since the league expanded to 12 teams. They reflect the probability that a team wins a series depending on the results from previous games and the number of home and away games remaining. This formula ensures that the in-game probabilities for each game match up at each of the "knots" between each game.

Figure 4 shows the results of the playoff series win probability metric for the 2014 Western Conference Final. Prior to scoring six unanswered goals in Game 2, the Kings had just a 14% chance of advancing to the Stanley Cup Final. Their fate shifted dramatically by Games 5 and 6, when their probability of winning the series spiked as high as 95% and 92%. The Kings' fortunes shifted again very quickly. With about 8 minutes left in Game 6, LA had a 92% chance of winning the series. Prior to tying the game with 7 minutes remaining in Game 7, their probability had eroded to just 12%.

The win probability metric has numerous other applications, such as a stats-driven power ranking, for those who write about hockey. I now turn to an example of how win probabilities can be useful for those who want to evaluate the skills of individual players.




Figure 4: Playoff series win probabilities for the 2014 Western Conference Final

4 Using win probabilities to calculate player Added Goal Value

Imagine a team wants to add to its depth at forward and has the choice between two equally priced players who average about 20 goals per season. Now imagine that one of those players scored all of his goals in overtime, while the other player scored all of his when his team already had a large lead. It is clear that the first player is a more valuable asset. But what about in situations that are not nearly as clear-cut? By using the change in win probability when a goal is scored, we can evaluate how much a player's goal contributions impact their team's chance of winning the game. I call this new statistic Added Goal Value (AGV).

Not all goals have the same impact on the probability that a team wins the game. The intuition behind AGV is that the more a player's offensive production comes during important moments during games, the more his team values him. The statistic is similar to Jeff Sagarin's Player Win Averages in baseball [11] and Win Probability Added metrics in baseball [12] and football [13]. These statistics take the change in a team's win probability that results from an at-bat or play and aggregate the changes over a game or season. AGV is similar in that it is measuring the impact that a goal has on the outcome of a game. Like PWA and WPA, AGV discounts events that occur when the outcome of a game is a foregone conclusion and upweights events during high pressure moments.

In this way, AGV is an improvement over traditional goals or points statistics which treat all goals as equal, regardless of the closeness of the game or time remaining. AGV is also an improvement over the game-winning-goals statistic. Players are awarded with a GWG if their goal was decisive in the final score. This means that when a team builds a 5-0 lead but wins the game 5-4, the fifth goal scorer receives the game winner even though his goal seemed inconsequential at the time he scored it. In this way AGV accounts for how players respond to the pressure of the moment.

AGV also nicely complements popular statistics like Corsi and Fenwick, which are a proxy for how well a player impacts puck possession when they are on the ice. While Corsi and Fenwick have their own merits, they do not correlate especially well with individual goals or points statistics [14, 15]. For a team that is looking to add a goal-scorer to their roster, AGV is a more useful tool.

$$AGV_i = \sum_{k=1}^K \Delta_k - \frac{\sum_{j=1}^J \Delta_j}{J} \cdot K$$

$$\Delta_k = P_{t_{k+1}}(w) - P_{t_k}(w)$$

$$\Delta_j = P_{t_{j+1}}(w) - P_{t_j}(w)$$

- i : player
- k : denotes a goal scored by player i
- K : the total number of goals scored by player i
- j : denotes a goal scored by a player who is not i
- J : the total number of goals scored by all players other than i
- w : the event that the goal scoring player's team wins the game
- t_k, t_j : time (in seconds) at which goal k or j was scored

The formula for calculating Added Goal Value is presented above. The first term on the right hand side calculates the total change in win probability that occurred from each of player i 's goals. The second term calculates the average change in win probability from goals by every other player in the NHL and then multiplies it by the total number of goals scored by player i .⁶ Since the 2004/2005 lockout, the average increase in probability from a goal is 17%; a team's probability of winning the game when they score a goal increases by 17% on average. Multiplying this by K yields the expected total probability increase for a player whose goals make average contributions to their team's odds of winning. The difference between these two numbers tells us how much total value a player's goals added to their team's probability of winning their games, compared to the rest of the league.




Figure 5: Relationship between AGV and goals per game for all forwards since 2004/2005 lockout

Figure 5 displays the relationship between AGV and goals per game for all NHL forwards who have played at least ten games and scored at least one goal since the 2004/2005 lockout. The horizontal line denotes the average AGV throughout the league. As the graph shows, the players with the highest AGV are Saku Koivu, Pavel Datsyuk, and Mikko Koivu. Alex Ovechkin and Ilya Kovalchuk are perhaps more valuable to their team because not only do their goals generally have a high impact on their teams' probability, they also score goals at exceptionally high rates.

⁶ Whether or not I include player i 's goals in the calculation of the league-average is inconsequential since no one player scores enough goals in relation to the total number of goals scored league-wide.

This graph also shows a limitation of the AGV metric. Any general manager would love to have Steven Stamkos, Sidney Crosby, or Evgeni Malkin on their roster, yet all three players have an AGV that is average. This highlights an important point about utilizing the metric for player evaluation. AGV is most useful in providing comparisons between players of similar rates of goal scoring per game. It does not matter that goals by Stamkos, Crosby, and Ovechkin only boost a team's win probability by 17% if they are going to score once every two games. Thus the AGV metric is most helpful in comparing players with similar goal scoring rates.



Figure 6: AGV for players who scored between 19 and 21 non-empty net goals in 2013-2014

Where the statistic can be incredibly useful is in evaluating the skill of non-elite players. Figure 6 displays the AGV for all players who scored between 19 and 21 non-empty net goals in 2013/2014. For a general manager looking to add scoring depth, these are the players to whom he may look. Based on AGV, Shea Weber and Alexander Semin scored goals which were more consequential in terms of win probability than Mats Zuccarello and Vladimir Tarasenko.




Figure 7: Between-year correlation of AGV for each pair of consecutive seasons from 05/06 to 13/14

A player's AGV also tends to be strongly correlated ($\rho=0.647$) across seasons. Figure 7 plots a player's AGV in one season against their AGV in the following season. The red line in the figure represents the results of

regressing y on x . As the graph indicates, player AGV this season is highly predictive of their AGV in the following season. The line nearly approximates a 1 to 1 relationship (β : 0.63499; se: 0.01145). This means that players who had a high AGV last season are likely to have a high AGV this season.

$$AGVG_i = \frac{\sum_{k=1}^K \Delta_k}{g_i}$$

g_i : number of games played by player i

The formula for AGV can also be slightly tweaked to provide estimates of the average increase in win probability provided by a player in each game they play. The formula for calculating Added Goal Value per Game takes into account both the total number of goals scored by a player, as well as the average win probability added by those goals. Dividing by the number of games played allows us to interpret the result as the average boost in win probability a team receives each game from a player's goals.

Table 2: Players with the best AGVG in 2013/2014

Player	AGV per game	Goals	Games played
Alex Ovechkin	13.68	51	78
Steven Stamkos	13.03	25	37
Jeff Skinner	10.28	33	71
Corey Perry	10.09	43	81
James Neal	9.54	27	59
Gustav Nyquist	9.15	28	57
Sidney Crosby	8.93	36	80
Phil Kessel	8.67	37	82
Max Pacioretty	8.61	39	73
Kyle Okposo	8.45	27	71
Joe Pavelski	8.34	41	82
Jeff Carter	8.2	27	72
Mike Cammalleri	8.06	26	63
Evgeni Malkin	7.85	23	60
Pavel Datsyuk	7.68	17	45

Instead of ranking players based on the total number of goals they scored in a season, Table 2 ranks them based on the average boost in win probability their goals provided their team each game. Alex Ovechkin, for example, increased Washington's chance of winning by about 14% based on the composite effect of the number of goals he scored and the average contribution his goals had on their win probability.⁷ Steven Stamkos, despite missing half the season with a broken leg, comes in second because of his 25 goals had a very large impact on the Lightning's chances of winning games.

5 Future Work and Conclusions

The paper has presented a new framework for assessing the flow of an NHL game and evaluating players using a new win probability metric. The method also provides several avenues for future research. In addition to incorporating other statistics and sources of data into the metric, there are numerous other ways it can be applied to the work of those writing about hockey and those evaluating talent. Just as the metric can be applied to playoff series, it could also be used to estimate second-by-second probabilities that a team makes the playoffs throughout the regular season.

The natural extension of AGV is to create a shootout statistic by treating attempts as similar to at-bats in analogous baseball methods. A metric like Average Goal Value can also be applied to goalies to evaluate how much their goals against and saves contribute to the game's outcome. Also as the metric becomes more rich in its use of spatial and other data sources, a metric could be developed which utilizes the changes in win probability that occur during a player's shift, rather than just the changes from the goals they score.

⁷ As a point of reference, the median AGVG for all forwards was 2.33.

6 References

- [1] Bradbury, J.C. "Putting A Dollar Sign on the Muscle: Valuing Players." Hot Stove Economics: Understanding Baseball's Second Season. New York: Springer, 2011. 69-97. Print.
- [2] Stern, Hal S. "A Brownian Motion Model for the Progress of Sports Scores." Journal of the American Statistical Association 89.427 (1994): 1128-134. Print.
- [3] Berri, David, and Brian Burke. "Measuring Productivity of NFL Players." The Economics of the National Football League. New York: Springer, 2012. Print.
- [4] O'Shaughnessy, Darren. "Possession Versus Position: Strategic Evaluation in AFL." Journal of Sports Science and Medicine 5.4 (2006): 533-40. Print.
- [5] Allen, et al. "StatsMonkey: A Data-Driven Sports Narrative Writer" AAAI Fall Symposium Series (2010). Web. 1 Dec. 2014
- [6] "Win Expectancy." FanGraphs Sabermetrics Library. Web. 1 Dec. 2014.
<<http://www.fangraphs.com/library/misc/we/>>.
- [7] "Win Prob Calculator." Advanced Football Analytics. Web. 1 Dec. 2014.
<<http://wp.advancedfootballanalytics.com/winprobc1.php>>.
- [8] A.C. Thomas and Samuel L. Ventura (2014). nhlsrapr: Compiling the NHL Real Time Scoring System Database for easy use in R. R package version 1.8. <<http://CRAN.R-project.org/package=nhlsrapr>>.
- [9] Andrew C. Thomas. "Inter-arrival Times of Goals in Ice Hockey." Journal of Quantitative Analysis in Sports 3.5 (2007). Print.
- [10] Vecer, et al. "On Probabilistic Excitement Of Sports Games." Journal of Quantitative Analysis in Sports 3.7 (2007). Print.
- [11] Sagarin, Jeff. "Player Win Averages for 1957-2006." 2007. Web. <<http://sagarin.com/mills/seasons.htm>>.
- [12] Studeman, Dave. "The One About Win Probability." 27 December 2004. Web.
<<http://www.hardballtimes.com/the-one-about-win-probability/>>.
- [13] Burke, Brian. "Win Probability Added (WPA) Explained." 27 January 2010. Web.
<<http://archive.advancedfootballanalytics.com/2010/01/win-probability-added-wpa-explained.html>>.
- [14] Johnson, David. "Goal Rates Better than Corsi/Fenwick in Player Evaluation." HockeyAnalysis. 30 May 2011. Web. <<http://hockeyanalysis.com/2011/05/30/goal-rates-better-than-corsifenwick-in-player-evaluation/>>.
- [15] Pettigrew, Stephen. "Why We Should be Trying to do Better than Corsi and Fenwick." Rink Stats. 28 May 2014. Web. <<http://www.rinkstats.com/2014/05/corsi-and-fenwick-suck-or-why-we-should/>>.