Expect the Expected: Approximating the Caliber of Possession Using Shot Quality

James McCorriston
jamie_m_12@hotmail.com

Connor Reed
connor.reed.92@gmail.com

Abstract

The NHL has experienced rapid growth in analytical metrics and advanced statistics in recent years. While popular statistics like Fenwick and Corsi act as good approximations for puck possession, they are limited in what they tell about scoring opportunities as they do not consider shot quality. In this study, we consider shot distance as an approximation of shot quality, and we combine Fenwick and NHL play-by-play shot distance data to develop a series of new statistics: Expected Goals (xGoals), Expected Differential (xDiff), and Goals-Above-Expected (GAE) for skaters, as well as Expected Save Percentage (xSv%) and Adjusted Save Percentage for goaltenders. As a basis for these new metrics, we first show that shot distance serves as a good approximation for shot quality, and that we can reverse-engineer scoring probabilities for each shot taken by a player. The concept of approximating shot quality is extended to analyze the performance of players, teams, and goaltenders. Using NHL play-by-play data from the 2007-08 season to the 2014-15 season, we show that xGoals are the best indicator of how many goals a player should be scoring, and we show that it stays more consistent for an individual from year-to-year than other comparable statistics. Finally, we show that on a single-game resolution, xGoals are the best indicator for which team should have won a particular game. The novel set of metrics introduced in this paper offer a more reliable and indicative tool for assessing the ability of skaters, goaltenders, and teams and provides a new basis for analyzing the game of professional hockey.

1 Introduction

In today's world of sports analytics, there is a delicate balance that must be maintained between the complexity of a new analytical method and its ability to be understood. When using a mathematical approach to evaluate players or teams, computational analysis is most effective if it can be quickly understood and utilized by key decision makers such as managers and coaches. In the NHL, puck possession approximations have become the go-to measurement used by the analytics community to assess the quality of a player or team. While puck possession has been linked to success over large samples ([1], [2], [3]), it is also clear that it does not tell the whole story when it comes to the performance of a player or team, or in determining the outcome of a game. Other more complex metrics including Total Hockey Rating ([4]) or even a previous expected goals model ([5]) have been explored but are difficult to translate into an actionable narrative due to their lack of a simple foundation. In this paper, we introduce a new series of metrics that combine a shot quality approximation with existing puck possession metrics. These new metrics provide a more accurate evaluation of player and team performance while creating easily understood narratives.

Previous research has demonstrated that possession metrics are linked to winning over large samples ([6]), and that shot quality paints a more detailed picture for converting on possessions ([7]). However, each of these approaches to analyzing the game come with substantial drawbacks.

Possession models are unintuitive and not clearly related to goal scoring. Shot quality models require heavy man-hours of charting, are subject to the charter's bias, and feature discrepancies between annotators ([8]). In an effort to draw from the objectivity of possession models and the intuition of shot quality models, we propose a method combining both types of models into one.

In this paper, we use NHL play-by-play data to study the relationship between shot distance and shot accuracy, and model this relationship as a function. We then use this function as a basis for a series of new metrics that approximate expected scoring rates for players and Expected Save Percentages for goaltenders. We proceed to show how these new statistics compare to other similar metrics in season-to-season robustness, in correlation with observed goal scoring, and in determining who should have won a game. Finally, we discuss the many types of narratives that can be drawn from this new series of metrics.

By combining shot quality and possession metrics, we are able to leverage the complexity of two dimensions (quality approximation and frequency) while maintaining a simple narrative with one-dimensional values. To communicate the results of our model, we created a set of metrics that provide continuous representations of expected player goal scoring, expected goal percentage, expected team goal percentage, player shooting ability, and team-independent save percentages for goalies. These metrics provide new ways to evaluate players, goaltenders and teams.

2 Dataset

All player and team data was manually collected from the NHL's website. Specifically, shots, goals, shot attempts, shot distances, as well as which players were on the ice for each event were gathered from the play-by-play data available on the website. As the NHL started recording shot attempts in the 2007-08 season, our data spans the beginning of the 2007-08 season to the end of the 2014-15 season and only includes regular season game data. Player time on ice and team powerplay and penalty kill times were also taken from the NHL's website. *All data used is available for download on Github, but is not referenced here for anonymity purposes during the review process.*

3 Methods

3.1 The Distance-to-Accuracy Relationship

The basis for the statistics developed in this study is grounded in the relationship between shot distance and shot accuracy. While the general concept that shots closer to the net are more likely to become goals may be intuitive, the actual relationship has not been thoroughly explored ([9]). In fact, with commonly used advanced hockey statistics, shot distance is either completely ignored, or taken into consideration only for shots on goal rather than shot attempts ([10]). Investigating this relationship, we found that there is valuable information in the distance from which a shot is taken. Looking at all unblocked shot attempts from the 2007-08 to 2014-15 play-by-play data, we found that unsurprisingly, as shot distance decreases, shot accuracy increases. However, we also found that this relationship is non-linear and that situational strength (e.g. 5v5, 5v4) is a factor that determines shooting accuracy ([11]). Given that the distance data was a sample of the true distance-accuracy function and that certain distances were less frequently represented than others in the sample, we fit the observed results to an exponential function using weighted exponential regression ([12]).[1] This exponential function was then used for all distance to accuracy mappings done in this study.

---

[1] Our data is considered to be heteroscedastic so weighted regression is preferred to unweighted regression.

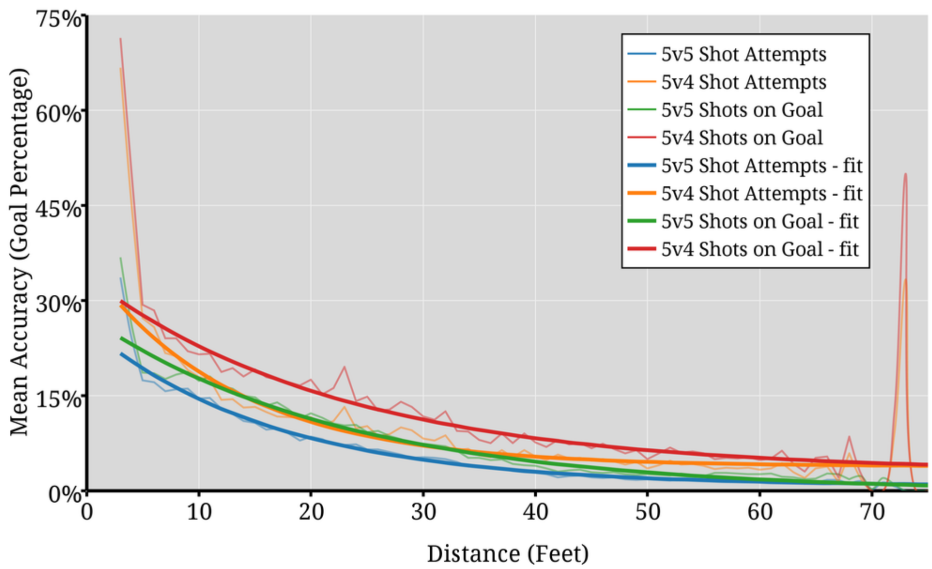

Figure 1: The relationship between shot distance and mean shot accuracy. Thicker smooth lines show the exponential functions to which the sampled data (faded colors) were fit.

In Figure 1, the weighted exponential functions are shown along with the sample data used to fit them. Shot accuracies range between ~0-30%, and it is observed that 5v4 situations offer elevated shooting percentages at all distances.[2] We hypothesize that this difference stems from the fact that players have more time and space with fewer defenders on the ice, and can make more calculated, unobstructed shots. The range in accuracy demonstrated in this relationship stresses the importance of where a shot is taken from. The models presented in this paper are founded on this relationship and dictated towards unblocked shot attempts taken in 5v5 or 5v4 situations.[3]

3.2 Expected Goals (xGoals)

The foundation of this study and the signature metric developed in this work is the expected goal, or xGoal. Rather than looking at discrete events such as goals or shot attempts, the xGoal was developed to more accurately reflect goal scoring at the 5v5 and 5v4 game state. The xGoal represents a shot-quality-weighted Fenwick that uses shot distance as an approximation of shot quality. By using game state (e.g. 5v5, 5v4) to determine the appropriate distance-accuracy function from Section 3.1, we can pass the distance of a shot as input, and retrieve a shot quality approximation in the form of an expected shooting percentage for that shot. This function is used to weight each shot attempt, giving greater value to shots taken from distances with historically better conversion percentages. For each unblocked shot attempt, the xGoal is calculated with the following equation:

$$xGoal \;=\; P(shot\_distance \mid game\ state) \tag{1}$$

---

[2] The outliers have very little influence on the fit function due to their low frequency in the sample. See Figure 4 in the appendix for the full frequency distribution.

[3] Shots and goals scored at 5v3, 4v5, 4v4, 4v3, 3v5, 3v4, 3v3, any situation with a pulled goalie, and penalty shots have been omitted, as we do not have adequate data to determine the distance-accuracy relationship

To find the total amount of goals we should expect a player or team to score over a sample of multiple shots, we can simply sum the individual probabilities as such:

$$xGoals \ = \sum_{i=1}^{m} P(shot\_distance_i \mid 5v5) \ + \sum_{j=1}^{n} P(shot\_distance_j \mid 5v4) \tag{2}$$

As an example, a shot attempt taken from 5 feet away from the goal at 5v5 is expected to result in a goal 19.3% of the time. When a player takes a shot from this distance, he is credited with 0.193 xGoals. A shot taken from 35 feet away from the goal results in a goal 3.8% of the time when 5v5 and 6.1% of the time when on a 5v4 powerplay. When a player takes a shot from this distance, he is credited with 0.038 or 0.061 xGoals, depending on the game state. If one player takes unblocked shot attempts from 5ft and 35ft at 5v5 and another from 35ft at 5v4, their total xGoals is calculated as follows:

$$\begin{aligned} xGoals &= P(5ft \mid 5v5) \ + \ P(35ft \mid 5v5) \ + \ P(35ft \mid 5v4) \\ &= 0.193 \ + 0.038 \ + \ 0.061 \\ &= 0.292 \end{aligned} \tag{3}$$

Over this game, we can now say that this player recorded a discrete sum of 3 unblocked shot attempts and was expected to score 0.292 goals.

While goal scoring is discrete and extremely stochastic in nature, xGoals paint a smoother picture, taking away some of the high variability that is associated with goal scoring, while recovering the value of better scoring opportunities that is lost by only looking at shot attempts. The xGoal forms the basis from which all other metrics in this paper are extended.

3.3 Expected Goal Difference (xDiff)

To assess a player's overall effectiveness we analyzed both their offensive and defensive performances. By applying the same methodology of xGoals to shot attempts both for and against a player's team when they are on the ice, we approximate both the offensive and defensive ability of this player in one statistic, xDiff. On a player level, xDiff is calculated as the sum of all xGoals recorded for a player and his teammates while the player is on the ice divided by the sum of all xGoals both for and against the player's team while he is on the ice:

$$xDiff_{(Team\ X, Player\ 1)} \ = \frac{\sum_{i=1}^{m}(xGoal\ of\ shot_i\ by\ Team\ X \mid Player\ 1\ on\ Ice)}{\sum_{j=1}^{n}(xGoal\ of\ shot_j\ by\ either\ team \mid Player\ 1\ on\ Ice)} \tag{4}$$

xDiff is recorded as a fractional metric, displaying the ratio (or percentage) of goals we can expect a player's team to score while he is on the ice. This is analogous to the shot-attempts-for percent commonly used to approximate puck possession.

This idea can be extended to the team level, where the only difference is that the condition of a certain player being on the ice is no longer necessary, and instead xDiff is calculated as the ratio of all xGoals for and against a team:

$$xDiff_{(Team\ X)} \ = \frac{\sum_{i=1}^{m}(xGoal\ of\ shot_i\ by\ Team\ X)}{\sum_{j=1}^{n}(xGoal\ of\ shot_j\ by\ either\ Team\ X\ or\ Team\ Y)} \tag{5}$$

Both player and team xDiffs feature a similar formula, albeit with one notable difference. The player version only includes 5v5 shot attempts so we do not unfairly penalize players who appear on the penalty kill, nor unfairly reward players who appear on the power play. Team-based xDiff includes shot attempts from our entire dataset (5v5 and 5v4), and thus penalizes teams who take a surplus of penalties and rewards teams who draw more penalties than they take.

3.4 Goals Above Expected (GAE) & Goals Per Expected (GPE)

Due to a variety of factors (including, but not limited to, issues with sample size, stochasticity, the discrete incrementation of goals, and a player's individual shooting ability), the number of observed goals can often waver from the expected goals. To better analyze these deviations, we have created a pair of metrics, Goals Above Expected (GAE) and Goals Per Expected (GPE). GAE and GPE include xGoals at both 5v5 and 5v4 game state, which is more inclusive than currently used "luck"-measuring practices such as PDO ([13]). GAE is measured simply by subtracting the expected goals from the observed goals:

$$GAE = Goals\_observed - xGoals \tag{6}$$

Since the magnitude of GAE is primarily a function of shot attempt volume (and thus a function of time on ice), we present GPE as the rate equivalent of GAE, measured by dividing observed goals by xGoals:

$$GPE = \frac{Goals\_observed}{xGoals} \tag{7}$$

For example, Alex Ovechkin amassed 43 goals at the 5v5 and 5v4 game state in 2013/14. His 501 shot attempts resulted in 28.385 xGoals. Therefore, his GAE value is measured as (43 - 28.385 = +14.615) while his GPE is measured as (43 / 28.699 = 1.515). GAE and GPE present the same idea in two different fashions, as GAE measures total surplus value while GPE measures rate of surplus value.

3.5 Expected Save Percentage & Adjusted Save Percentage (for usage with Goaltenders)

In hockey, skaters get most of the attention when it comes to advanced statistics. However, in this study we propose a new metric for goaltenders called Adjusted Save Percentage which takes into account the quality of shots a goaltender faces. In order to calculate this metric, we first calculate another statistic called the Expected Save Percentage (xSv%). The xSv% of a goalie is calculated in a similar manner to the expected goals of a skater. The logic follows the same as with expected goals where shot distances are used to determine expected goal values for shots, but in this case, we are interested only in the shots on goal against a specific goalie. It should be noted that the xGoals calculation for this metric is made using the shots on goal probability function due to the fact that goaltenders cannot save shots that miss the net. Expected Save Percentage is calculated by the following equation:

$$xSv\% = 1 - \frac{xGoals\ against}{number\ of\ shots\ against} \tag{8}$$

The xSv% of a goalie actually tells us about the team in front of him, and is unrelated to the ability of the goalie himself. This metric can be thought of as an approximation of the quality of shots that a goalie's team allows against him.

Using the Expected Save Percentage of a goalie and the league average Expected Save Percentage, we can normalize our data and adjust the actual save percentage of a goalie to take into account the quality of shots that he faces. This is calculated using the following equation:

$$Adjusted\ Save\ \% \ = \ Observed\ Sv\% \ * \ \left(\frac{League\ Avg.xSv\%}{xSv\%}\right) \tag{9}$$

As an example, from 2007-08 to 2014-15 Craig Anderson faced 5,843 shots in road games at either 5v5 or 5v4 split between games played for the Florida Panthers, Colorado Avalanche, and Ottawa Senators. Of these shots, he allowed 418 goals against (a raw 0.928 Sv%), but after weighting those shots by their distance-related shooting probability, we would expect 484.62 goals to be scored against a league average goalie (an xSv% of 0.917). Given that the league average xSv% is 0.919, we can calculate Anderson's Adjusted Sv% as follows:

$$Adjusted\ Save\ \% = \ 0.928 * \left(\frac{0.919}{0.917}\right) \ = \ 0.930$$

The idea behind this adjustment is that it provides a boost to goaltenders who face more difficult shots on average and penalizes goaltenders who face easier shots (a large factor in small sample swings of goaltender performance [14]), giving them a team-independent evaluation.

4 Results

4.1 Comparing xGoals to Other Metrics

To validate our hypothesis that we are gaining information by adding weight to Fenwick events, we performed comparisons between xGoals and several commonly used statistics. First, we looked at the relationship between xGoals and actual goal scoring. Second, we looked at the association between a team having higher xGoals and winning games. Finally, we analyzed the year-to-year sustainability of an individual's xGoals.

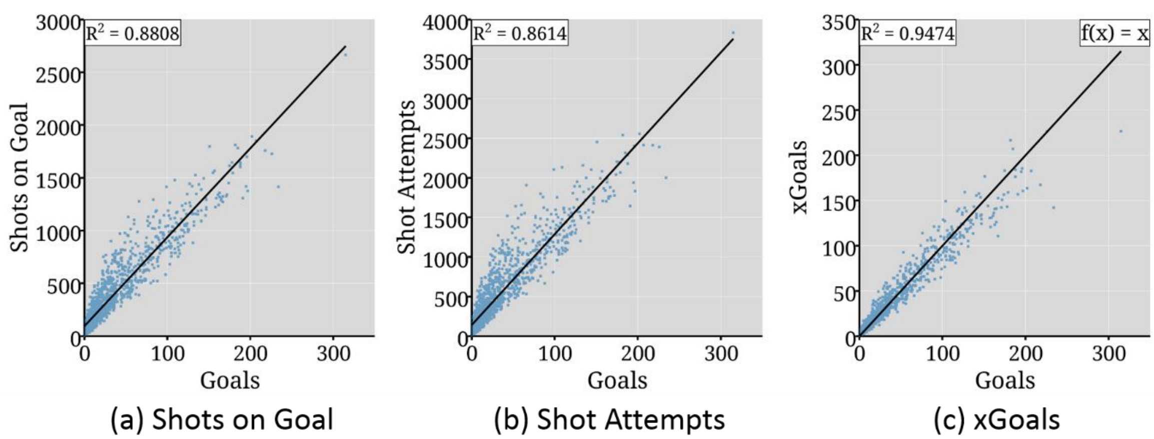

Figure 2: Correlations between observed goal scoring and each of (a) shots on goal, (b) shot attempts, and (c) xGoals. We observe that xGoals have the strongest linear correlation with actual goal scoring.

First, we studied the relationship between xGoals and actual (observed) goal scoring on an individual player basis. In Figure 2, we see the correlation between goal scoring, and each of shots on goal, shot attempts, and xGoals. Note that in an effort to normalize the data, rates per 60 minutes were used. From these analyses, we found that xGoals and observed goals were linearly correlated with an R^2 value of 0.95, while shots on goal and goals were linearly correlated with an R^2 of 0.88, and shot attempts and goals were correlated with an R^2 of 0.86.

		Winning Percentages		
Season	$n^{\dagger}$	***xGoals***	***Shot Attempts***	***Shots on Goal***
2014-2015	1054	56.0%	47.7%	51.0%
2013-2014	1051	54.1%	47.3%	48.7%
2012-2013	616	54.9%	48.1%	51.5%
2011-2012	1033	55.5%	46.7%	49.2%
2010-2011	1066	49.5%	43.2%	45.3%
2009-2010	1039	53.6%	45.0%	47.0%
2008-2009	1063	57.5%	50.1%	52.7%
2007-2008	1056	55.6%	47.4%	51.2%
Totals	$7978^{\ddagger}$	54.4%	46.9%	49.5%

† $\boldsymbol{n}$ is the total games from a given season, excluding those which are decided by shootout

‡ an average of 997.25 games per season

Table 1: Winning percentage of teams with higher shot metrics in a given game. Teams who produced more xGoals than their opponent in a given game won 54.4% of the time over an 8 year sample, while teams who produced more shot attempts won 46.9% of their games, and teams who produced more shots on goal won 49.5% of their games.

To evaluate xGoals as an indicator of team performance, we looked at the winning percentage of teams with higher xGoals than their opponents on a game-by-game level, and we compared these results to the winning percentages of teams with higher shot attempts, or higher shots on goal (Table 1). We found that of all games that were resolved before a shootout, the team with the higher xGoals won 54.4% of the time, whereas teams with more shots on goal and teams with more shot attempts only won 49.5% and 46.9% of the time, respectively. This demonstrates that xGoals is a better indicator of who "should have won" a game. While it has become relatively commonplace for team performance to be evaluated using shot attempts or shots on goal, this evidence shows that xGoals are a better metric on a single game resolution. This is likely an artifact of the strong linear relationship between xGoals and observed goals.

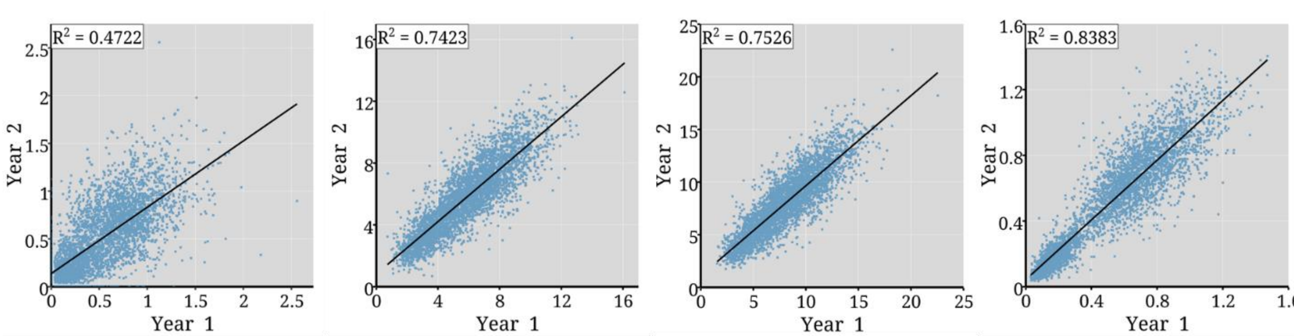

(a) Goals (b) Shots on Goal (c) Shot Attempts (d) xGoals

Figure 3: Correlation of (a) goals, (b) shots on goal, (c) unblocked shot attempts, and (d) xGoals in consecutive seasons. Each statistic is normalized to unit rate per 60 minutes played. xGoals have the highest R^2 value and remain the most consistent from year to year.

In an effort to examine the sustainability of different statistics across multiple seasons, we compared each individual player's actual goal scoring, shots on goal, shot attempts, and xGoals from one year to the next. In Figure 3, the data was normalized by looking at rates per 60 minutes of play. From this experiment, two observations are made: goal scoring varies quite drastically from year to year, while xGoals stay the most consistent. Goal scoring is known to be stochastic ([15]), so this is an expected result. The year-to-year xGoals linear correlation R^2 value of 0.84 suggests that xGoals per 60 minutes of play stays more consistent from year to year than any of shots on goal, shot attempts, or goal scoring of a given player. This shows that xGoals are the most reliable measure of a player's offensive production.

4.2 Constructing Narratives

One of the strengths of our new series of metrics is that it offers the opportunity to build several new types of narratives. These metrics can be used to generate narratives over any time frame, from a single game to an entire career.

The most unique individual player narratives that can be drawn from these new metrics come from the Goals Above Expected (GAE) and Goals Per Expected (GPE) statistics. For the purpose of this analysis, GAE can be thought of as a measure of volume, and GPE a measure of efficiency. Depending on the size of the sample in question, GAE and GPE together offer two new narratives: one describing the shooting ability of a player, and the other describing the amount of "luck" that a player might be experiencing at their current goal scoring pace.

Rank	*Player Name*	*GAE/60 min.*†	*Rank*	*Player Name*	*GAE/60 min.*†
			⋮	⋮	⋮
1	Steven Stamkos	+0.583	395	Justin Abdelkader	-0.227
2	Alex Ovechkin	+0.420	396	Nate Thompson	-0.227
3	Ilya Kovalchuk	+0.357	397	Sean Bergenheim	-0.238
4	Nathan Horton	+0.291	398	David Moss	-0.254
5	Alexander Semin	+0.269	399	Travis Moen	-0.258
6	Tyler Seguin	+0.262	400	Chris Neil	-0.265
7	Brad Marchand	+0.261	401	Scott Gomez	-0.279
8	Mike Cammalleri	+0.251	402	Boyd Gordon	-0.294
9	Jarome Iginla	+0.249	403	Ryan Smyth	-0.328
10	Alex Tanguay	+0.240	404	Jason Blake	-0.391
⋮	⋮	⋮			

† Minimum 5000 minutes played over this timespan.

Table 2: Top and bottom 10 players ranked by Goals Above Expected per 60 minutes (GAE/60 min) from the 07/08 season through the 14/15 season. Highly positive GAE/60 indicates strong shooting ability, while negative values indicate weak shooters.

When looking at a large sample, a positive or negative deviation from a GAE of 0 can be an indicator of a player having a superior or inferior shooting ability, respectively. In Table 2, the top 10 and bottom 10 players ranked by GAE/60 minutes played are listed. Both lists are populated with familiar names, but the top 10 boasts a list of players including Steven Stamkos, Alex Ovechkin, and Ilya Kovalchuk, typically known for their shooting ability, while the bottom 10 include names more closely associated with the "grinder" role.

When looking at a small sample, if a player's GPE is close to 1 (GAE close to 0), this indicates that they are scoring at a sustainable rate. If their GPE is much greater than 1, they are likely getting lucky and can be expected to regress, while if their GPE is much lower than 1, they are probably having bad luck and can be expected to score more. As an example, in 2007-08, Brad Boyes finished the season with 40 goals scored at 5v5 and 5v4, but only tallied 17.9 xGoals, totaling a GPE of 2.23 and GAE of +23.1. Without a bigger sample, or unless we believe Boyes to be a historically great shooter, he would not have been expected to maintain this observed goal scoring pace. As expected, eight seasons later, Boyes has regressed to a career GPE of 1.10.

Rank	*Player Name*	*xDiff*†	*Player Name*	*FF%*
1	Pavel Datsyuk	58.0%	Pavel Datsyuk	58.3%
2	Justin Williams	57.8%	Justin Williams	58.0%
3	Joe Pavelski	57.1%	Jonathan Toews	57.1%
4	Zach Parise	57.0%	Brian Rafalski	57.1%
5	Patrice Bergeron	56.8%	Patrice Bergeron	57.0%
6	Jonathan Toews	56.7%	Brad Marchand	56.3%
7	Brian Rafalski	56.5%	Johan Franzén	56.3%
8	Jaromir Jagr	56.5%	Henrik Zetterberg	56.1%
9	Henrik Zetterberg	56.4%	Joe Pavelski	56.1%
10	Sidney Crosby	56.2%	Nicklas Lidstrom	56.1%

† Minimum 5000 minutes played over this timespan.

Table 3: Top 10 players ranked by Expected Goal Differential Rate (xDiff) and unadjusted Fenwick-for percentage (FF%) from the 07/08 season through the 14/15 season. High xDiff is representative of strong two-way play.

Similar to possession metrics such as Fenwick-For Percentage (FF%), over a large sample xDiff can be used to measure the two-way ability of a player. Given that observed goal scoring is more closely associated with xGoal scoring than with shot attempts, xDiff provides a more comprehensive analysis of a skater's overall performance. Comparing the top 10 players in each column of Table 3, we note that the top players ranked by xDiff and FF% are similar but a few players do differ (most notably, Zach Parise and Sidney Crosby in for Brad Marchand and Johan Franzen). While the numbers between the complete lists are also quite similar, some players vary by 2-3%, which, over large samples, is significant. This suggests that using possession metrics alone, some players may be receiving unfair or inaccurate evaluations.

Season	*Team*	*xDiff*	*FF%*	*5v4 xGoals For*	*4v5 xGoals Against*
2014/15	NY Rangers	51.0%	49.0%	41.02	36.00

Table 4: FF%, xDiff, and specials teams xGoals of the New York Rangers in the 14/15 season. A net positive special teams xGoals, leads to xDiff being higher than FF%.

The application of xDiff on a team level is analogous to its application at a player level. Similar to possession metrics, xDiff condenses both the offensive and defensive play of a team into a single value, but also includes powerplay and penalty kill components giving a more complete analysis of a team's ability. Table 4 shows the FF% and xDiff of the New York Rangers in 2014-15, as well as their powerplay xGoals for and their penalty kill xGoals against. Their net special teams xGoals were positive and hence their xDiff was greater than their FF%. The inclusion of power play and penalty kill allows for a more organic evaluation of a team's overall performance.

Rank	*Goalie Name*	*Expected Save %*†	*Observed Save %*	*Adj. Save %*
1	Cory Schneider	0.912	0.927	0.934
2	Craig Anderson	0.917	0.929	0.931
3	Devan Dubnyk	0.914	0.925	0.930
4	Corey Crawford	0.917	0.927	0.929
5	Sergei Bobrovsky	0.917	0.927	0.929
6	Tomáš Vokoun	0.917	0.925	0.927
7	Jimmy Howard	0.916	0.924	0.927
8	Tim Thomas	0.919	0.928	0.927
9	Kari Lehtonen	0.918	0.925	0.926
10	Jonas Hiller	0.915	0.922	0.926

† Expected Save Percentage is the percentage of shots we expected an average goalie to save based on the quality of shots faced.

* Only road data was included to account for home scorer bias. Minimum 3000 road shots faced.

Table 5: Top 10 goalies ranked by Adjusted Save Percentage from the 07/08 season through the 14/15 season. By facing more difficult shots, goaltenders with a low Expected Save % see increases in their Adjusted Save % over their Observed Save %. The opposite is true for goaltenders facing easier shots (i.e. high Expected Save %).

For goaltenders, the Adjusted Save Percentage statistic allows for a team-independent evaluation of performance. Table 5 shows the list of top performing goalies ranked by Adjusted Save Percentage, using only road game data to correct for scorekeeper's bias. On this list, we note a set of names not typically as flashy as may be expected. The likes of Tim Thomas and Henrik Lundqvist take a back seat to Cory Schneider, Craig Anderson, and Devan Dubnyk at the top of the list.[4] While one might assume that Dubnyk's numbers playing with Minnesota in the 2014-15 season affect his ranking on this list, it should be noted that his Adjusted Save Percentage in 2014-15 was 0.931, virtually the same as his career average of 0.930. Despite the fact that his raw save percentage leaped from a career average of 0.919 prior to 2014-15, to 0.934 last season, Dubnyk's Adjusted Save Percentage suggests that the numbers he put up last season are sustainable. To our knowledge, this is the first metric that evaluates a goaltender's team-independent puck stopping ability and offers a unique opportunity to correctly value underrated goalies.

4.3 Further Results

For an interactive and complete set of results from this study (both reported and unreported), visit http://www.fourthlineheroes.com.

[4] Carey Price finished with the highest Adjusted Save Percentage in 2014-15. However, his performance in previous seasons has him ranked outside the top 10 from 2007-08 to 2014-15.

5 Conclusions and Future Work

This paper has presented an entirely new set of statistics based on expected goals (xGoals), a metric derived from distance-weighted unblocked shot attempts. We showed that xGoals are more closely associated with observed goal scoring and winning games than standard shot or possession metrics. Similarly, a player's xGoals stay more consistent year to year than any of their observed goals, shots on goal or unblocked shot attempts. We also introduced xDiff, Goals Above Expected (GAE), Goals Per Expected (GPE), and Adjusted Save Percentage, each of which are based on the concept of xGoals. The xDiff statistic acts as a good evaluation of a player or team's combined offensive and defensive performance. GAE and GPE are good measures of either goal-scoring "luck" or shooting ability, depending on the sample size. Finally, Adjusted Save Percentage corrects for the difficulty of shots that a goaltender faces, acting as a team-independent measure of their ability.

One of the major contributions offered by the metrics presented in this study is the generation of new narratives that provide unique evaluations of players, teams, and goaltenders that are relatively easy to understand. By compressing shot distance and shot attempt frequency into a single statistic, and by calculating "expected" performances of players, we are able to add complexity to the widely accepted possession metrics without sacrificing the simplicity of the narrative. This makes these metrics more applicable as the analyses they provide can be grasped by those in management and coaching positions.

The concepts behind the metrics introduced in this paper can easily be extended as more complex data becomes available. We showed that adding shot distance to shot attempts clearly adds information to the evaluation of a player, and we believe that two-dimensional shot location, as well as shooting conditions (e.g. rush, one-timer, or stationary for players; rush, screen, or rebound for goaltenders) can add even more. As this data becomes available ([16]), only the probability weight function will have to change; the theory behind xGoals and its derivatives will be preserved.

6 References

1. Johnson, D. *Intro to Advanced Hockey Statistics / Hockey Analytics*. August 5, 2014; Available from: http://www.mapleleafshotstove.com/2014/08/05/introduction-to-advanced-statistics/.
2. Tulsky, E. *Understanding Advanced Stats*. September 30, 2013; Available from: http://www.sbnation.com/nhl/2013/9/30/4784288/nhl-season-preview-2013-14#38-.
3. (JLikens), T.P. *Loose Ends - Part I: Predicting Future Success*. March 17, 2011; Available from: http://www.objectivenhl.blogspot.com/2011/03/loose-ends-part-i-predictive-validity.html.
4. Shuckers, M. and J. Curro, *Total Hockey Rating (THoR): A comprehensive statistical rating of National Hockey League forwards and defensemen based upon all on-ice events.* MIT Sloan Sports Analytics Conference. **2013**: p. 10.
5. Macdonald, B., *An Expected Goals Model for Evaluating NHL Teams and Players.* MIT Sloan Sports Analytics Conference, 2012: p. 8.
6. Boyle, C. *Why possesion matters: A visual guide to Fenwick*. April 4, 2013; Available from: http://www.habseyesontheprize.com/2013/4/4/4178716/why-possession-matters-a-visual-guide-to-fenwick.
7. Desjardins, G. *A Tale of Two Talents: One True and One Not*. May 24, 2010; Available from: http://www.arcticicehockey.com/2010/5/24/1482659/a-tale-of-two-talents-one-true-and?_ga=1.230070619.611281858.1447141315.
8. Zona, D. *Join The Scoring Chance Project*. September 6, 2011; Available from: http://www.coppernblue.com/2011/9/6/2407259/join-the-scoring-chance-project.
9. Tulsky, E. *How to bring shot quality into player evaluation*. December 12, 2013; Available from: http://www.sbnation.com/nhl/2013/12/12/5114366/nhl-stats-shot-quality-player-evaluation.
10. (DTMAboutHeart). *Expected Goals are a better predictor of future scoring than Corsi, Goals*. October 1, 2015; Available from: http://www.hockey-graphs.com/2015/10/01/expected-goals-are-a-better-predictor-of-future-scoring-than-corsi-goals/.
11. (SnarkSD), P.D. *Special Teams II: What Drives Power Play and Penalty Kill Success?* July 14, 2013; Available from: http://www.fearthefin.com/2013/7/14/4520832/special-teams-part-ii-what-drives-special-teams-success.
12. NIST/SEMATECH. *NIST/SEMATECH e-Handbook of Statistical Methods - Section 4.1.4.3. [Weighted Least Squares Regression]*. 2003 2013; Available from: http://www.itl.nist.gov/div898/handbook/pmd/section1/pmd143.htm.
13. (TheNeutral). *Non-Traditional Metrics Glossary: PDO*. August 2, 2012; Available from: http://www.fearthefin.com/2012/8/2/3215351/non-traditional-metrics-glossary-pdo.
14. Desjardins, G. *Fooled by Randomness: Goaltender Save Percentage*. April 20, 2010; Available from: http://www.arcticicehockey.com/2010/4/20/1429278/fooled-by-randomness-goaltender.
15. Tulsky, E. *Factoring regression into analysis*. June 16, 2013; Available from: http://www.broadstreethockey.com/2013/6/16/4430864/shooting-percentage-regression-goals.
16. Buntin, C. and M. Javan. *SportLogIQ*. 2015; Available from: http://www.sportlogiq.com.

7 Appendix

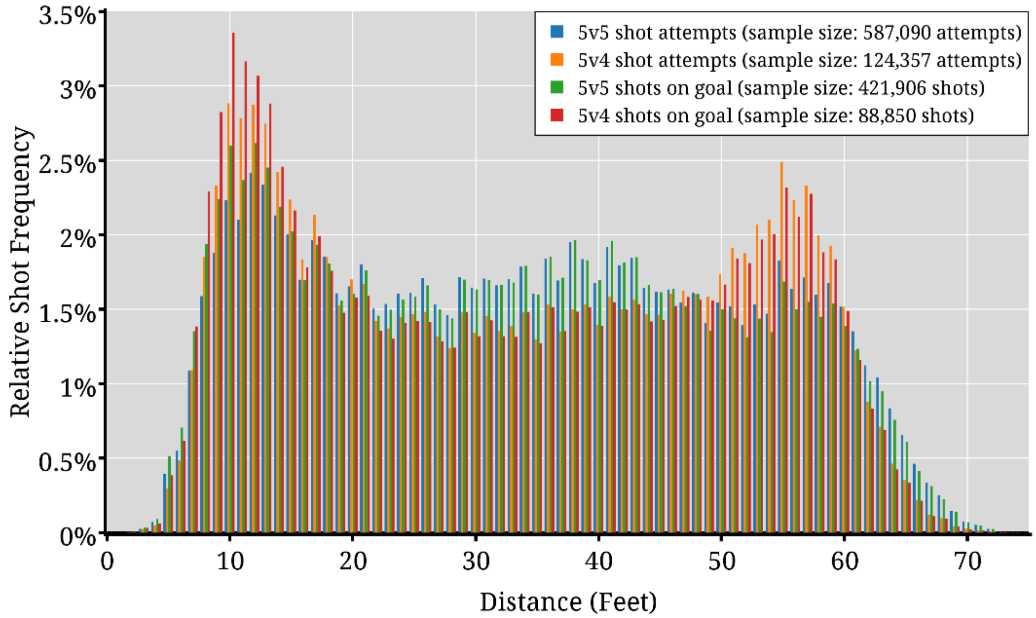

Figure 4: The relative shot frequency of all shot events considered in this study.

In Figure 2, the relative shot frequencies used as weights for each fit function are displayed. We note that the majority of shots are recorded between 5 and 70 feet, and that shots taken in the 5v5 game state have more of a bimodal distribution than shots in the 5v4 game state.