SHOT METRICS:
AN INTRODUCTORY GUIDE TO ~~ADVANCED~~ HOCKEY STATS

by

Mike McLaughlin
Patrick McLaughlin

July 2013

ABSTRACT

A primary goal of analysis in hockey and fantasy hockey is the ability to use statistics to accurately project the future performance of individual players and teams. Traditional hockey statistics (goals, assists, +/-, etc.) are limited in their ability to achieve this goal, due in large part to their non-repeatability.

One alternative approach to hockey analysis would use puck possession as its fundamental metric. That is, if a player or team is dominant, that dominance should be reflected in the amount of time in which they possess the puck. Unfortunately, the NHL does not track nor publish data related to puck possession. In spite of this lack of data, there are methods that can be used to track puck possession.

The purpose of this document is to introduce hockey fans (and fantasy hockey managers) to the topic of *Shot Metrics*. Briefly, *Shot Metrics* involves the use of NHL shot data to analyze individual players and teams. The shot data is used as a proxy for puck possession. Essentially, teams that are able to shoot the puck more often are doing so because they are more frequently in possession of the puck. It turns out that teams that are able to consistently outshoot their opponents typically end up winning games and performing well in the playoffs [1]. Thus, shot data can play an integral role in the way the game of hockey is analyzed.

Editor's Note: The type of analysis outlined in this manual is sometimes referred to as advanced statistics. Readers should be aware that the math involved in *Shot Metrics* is limited to basic addition, subtraction, multiplication, and division. The term advanced statistics is misleading (since there is **very little** math involved) and won't be used in this document.

CONTENTS

ABSTRACT . . . .

LIST OF FIGURES . . . . **iii**

LIST OF TABLES . . . . **iv**

NOTATION AND SYMBOLS . . . . **v**

CHAPTERS

1. INTRODUCTION **1**

1.1 A Simple Flip of a Coin 1
1.1.1 10 Coin Flips 1
1.1.2 100 Coin Flips 2
1.1.3 1000 Coin Flips 2
1.1.4 Summary 3
1.2 NHL Players as Coins 3
1.2.1 Phil Kessel 4
1.3 Thoughts on Sample Sizes 7
1.4 The Motivation for Analyzing Shot Data 8

2. THE SHOT METRICS **9**

2.1 Introduction 9
2.2 Notation and Terminology 9
2.2.1 Simple Shot Counting 9
2.2.2 Shot Differential 10
2.2.3 Adjusted Shot Differential 11
2.2.4 Fluctuations from League Average Performance 12
2.3 Assumptions of the Model 14
2.3.1 EV, PP, and PK 14
2.3.2 Score Effects 14

3. APPLIED SHOT METRICS **16**

3.1 Introduction 16
3.2 At the Team Level 16
3.2.1 Paper Tigers 16
3.3 At the Player Level 17
3.4 Summary 18
3.4.1 Final Thoughts 18
3.4.2 Shot Metrics Web Tool 18

APPENDICES

A. MAPPING SHOT METRICS TO ADVANCED STATS **19**

B. REGRESSION TOWARD THE MEAN **20**

REFERENCES **22**

LIST OF FIGURES

1.1 Results of a computer simulation of 10 coin flips run one million times. . . . 2
1.2 Results of a computer simulation of 100 coin flips run one million times. . . 3
1.3 Results of a computer simulation of 1000 coin flips run one million times. . 4
1.4 Results of a computer simulation of 60 SOG run one million times. 5
1.5 Results of a computer simulation of 275 SOG run one million times. 6
1.6 Results of a computer simulation of 1693 SOG run one million times. 7

LIST OF TABLES

2.1 League Average Data for Save Percentage and Shooting Percentage 12
3.1 Regression of FLAP Values: 2011-2012 Season to 2012-2013 Season 17

NOTATION AND SYMBOLS

G	Goals
SOG	Shots on Goal (note that these include Goals)
MS	Missed Shots: shot attempts that miss the net
BS	Blocked Shots: shot attempts that are blocked by the opposing team
SF	Shots For: the sum of all shots (SOG MS, BS) taken by a team
SA	Shots Against: sum of all shots (SOG, MS, BS) taken against a team
SD	Shot Differential: computed as SF - SA
$SD/60$	The shot differential expressed as a rate (per 60 minutes)
$SD\%$	Shots For expressed as a percentage of the shots attempted by both teams
ASF	Adjusted Shots For: sum of (SOG, MS) taken by a team
ASA	Adjusted Shots Against: sum of (SOG, MS) taken against a team
ASD	Adjusted Shot Differential: computed as ASF - ASA
$ASD/60$	The adjusted shot differential expressed as a rate (per 60 minutes)
$ASD\%$	Adjusted Shots For expressed as a percentage of the shots attempted by both teams
$onSH\%$	On-ice Shooting Percentage
$onSV\%$	On-ice Save Percentage
$FLAP$	Fluctuations from League Average Performance

CHAPTER 1

INTRODUCTION

1.1 A Simple Flip of a Coin

A natural place to start a discussion of statistics is with the simple idea of a coin flip. We'll consider an ideal coin in which there is a 50% chance of getting a "heads" result if you flip the coin. What we would like to explore is the following: if we have a set of coin flip results, can we trust these results to be a good predictor of future coin flip results.[1]

1.1.1 10 Coin Flips

Let's start with an experiment in which we flip a coin 10 times and record the number of times the coin lands head side up. Since doing this experiment only once doesn't give us any worthwhile data, we'll perform this experiment one million times. That is, we'll flip the coin 10 times and record how many heads we see - and we'll do this one million times. This type of simulation can be done fairly quickly on a computer and the results are shown below in Fig. 1.1. But before you take a look at these results, try running two or three experiments yourself (remember, each experiment is only 10 flips).[2]

The results of this experiment match common sense: the most likely result is to see five heads. What should also be clear to you (both from personal experience and the results) is that it is not uncommon at all to see four heads or six heads. In fact, even three or seven heads isn't unrealistic. If you performed this experiment just once and got three heads as the result, what would you expect to happen in the next 10 flips?

What you should take away from these results is that the results of one experiment involving 10 coin clips does not have much predictive benefit on future experiments.

[1] Another interesting question to ask here is at what point can you determine whether or not the coin flipping is rigged. If you get six heads after 10 flips, is the coin rigged? What about 60 heads after 100 flips? And 600 heads on 1000 flips? Be sure to check out the results of our three coin flipping simulations to see the answer.

[2] If you don't have a coin handy, there is a virtual coin flipper available at: http://www.random.org/coins/.

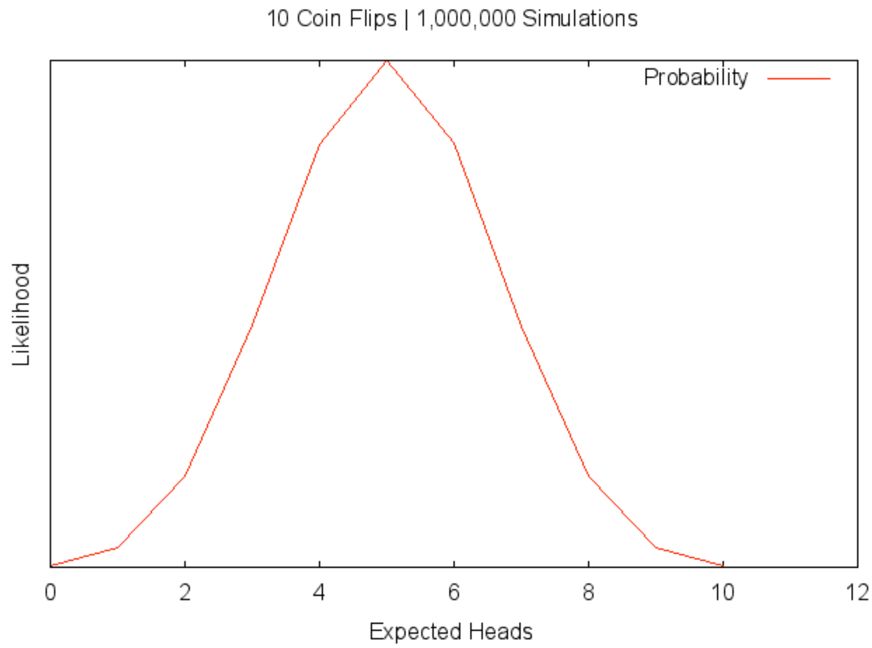

Figure 1.1. Results of a computer simulation of 10 coin flips run one million times.

Furthermore, using a single experiment as the basis for future projections could lead you to make rather erroneous statements regarding the coin.

1.1.2 100 Coin Flips

Next, we'll change the simulation so that each single experiment involves flipping a coin 100 times. And the simulation will again run one million experiments. The results are shown in Fig. 1.2. How often does an experimenter see 40% or 60% heads? Not too often at all. What about 30% or 70% heads? Something has changed. We've created more events by flipping the coin 100 times instead of 10 and the spread (or variance) of our results has narrowed. It is becoming less likely for us to see the results that were quite common in the 10 flip experiment.

1.1.3 1000 Coin Flips

Let's run one more simulation. Each experiment this time will involve 1000 flips of a coin. And the simulation will run one million times. Fig. 1.3 reveals the outcome of the simulation. Now we see a dramatic change in the data: there is virtually no chance at all of seeing a result in which 40% or 60% of the coin flips are heads. And 30% and 70%

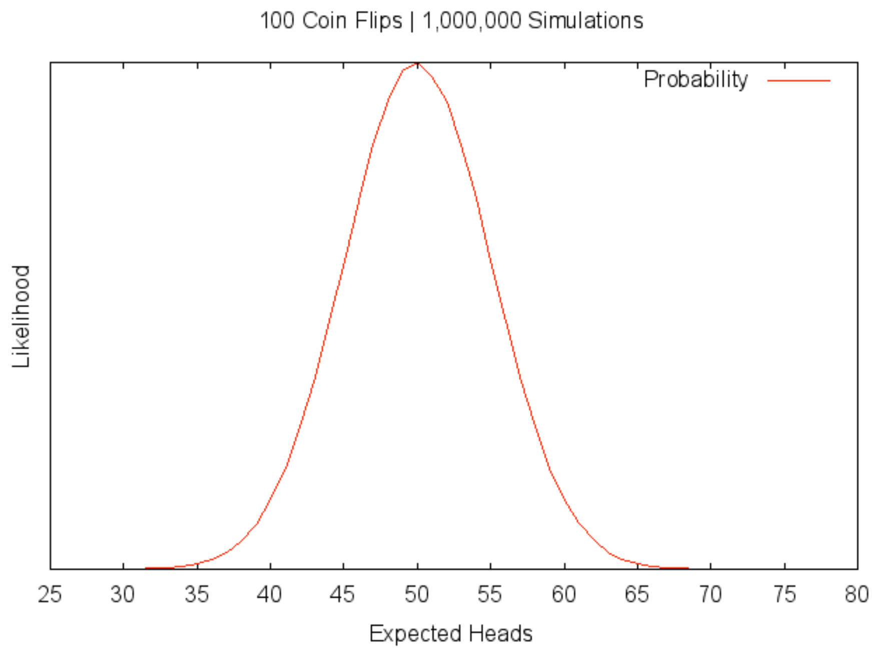

Figure 1.2. Results of a computer simulation of 100 coin flips run one million times.

don't even show up on the graph. We now have a fairly narrow range of expected results from this simulation.

1.1.4 Summary

Our original goal was to explore the relationship between sample sizes and their predictive abilities. With only 10 coin flips in your sample size, using the number of heads from a single experiment would be practically useless in determining how that same coin would behave in future experiments. As we increased the sample size to 100 coins, we gained additional predictive power in that our range of expected results narrowed to a window of about 40% - 60% heads. Finally, we boosted our sample size to 1000 coin flips and the benefit was immediately recognized: expected results narrowed to about 47% - 53%. Given the three simulations above, it is clear which one provides your best chance at predicting future performance: it is the simulation in which you have more data.

1.2 NHL Players as Coins

In the previous section, we ran simulations on the flipping of ideal coins - that is, coins with a 50% probability of landing with the head side up. What we'll do in this section is

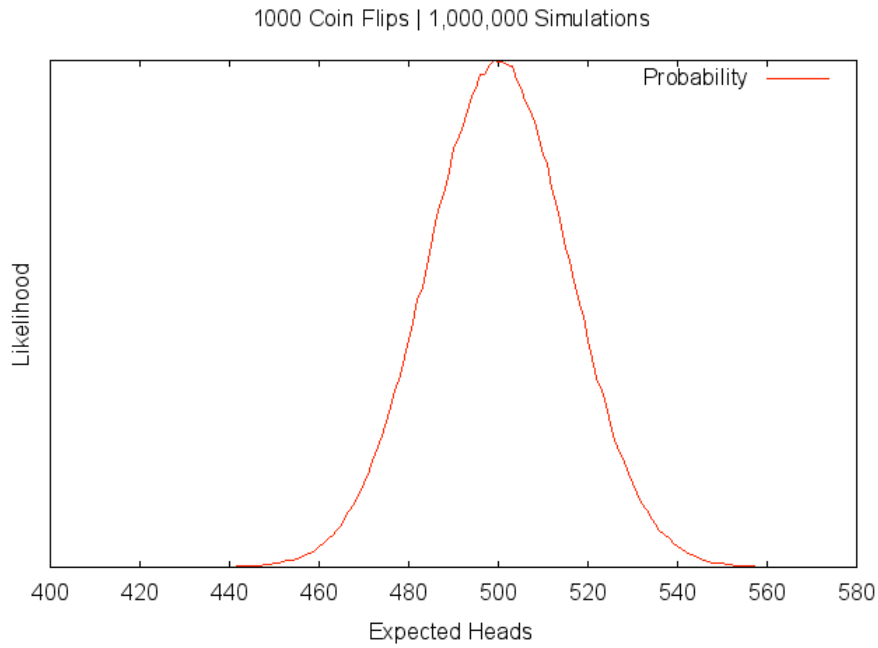

Figure 1.3. Results of a computer simulation of 1000 coin flips run one million times.

treat NHL players as coins. Instead of ideal coins though, we'll be using weighted coins (or biased coins). Here, the probability of landing on one particular side will not be 50%, but will instead be determined by how often that particular player scores goals.

1.2.1 Phil Kessel

Phil Kessel, a forward for the Toronto Maple Leafs, has a career shooting percentage (SH%) of 10.9%. What exactly does it mean to have a SH% of 10.9%? To explore this idea, we'll model Phil Kessel as a weighted coin. Instead of using heads and tails, we'll call the sides of the coin goals and non-goals. Kessel's coin model works in the following manner: when we flip the coin, there is a 10.9% chance that it will land with the goals side up and an 89.1% chance that it will land with the non-goals side up.

We're going to run three simulations for the Phil Kessel coin:

- a 60 flip experiment
- a 275 flip experiment
- a 1693 flip experiment.

You might be wondering why we have chosen such odd numbers for our experiments. The 60 flip experiment will represent Kessel's shot total in his first 15 games played in the 2012-2013 season; the 275 flip experiment represents the average number of shots that Kessel takes in a complete season; the 1693 flip experiment represents the total number of shots Kessel has taken in his NHL career.

The first experiment is a fascinating one; Kessel, a five-time 30+ goal scorer[3], had managed only two goals in his first 15 games last season (he had taken 60 shots on goal during this time frame). We'll use one million identical Kessels and have them take 60 shots on goal each. The results are shown in Fig. 1.4. Notice that the likelihood of

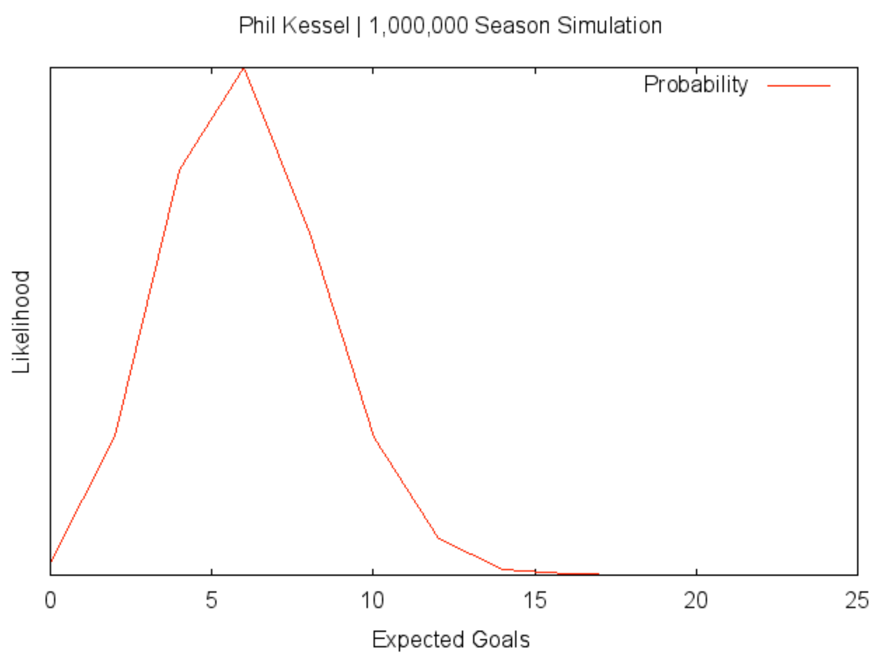

Figure 1.4. Results of a computer simulation of 60 SOG run one million times.

Kessel scoring only two goals in a 60 shot span (3.3% SH%) is not that high, but it's certainly possible. Had you tried using this 60 shot span to make assumptions about Kessel's future goal scoring, you would have been sadly disappointed. Kessel went on to score 20 goals on 161 shots, giving him a SH% of 12.4% for the season. Much like the 10 flips coin experiment, you can't use small sample sizes to make accurate projections of

[3]Kessel's 20 goal shortened season in 2012-2013 projects to 34 goals over 82 games.

future performance. Fantasy hockey managers who traded Phil Kessel early in the season became acutely aware of this notion.

What happens if we use a shot sample size of 275 shots on goal (recall this is the average number of shots on goal by Kessel over an 82 game span)? Fig. 1.5 provides us with a glimpse of the range of expected goals scored by Phil Kessel. Taking a look

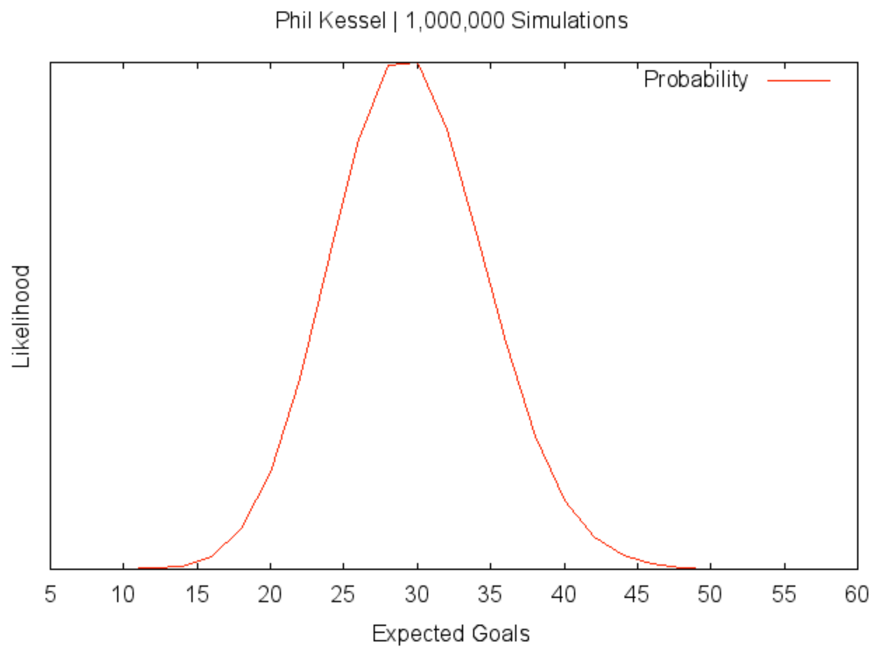

Figure 1.5. Results of a computer simulation of 275 SOG run one million times.

at these results should convince you of one thing: the more shot data you have on a player, the better your ability to project his future shooting percentage. Notice in the 60 shot experiment, that shooting percentages ranging from 0%-20% were in the realm of possibility. In the 275 shot experiment, that range narrows to about 5% - 15%. As it turns out, all of Phil Kessel's seven NHL seasons have produced shooting percentages within this exact range.

Finally, we look at Phil Kessel's career numbers. He has taken 1693 shots over the course of his career. With that data in mind, what happens to the range of expected SH%? It is now only about 5% points wide, ranging from about 8% to 13%. Incidentally, this gives us some insight into the talent level of Phil Kessel: given the 1693 recorded shots on goal, it would be very unlikely to see him record a SH% outside of the 8% - 13% range for an extended period of time.

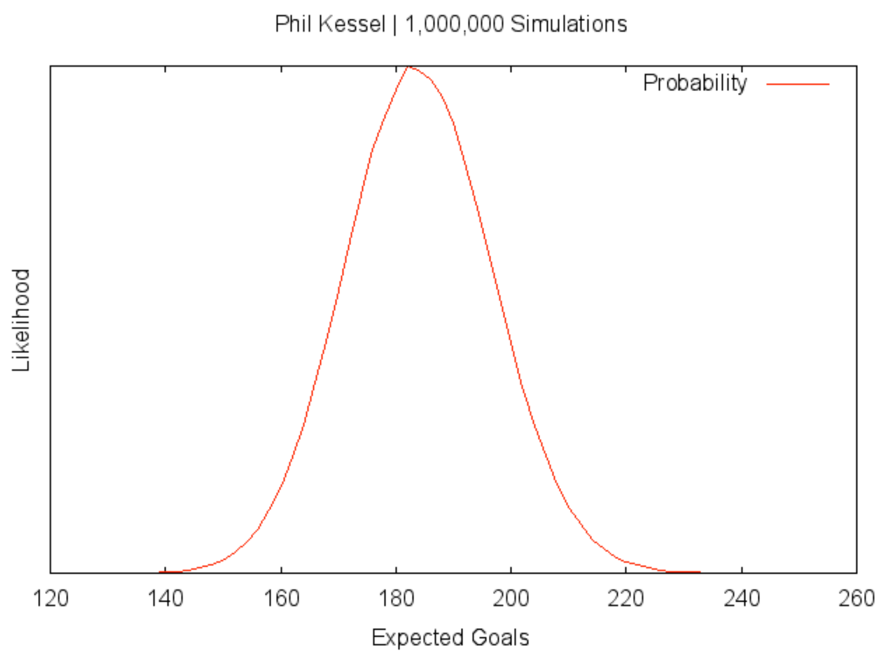

Figure 1.6. Results of a computer simulation of 1693 SOG run one million times.

1.3 Thoughts on Sample Sizes

It should now be quite clear to you that using small sample sizes gives you virtually no predictive power when projecting the future number of heads in a coin flip experiment or the number of goals scored by an NHL player. In a typical NHL season, approximately five players will score more than 40 goals. That means, for the overwhelming majority of NHL players, the sample size for goals in a season is less than 40. Given what you've just learned from the coin experiments, using goal data to project future goal scoring is going to produce inaccurate results.

Using shots on goal data should provide a significant boost. But, it would take an entire season to see sample sizes in the 200-300 event range (and that's for the high-end NHL shooters). The problem with projecting the performance of NHL players is now front and center: sample sizes of in-season goal and shot totals are too small to yield useful predictive benefits.

1.4 The Motivation for Analyzing Shot Data

If the number of goals or shots on goal isn't enough to form a useful data set, then how can this data set be constructed? The solution is to include all shot data in your data set. By all shot data, we mean goals, shots on goal that were saved, missed shots, and shots that were attempted but ended up being blocked by the opposing team.

You might be asking yourself how much you can really gain by switching from a goals analysis to a shots analysis. The question is fair and the answer may surprise you: for Phil Kessel's 2012-2013 season, instead of having only 20 events (goals) in your sample size, you end up with 313 events (all shots combined). And how about that 15 game stretch where Kessel scored only two goals? You'd have 114 events in your sample if you included all of his shot attempts.

The method of using all shot data to analyze NHL teams and players will be called *Shot Metrics*. An immediate benefit of using all shot data is that your sample sizes can be grown significantly, thereby allowing you to make more accurate player projections using data over shorter time intervals. In the next chapter, we'll introduce *Shot Metrics* to you and reveal other benefits of the system to assist you in your pursuit of hockey and fantasy hockey analysis.

CHAPTER 2

THE SHOT METRICS

2.1 Introduction

The idea of using all shots (and not just goals or shots on goals) to analyze NHL teams has been around for decades. Rather than analyze skaters (or teams as a whole), Jim Corsi first used shots to measure the workload of his goalies [2]. The first references to this type of analysis being used on skaters was the work done by Vic Ferrari[1] in a 2007 article about the Edmonton Oilers [3]. This analysis has since evolved into a greater understanding of which players and teams are driving play in the NHL.

2.2 Notation and Terminology

Currently, modern hockey analysis uses meaningless names[2] to represent the rather meaningful statistics that form the foundation of the field. In an effort to avoid this trap, we've chosen names that should give readers an immediate clue as to the definition of the statistic they are reading about. It is worthwhile to discuss the new notation and terminology here.

2.2.1 Simple Shot Counting

We'll start with the most basic statistic of *Shot Metrics*: shots for. Shots for (SF) is simply the total number of shots taken by a player or team. The total number of shots taken by a team (shots for) is found as follows:

$$SF = SOG + MS + BS, \tag{2.1}$$

where SOG represents shots on goal (both those that are saved and those that result in goals), MS represents shots taken that missed the net, and BS represents shots taken

[1] This is the blogger's pseudonym.

[2] In some instances, they have renamed basic hockey concepts that have existed since the beginning of the game. For example, the term Corsi is being used to represent the total number of shots a team has taken. We already have a word for that: it's called shots.

that were blocked by the opposing team.[3] An equivalent metric can be defined for shots taken by the opposing team. This sum is known as shots against (SA) and is computed in the same manner:

$$SA = SOG + MS + BS. \tag{2.2}$$

Together, Eqs. (2.1) and (2.2), form the basis for most of the statistics used in *Shot Metrics*. All future terminology in this document will depend on these simple definitions.

2.2.2 Shot Differential

A more useful look at shot counting involves looking at a team's shot production compared to their opponent's shot production. The shot differential SD is defined as shots for minus shots against. It is a simple measure of whether or not a team is outshooting their opponent. The shot differential is computed as follows:

$$SD = SF - SA. \tag{2.3}$$

Alternatively, one can look at a team's shot production relative to the total number of shots in a game (or season). By expressing the shots for as a percentage of the total shots for and against, it becomes immediately clear if a team is outshooting their opponents. The shot differential percentage $SD\%$ is defined as:

$$SD\% = 100 \times \frac{SF}{SF + SA}. \tag{2.4}$$

A value of 50% for the $SD\%$ would indicate that the team played in a balanced game (or season) as measured from a shot differential perspective. For the 2012-2013 season, the top four teams in $SD\%$ were:

- Los Angeles Kings (56%)
- New Jersey Devils (56%)
- Boston Bruins (54%)
- Chicago Blackhawks (54%).

[3]This is a good place to point out that these individual shot definitions are mutually exclusive. That is, for example, a shot cannot be defined as both a missed shot and a blocked shot. Each shot taken falls into no more than one of the three types: shots on goal, missed shots, and blocked shots.

Given that three of the teams listed above played in the conference finals, it would seem that $SD\%$ might be a reasonable approach to discovering teams destined for playoff success.

You can also compute SD and $SD\%$ for individual players. The formulas are identical to those for the team calculations. You simply compute the SF and SA for events while that particular player is on the ice. There is one useful modification you can make when analyzing individual players: correcting for the differences in time-on-ice. Since some players see 25 minutes of ice time and others only see 8 minutes, their shot totals would be wildly different. To account for this, we simply express the shot differential as a rate statistic. We've chosen 60 minutes as our time frame. This stat is written as $SD/60$ and represents the expected shot differential for a player over a 60 minute interval of time. For reference, the top five skaters (as measured by $SD/60$ in the 2012-2013 season) were:

- Justin Williams (LA)
- Jake Muzzin (LA)
- Patrice Bergeron (BOS)
- Anze Kopitar (LA)
- Tyler Seguin (BOS).

2.2.3 Adjusted Shot Differential

As mentioned earlier, the core reason for using *Shot Metrics* is that they serve as a reasonable proxy for a team's puck possession. If a team is consistently outshooting an opponent, then that team has possession of the puck more often. An argument has been made [4] that these *Shot Metrics* can be used as a proxy for scoring chances and therefore the inclusion of blocked shots is really not necessary at all.[4] Thus, a simple adjustment can be made to the *Shot Metrics* to account for this change. We'll redefine the basic SF and SA metrics by removing the blocked shots component. Thus, we have the adjusted shots for and adjusted shots against defined as:

$$ASF = SOG + MS \tag{2.5}$$

[4]Furthermore, blocked shot totals suffer from rink bias - an apparent inflation of certain stats by home arena statisticians.

and

$$ASA = SOG + MS. \tag{2.6}$$

With these adjusted definitions for ASF and ASA, we can easily compute the remaining modified metrics. For completeness, these metrics are:

$$ASD = ASF - ASA \tag{2.7}$$

and

$$ASD\% = 100 \times \frac{ASF}{ASF + ASA}. \tag{2.8}$$

2.2.4 Fluctuations from League Average Performance

During the 2012-2013 season, NHL goalies posted an average save percentage of .920, while NHL skaters posted a .080 shooting percentage.[5] League averages for both save percentage and shooting percentage have remained steady over the past six seasons as Table 2.1 reveals. Imagine if you were to take the SV% and SH% values above and add

Table 2.1. League Average Data for Save Percentage and Shooting Percentage

Season	EVSV%	EVSH%
2012-2013	.920	.080
2011-2012	.921	.079
2010-2011	.921	.079
2009-2010	.919	.081
2008-2009	.919	.081
2007-2008	.920	.080

them together (and, for convenience only, multiply that sum by 1000). It turns out, that the result of this quick calculation is 1000. That is, if you sum the league averages for SV% and SH% you'll get the number 1000.[6]

Next, imagine calculating this same sum for each of the 30 NHL teams individually. We'll post a few of these here for you from the 2012-2013 season:

[5] Numbers computed using Left Wing Lock internal data.

[6] Intuitively, this should make sense to you since every shot on goal must result in either a save or a goal. So, the sum of SV% and SH% should be 1 (which becomes 1000 after we perform our multiplication of convenience).

- Chicago Blackhawks - 1019
- New York Rangers - 1009
- Ottawa Senators - 994
- Calgary Flames - 972.

These are just a few examples for random NHL teams. What is particularly interesting about these sums is that teams are unable to sustain (for very long times) sums that stray too far from 1000. In 2008, Tyler Dellow performed a study on these sums [5] that revealed that teams generally regress back to a sum of 1000 given enough time.[7] In Dellow's simple and powerful study he looked at five years of NHL data and broke each season into four quarters. He computed the sums (SV% + SH%) for each NHL team at the end of the first quarter of each season and then computed the sums again for the remaining portion of each season. The results of the study are stunning:

- the 20 best teams from the first quarter had average sums of 1031;
- the 20 best teams dropped to 1005 in the remaining three quarters of the year;
- the 20 worst teams from the first quarter had average sums of 970;
- the 20 worst teams jumped to 998 in the remaining three quarters of the year.

These sums (SV% + SH%) will be called *Fluctuations from League Average Performance*, or FLAP for short. Briefly, these FLAP values provide a measure for how much a team is straying from their expected performance. The take-away is this: a sample of teams with a high average FLAP is likely to see their future performance decline; on the other hand, a sample of teams with a low average FLAP will likely see a rise in their future performance. This idea of changing FLAP values is an example of a statistical phenomenon known as regression toward the mean.[8]

Like all other *Shot Metrics*, FLAP can be computed for individual players as well. To perform the calculation for an individual player, you want to find the team's SV% and SH% only for the times when that particular player was on the ice. These two percentages are known as *onSV%* and *onSH%* respectively. To be clear, you are not simply using the SH% of the individual player.

[7]The original reference to these sums is generally credited to Brian King.

[8]See the Appendix for a detailed explanation of regression to the mean.

2.3 Assumptions of the Model

2.3.1 EV, PP, and PK

You have probably noticed that both goal scoring rates and shot production rates increase dramatically when a team is on the power play (PP) as compared to when a team is at even-strength (EV).[9] Likewise, when a team is on the penalty kill (PK), the number of shots they face rises. Rather than allow these lopsided shooting situations to affect the overall shot data, it is common practice to remove this data entirely from the set. Therefore, when we speak of *Shot Metrics*, we will always be referring to game situations involving even-strength hockey (more precisely, five skaters vs. five skaters)[10].

2.3.2 Score Effects

Another aspect of a hockey game that can affect shot differentials is the score. Consider a hockey game that is close in score (perhaps tied) in the third period. Both teams are motivated to score and continue to throw shots at the opposing net in an effort to win. Contrast that game situation with one where a team has a two or three goal lead in the third period. The team with the large lead is likely to employ a defense-first strategy, thereby reducing their shot output. Not only does a defense-first strategy lead to less shots by the team with the lead, it also yields more shots by the team without the lead. The shot reduction by the leading team combined with the shot increase by the lagging team skews any measurement of shot differentials.

There are a number of ways to account for these effects. The most rigid approach is to remove them from the data set entirely. In this instance, one would compute shot differentials only when the game is tied. But, this approach seems contrary to one of the major advantages of using *Shot Metrics*: the benefit of larger sample sizes over small time intervals. An alternative to throwing out all of the non-tied game situation data is to use game data when the score is tied or *close*.[11] A slightly more complicated approach

---

[9]Teams can generally stop 92% of SOG at even-strength. That number drops to 87% when a team is on the penalty kill.

[10]Note that the NHL considers the end moments of a game with a goalie pulled to be a 5-on-5 EV situation (even though one team has six skaters on the ice). The data at Left Wing Lock removes these extra-attacker situations (which occur any time a goalie is pulled) because we believe they don't accurately reflect what we're trying to capture when we say "even-strength" situation. In summary, *Shot Metrics* data is computed using 5-on-5 EV situations with all extra-attacker situations removed.

[11]The term close here refers to game situations when the lead is no more than one goal in the first two periods or the score is tied in the third and overtime periods.

has been suggested recently [6] which adjusts the shot differentials for situations when a game is not tied. This proposal has the benefit that it allows you to keep all of the 5-on-5 data from a game no matter the score. It should be pointed out that the last of these three suggestions appears to have the greatest predictive value (based on testing). This should not surprise you since it is the method with the largest sample size. In the end, whichever score effect adjustment you prefer, know that they ultimately will agree with one another the deeper you get into the season.

CHAPTER 3

APPLIED SHOT METRICS

3.1 Introduction

Now that you've made the effort to become familiar with the basic stats of *Shot Metrics*, it is time to learn how to use them to your benefit. In this section, we'll present some of the major concepts of *Shot Metrics* in the form of case studies.

3.2 At the Team Level

3.2.1 Paper Tigers

The Merriam-Webster dictionary [7] defines paper tiger as: one that is outwardly powerful or dangerous but inwardly weak or ineffectual. This is the perfect term to describe NHL teams that (for a time) live high in the standings despite the fact that they simply aren't very good.

The appearance of paper tigers in the NHL is typically a once-per-season event with the classic example being the Minnesota Wild of 2011-2012. The Wild, according to the NHL standings, were one of the best teams in the league. On November 30, 2011 the NHL standings had the Wild ranked number one in the league with 33 points on a 15-7-3 record [8]. Wild fans and bloggers alike [9] became believers overnight and invented narratives surrounding the meteoric rise of their franchise.

Just 70 days later, the Minnesota Wild would find themselves out of the playoff picture and sitting 10th in the Western Conference. How could this happen and was there a way to predict the likelihood of such a catastrophic drop in the standings? The answer to the first question is easy: the Minnesota Wild simply weren't very good. The answer to the second question requires more finesse: yes, there was a way to know that the Wild weren't as good as the standings indicated but, no, there was not a way to predict the enormity of the collapse.

To understand why the 2011-2012 Wild weren't a very good team, we'll use two *Shot Metrics* in concert: ASD (Adjusted Shot Differential) and FLAP (Fluctuations from

League Average Performance). It pays off to remember that shot differential statistics are currently the best indicators of true team talent (that is, projecting future team performance is most accurately done by an assessment involving shot differentials). It turns out that the Wild had the lowest ASD value in the league [10]. This means the Wild were consistently being outshot when on the ice for even-strength hockey; the opposing teams were dominating play. In the short term (Oct - Nov) an outshot team can get lucky, but over the course of a full season, teams that consistently get outshot at even-strength generally don't make the playoffs. And what about their FLAP? The Wild maintained a FLAP of 1021 through mid-December of 2011 [11]. As you'll recall, FLAP values regress heavily to the mean. Generally, if a team maintains a high FLAP value over a given time interval, they are likely to see their performance drop in the future. And boy the did the performance of the Wild drop.[1]

The important point to take away from the 2011-2012 Wild is: the FLAP and ASD values of a team can provide you with insights into a team's talent level that ordinary standings or a goals-based analysis gets wrong (at least early in the season).

3.3 At the Player Level

In 2011-2012, the top five player values for FLAP belonged to Chris Kelly, Tyler Ennis, Rich Peverley, Jacob Josefson, and Alexander Steen. These five players would have been among the group of players most likely to experience a regression in their performance for the following season. Table 3.1 compares the FLAP values for these five players in the two relevant seasons.

Table 3.1. Regression of FLAP Values: 2011-2012 Season to 2012-2013 Season

Player	2011-2012 FLAP	2012-2013 FLAP
Chris Kelly	1056	960
Tyler Ennis	1056	987
Rich Peverley	1056	957
Jacob Josefson	1053	902
Alexander Steen	1045	1006

[1] As noted above, the actual drop for the Wild was much more significant than what was to be expected from a regression to the mean standpoint. The Wild's FLAP value for the remainder of the season was 972.

For the 2012-2013 season, Chris Kunitz of the Pittsburgh Penguins scored a FLAP of 1074. In recent years, his FLAP values have been: 1074, 1002, 1030, 995, and 1005. Since his 2012-2013 value differs from 1000 significantly, Kunitz will be in the group of players most likely to suffer a drop in production next season.

3.4 Summary

3.4.1 Final Thoughts

We hope this brief introduction to *Shot Metrics* sets you on a path for a deeper understanding of the game of hockey. We've carefully chosen stat names in an effort to make this type of hockey analysis more accessible to all fans of hockey. We've also taken great care in writing the definitions for each metric you'll need as you learn how to perform a shots-based analysis of an NHL team or player. While we only had the space to share a few applications of these *Shot Metrics*, you should look forward to many articles at our website on the subject in the coming months.

3.4.2 Shot Metrics Web Tool

We're excited to add *Shot Metrics* to the list of tools available at our website. As with our other tools for hockey fans and fantasy hockey managers, this one will be completely free to access. Enjoy the data and don't hesitate to contact us with questions or suggestions. The new tool is located at: http://www.leftwinglock.com/shot-metrics/.

APPENDIX A

MAPPING SHOT METRICS TO ADVANCED STATS

Preliminaries

Data for these *Shot Metrics* are now available at the Left Wing Lock website. There are other websites [12, 13] that publish this data under a different naming convention. Below, you'll find the appropriate mappings for the different naming conventions for your reference.

SD	Presented as Corsi Number on other sites
$SD\%$	Presented as Corsi For Percentage on other sites
$SD/60$	Presented as On-Ice Corsi (or Corsi On) on other sites
ASD	Presented as Fenwick Number on other sites
$ASD\%$	Presented as Fenwick For Percentage on other sites
$ASD/60$	Presented as On-Ice Fenwick (or Fenwick On) on other sites
$FLAP$	Presented as PDO on other sites

APPENDIX B

REGRESSION TOWARD THE MEAN

Shooting Percentage

In the fantasy hockey draft kit that Left Wing Lock, Inc. publishes every season [14], a subset of NHL players most likely to see a large decrease in scoring production is selected from over 700+ NHL players. This particular chapter of the draft kit is always well-received and provides managers with a huge advantage going into their fantasy hockey drafts. For the 2012-2013 season, our staff hand-picked 18 NHL players (names ranging from Nathan Horton to Radim Vrbata to Scott Hartnell). At the end of the season, we reviewed these 18 selections and found that our predictions were correct in 17 cases - a success rate of 94%. In fact, in all of the seasons that we've made these predictions, our lowest success rate has been 89%. How do we do it?

The basic idea behind our selections is to compile a list of NHL players with season-long shooting percentages significantly different from their career shooting percentages. We are essentially creating a subset of NHL players who all suffer the same defect: they exhibited an extreme value for the variable in which we were interested (SH%).

An individual player's SH% for a season is a marriage of skill and luck (or randomness, if you prefer). Sometimes that luck is unfavorable and the player has an "off" year. Other times that luck is favorable and the player has a "career" year. Most of the time, these factors are reasonable in size and a player fluctuates about some career SH% average. But sometimes these luck factors are large and lead to extreme SH% values (when compared to the player's career SH%). It is these extreme values that lead to a player being selected by the Left Wing Lock staff.

Since this subset of players had to be both good and lucky to experience the wild SH% swing, then it is clear that they are playing above their true talent level. If you were to continue to watch them play (i.e. let them take more shots), the randomness in the SH% of this subset of players will push them closer to their career averages. This swing in SH% (from extreme toward average) is known as *regression toward the mean* and is a phenomenon witnessed in fields as varied as genetics, sports, finance, and coin flipping.

Some factors of regression toward the mean that are important include:

- it can be positive or negative depending upon which side of the average your sample is on
- the amount of regression your subset experiences depends upon how extreme your subset is
- the phenomenon reflects upon the entire subset and doesn't necessarily provide a correct prediction for an individual in that subset
- it is a statement about the likelihood, but not the certainty, of something happening
- any phenomenon that involves a level of randomness (large or small) will be subject to regression toward the mean.

Because luck (or randomness) by its very definition is non-repeatable, if you can isolate a subset of players who were subject to favorable luck during a short time interval, then you'll be able to predict (with good accuracy) which players will see a reduction in their overall numbers in the future.

REFERENCES

[1] Gabe Desjardins. Corsi and score effects. http://www.arcticicehockey.com/2010/4/13/1416623/corsi-and-score-effects, April 2010.

[2] Jesse Spector. Who is jim corsi? http://www.sportingnews.com/nhl/story/2013-04-23/jim-corsi-hockey-definition-nhl-advanced-stats-primer-behind-the-net-fenwick-pdo, April 2013.

[3] Vic Ferrari. Corsi numbers. http://vhockey.blogspot.ca/2007/10/corsi-numbers.html, October 2007.

[4] Matt Fenwick. Driving possession. http://vhockey.blogspot.com/2007/11/driving-possession.html, November 2007.

[5] Tyler Dellow. Pdo numbers. http://www.mc79hockey.com/?p=2996, November 2008.

[6] Eric Tulsky. Adjusting for score effects to improve our predictions. http://www.broadstreethockey.com/2012/1/23/2722089/score-adjusted-fenwick, January 2012.

[7] Merriam-Webster Dictionary. Paper tiger. http://www.merriam-webster.com/dictionary/paper

[8] ShrpSports. Nhl standings. http://www.shrpsports.com/nhl, November 2011.

[9] JS Landry. Now and then through 30. http://www.hockeywilderness.com/2011/12/11/2628263/now-and-then-through-30, December 2011.

[10] Daniel Wagner. The minnesota wild regressed. http://blogs.thescore.com/nhl/2012/02/23/the-minnesota-wild-regressed-what-about-the-nashville-predators/, February 2012.

[11] The Neutral. Non-traditional metrics glossary: Pdo. http://www.fearthefin.com/2012/8/2/3215351/non-traditional-metrics-glossary-pdo, August 2012.

[12] Gabe Desjardins. Behind the net. http://www.behindthenet.ca/, July 2013.

[13] David Johnson. Hockey analysis. http://hockeyanalysis.com, July 2013.

[14] Left Wing Lock. Fantasy hockey draft kits. http://www.leftwinglock.com/draftkits/, July 2013.