

# How Much Do Faceoffs Matter? Translating Faceoffs to Goals, Wins, and Championships in Hockey

Other Sports  
193940

## Abstract

Occurring around 60 times per game, hockey faceoffs uniquely start without any team possessing the puck but invariably end with one team in control. While the importance of faceoffs has long been acknowledged, their actual impact on scoring outcomes remains inadequately measured. It is acceptedly evident that a center winning 54% of their faceoffs outperforms one with a 51% success rate, but the tangible extent of this advantage in terms of goals, wins, and losses remains underexplored. This research fills the void by continuing the effort to translate faceoff results to scoring outcomes, measuring faceoff performance in goals, wins, and losses in a novel manner. We explore evidence that faceoffs are an undervalued championship-caliber market inefficiency and offer models enabling General Managers to see role-specific projections of how different personnel and usage could maximize offense, defense, and championship chances.

## 1. Introduction

In last year's tightly contested playoff series between the Toronto Maple Leafs and Tampa Bay Lightning, faceoffs proved to be a series-changing decider of outcomes. Punctuated by Morgan Rielly scoring a game-winning, overtime goal off an assist from victorious faceoff taker Ryan O'Reilly, the series involved four goals scored within 11 seconds of a faceoff victory where the faceoff win deterministically established the puck possession that offered the opportunity for game-changing offense. To be clear, there is no automatic relationship between faceoffs and goal scoring, but the example demonstrates how the newfound puck possession afforded by a faceoff win can be of immense offensive value. This paper analyzes the importance and value of faceoffs in the context of the incremental puck possession and offense they afford. We estimate the value of a faceoff win in goals and explore the implications of faceoff talent on optimal roster construction strategy and team success.

### 1.1. Existing Literature

Our research is inspired by the belief that faceoffs remain an underexplored topic capable of winning games. Several articles <sup>[1], [2], [3]</sup> have been written on the topic of the importance of faceoffs, with each basically arguing that they are not important because faceoff win percentage in its raw form is fairly uncorrelated with goal-scoring metrics or game win percentage. Beyond these articles, there are two key papers on the topic to date that relate to our approach and research question.

In "Winning Isn't Everything: A Contextual Analysis of Hockey Faceoffs," a product of the 2019 Sloan Sports Analytics Conference Research Paper Competition, by Czuzoj-Shulman et al of Sportlogiq <sup>[4]</sup>, the researchers focus on "how directionality, clean wins, and player handedness play a significant role in creating value." While this paper does consider "events following the face-off," it

primarily focuses on quantifying statistical correlations between different physical characteristics of how a faceoff is won and the events that follow. We build upon the concept of analyzing win value but consider it at a more macro level strictly centered around impact on scoring. Czuzoj-Shulman et al compute a probabilistic weighted average based on win value (implied by directionality, handedness, cleanliness) and, separately from model creation, compare it to expected goal differential. We use key performance indicators of expected goal creation as a direct objective for our models and train directly on these objectives to quantify the direct rather than coincidental impact of faceoffs on goals.

This difference may seem subtle, but the added value is substantial. Rather than quantifying how a victorious center can deliver value in ten different ways based on the mechanics of how he or she won the faceoff, we take a higher-level view of quantifying the number of goals a center is expected to gain or cost his or her team through faceoff performance in any deployment and game situation. We agree with and further strengthen Czuzoj-Shulman et al's argument that all faceoff wins are not created equal in a complementary manner. However, instead of relying on the mechanics of stick movement or the physical characteristics of how a faceoff is won, we focus on making a direct linkage between faceoffs and goals and quantify how much faceoffs are winning and losing teams games.

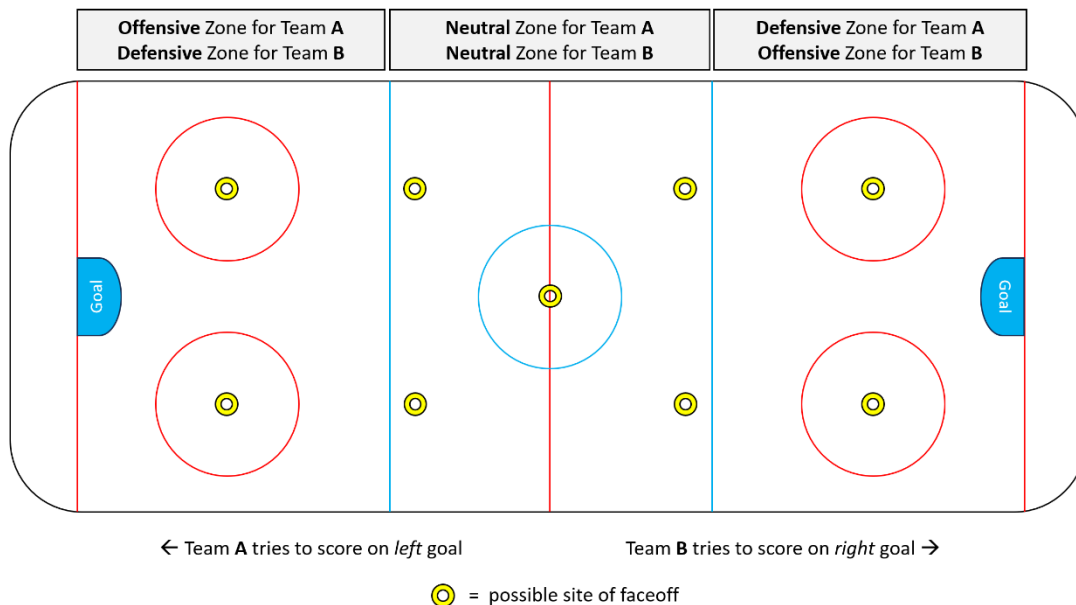
In "An Analysis of NHL Faceoffs" by Schuckers et al of St. Lawrence University [5], the researchers suggest that a center improving from "winning 50% of their faceoffs to winning 60% of them gains just over 12 goals per season which is equivalent to two additional wins."

We utilize a very similar approach to Schuckers et al but introduce several modifications that we believe build upon their model to offer potential for a stronger conclusion. First, we relax the assumption made by Schuckers et al that "after 20 seconds the impact of an individual event in the NHL is noise with the exception of penalties," instead creating a concept of faceoff attributability grounded in transitions in play to yield a more accurate mapping of faceoffs to scoring plays. Second, we frame our analysis not just in terms of goal-related metrics but also in terms of gains in zone time, aiming to consider possession time as a more projectable complementary proxy for goal generation. Third, and perhaps most significantly, we design our models not at the center-by-center level but rather at the line-by-line level, with the aim of more strongly accounting for the linemates a center is deployed with and against and their complementary role in goal generation in the aftermath of faceoffs. To reiterate, we consider linemates rather than just the two centers specifically. We consider a variety of prior-year performance statistics for every player on the ice and provide an architecture for understanding how a center's impact on goal scoring via faceoff performance would adjust if they were deployed with different linemates in different situations. We utilize these updates to our approach to build on and review the results of Schuckers et al and further explore the relationship between faceoff wins, goal scoring, and season outcomes.

## 2. Methods

### 2.1. Objectives

In developing objective characterizations of faceoff performance that drive our models, we consider the different locations at which a faceoff can occur, as reflected in the following diagram. We then breakdown our approach in different ways based on faceoff location.



### 2.1.1. Offensive Zone Faceoffs

A key dimension of our research is factoring offensive events that follow a faceoff into a valuation of each faceoff. However, clearly not every offensive event that follows a faceoff is related to faceoff performance. At one extreme, if every offensive event was considered to directly relate to its nearest preceding faceoff, then in effect every goal scored in a game would be considered to be a byproduct of faceoff performance, which would certainly be an inaccurate claim that defeats the purpose of a faceoff model by every analytic or eye test. At the other extreme, it would be reckless to consider faceoff outcomes and scoring entirely disjoint, as teams consistently prove by generating offense after faceoff wins before the puck clears the offensive zone.

To recognize the role of faceoffs in empowering offense within reason, we introduce the concept of *faceoff attributability*, evaluating faceoffs in terms of the incremental offense generated for the winning team and the offense taken away from the opposing team. We use zone changes—transitions of the puck between the offensive, neutral, and defensive partitions of the rink—as a barometer for faceoff attributability, considering any incremental offense generated by the faceoff-winning team before a zone change occurs. This appears to be novel at least in the public realm, and the premise is simple but powerful: much of the offense generated after a faceoff win but before a zone change would not have been possible without the faceoff win to establish offensive zone possession. We say “much of” to account for the fact that, even when the defensive team wins a faceoff, the puck will be successfully cleared out of their defensive zone only a subset of the time, meaning even after a faceoff loss the offensively positioned team will *sometimes* still get to generate *some* offense even after losing the faceoff. This is where the “incremental” element comes into play: by comparing the typical offense generated after an offensive zone faceoff win to the typical offense generated after an offensive zone faceoff loss, we have a mechanism for quantifying the incremental value in goals associated with winning the faceoff.

In this spirit, our study defines two objective functions on which we train our supervised machine learning models:

1. Net expected goals: the cumulative expected goal differential gained through offensive zone faceoff wins and forfeited through faceoff losses. This is very similar to the metric used by Schuckers et al in their analysis.
2. Incremental zone time: a goal-scoring-agnostic measure of the expected gain in offensive zone time achieved via a faceoff win relative to a loss. This appears to be a new consideration that we haven't seen before in the public realm.

### **2.1.2. Defensive Zone Faceoffs**

So far, our discussion of net expected goals and incremental zone time has been limited to offensive zone faceoff wins, which begs the important question of how we quantify faceoff wins where the defensively positioned team is victorious. In a manner analogous to Schuckers et al, we structure our analysis so that defensive-zone faceoff wins and offensive-zone faceoff wins are two sides of the same coin (and apply this methodology to our new incremental zone time metric as well). Specifically, for the defensive team, the net expected goals and incremental zone time are each considered to be the negative of that credited to the offensive team. The idea is that faceoff wins in defensive situations should be valued in terms of the opportunity cost of the net expected goals and the incremental zone time the offensively positioned team would have been expected to accumulate had they instead won the faceoff. It is not perfectly intuitive, but a key feature of this structure is that, despite typically involving negative numbers, net expected goals and incremental zone time for the defensive team are still frequently positive, indicating credit to the defensive center. For example, if a defensive center's team surrenders fewer expected goals or less zone time on average after winning a defensive zone draw than after losing a defensive zone draw, this framework would capture the incremental positive value of the defensive zone win and credit the defensive center for seizing would-be offense from the offensive team upon any defensive zone faceoff win.

### **2.1.3. Neutral Zone Faceoffs**

We also wish to mention neutral zone faceoffs, which are by design systematically excluded from our models. Since the top ingredient in most public expected goals models is shot distance, it is extremely hard, if not impossible, to generate any meaningfully nonzero expected goal shots from the neutral zone. To be clear, the expected stats here likely match the eye test... when is the last time you saw someone score a goal on a shot taken from the neutral zone? Moreover, unlike offensive zone time, there is not necessarily strong value to be gained from possessing the puck in the neutral zone for longer periods of time since it is largely a zone of transition between offense and defense. Accordingly, we omit neutral zone faceoffs from consideration.

## **2.2. Data**

A significant part of the lift of this research was creating a dataset from public data capable of mapping faceoffs to scoring outcomes through metrics like net expected goals and incremental zone time. This starts with utilizing play-by-play data from *Evolving Hockey* <sup>[6]</sup> but also requires introducing new data from a variety of origins to conceptualize net expected goals and incremental zone time, which rely on the ability to detect zone changes—an event that is pivotally omitted from the *Evolving Hockey* play-by-play data.

To address this, we connect the play-by-play data to hand-tracked “microstat” data by Corey Sznajder through his *All Three Zones* project [7], which includes manual annotation of zone entries and exits for a subset of NHL games each season. This is far more complex than merging two traditional datasets due to the stark difference in how the play-by-play and *All Three Zones* data are tracked, the involvement of human labeling with the latter, and the lack of consistent identifiers between the two datasets, so significant inference-oriented logic was designed to create this linkage.

To our knowledge, the *Evolving Hockey* play-by-play data and Sznajder *All Three Zones* data haven’t been integrated in existing public code repositories. Our research includes extensive code in R to make this linkage. This code, which could in time become an R package, in some ways could be considered a second source of value in our project: a tool to let at-home “armchair analysts” like us work with play-by-play data that in effect contains zone changes as events, which for many projects can be the missing ingredient that makes the impossible doable.

Our model is trained entirely on data from games played in the National Hockey League (NHL), but we anticipate that it would generalize with minor adjustments to other leagues, such as the Professional Women’s Hockey League (PWHL), American Hockey League (AHL), Ontario Hockey League (OHL), East Coast Hockey League (ECHL), Western Hockey League (WHL), Quebec Major Junior Hockey League (QMJHL), and international play as well if comparable data was provided. One of the key reasons we only use the NHL in training our model is data availability: the NHL has by far the most data available for research.

## 2.3. Models

### 2.3.1. Structure

We train at the faceoff level, meaning that each record in our training set is a faceoff labeled with the net expected goals or incremental zone time in its aftermath as the objective for supervised learning. That being said, we primarily care about our models’ performance at the seasonal level and evaluate it as such. The reason for this is several-fold. First, our training data is restricted to the subset of all games watched and tracked by Corey Sznajder in the *All Three Zones* data, which is an impressively large number of games but far from every game played during the season. Accordingly, our training and validation datasets are not full seasons in their own right, lending to training at the faceoff-by-faceoff level. Second, we want our model to capture the full distribution of possible outcomes of each faceoff even if these typically aggregate to a more confined set of outcomes over repetition throughout a full season.

### 2.3.2. Architectures

Over the course of various analyses, we have trained a combination of extreme gradient boosting (XGBoost), neural networks, and random forest models. To assist with reproducibility, we are including commentary here on what parameters we tuned and how we structured our models and design choices we made to help minimize the potential for overfitting. Our XGBoost models are trained with ten-fold cross validation to help ensure robustness with hyperparameters tuned via racing using ANOVA methods. Our neural networks utilize batch normalization to help prevent covariate shift, dropout to help curtail feature coadaptation, and ReLU activation and perform hyperparameter tuning via grid search. The random forest models featured tuning of the maximum

depth, maximum number of features, minimum number of samples on a leaf, minimum number of samples required for splitting, and number of estimators.

### 2.3.3. Features

In terms of features, our model gets nearly 100 fields of context about the players on the ice. These features encompass everything from typical box score stats to relative-to-teammate metrics, actual and expected metrics, on-ice play driving statistics, and team-level context info.

For each of these metrics, we consider performance in the prior year. To be explicit, if the faceoff is year  $n$ , we pull the statistics from year  $n - 1$ . We do this to prevent bias from infiltrating the dataset. First, we want to make sure events that occur in a game at or after the game of the faceoff don't influence our analytical conception of the faceoff. Second, we want to make sure that faceoffs that occur later in the season don't benefit from features stabilizing to a greater extent as the season has elapsed than they would have for faceoffs earlier in the season.

### 2.3.4. Feature Encoding

Initially, we aimed to provide contextual stats for every player on the ice, thinking that our model could benefit from as much unbiased information on players participating in a faceoff and its aftermath as possible. As is typically best practice in sports analytics, we withheld player IDs and names from the model to ensure the model learns about important underlying characteristics of players' effectiveness rather than simply the signature of the unique identifiers of players who are strong performers. However, a key issue that arrives with these two considerations is that, when you provide a measure of a given metric for each player on the ice, there needs to be some reasonable logic as to which player's metric is encoded in which feature.

We encoded a role for each player—one of F1, F2, F3, D1, D2, or G1 in typical even-strength situations indicating three forwards, two defensemen, and one goalie on the ice—where F1 is the top-ranked forward by *Evolving Hockey's* Goals Above Replacement (GAR) method, F2 is the second-ranked forward, and so on (with D and G indicating defensemen and goalies, respectively). We quickly realized that what player is encoded into what role is increasingly arbitrary with respect to any given metric even when encoding by positional GAR rank and that the model wasn't able to maximize its learning from such a structure with a significant element of randomness present as a result of the encoding structure of its features. Accordingly, we designed our final model to train on an aggregated summary of each metric by position group. Instead of encoding each metric as its own column for every player on the ice, we encoded the sum of each volume-based metric and mean of each rate-based metric by position group. For example, for goals, we have two columns per team: sum of goals by forwards and sum of goals by defensemen. The model benefited from this new way of feature encoding.

## 3. Results

We trained models aiming to project net expected goals and incremental zone time from a variety of contextual stats as a means for characterizing faceoff performance. Regarding net expected goals, a key realization during training was that it is very sparse, with many non-neutral-zone faceoffs quickly resulting in a zone change (or play stoppage) that prevents any faceoff-attributable expected goal generation. This is relevant because a sparse objective function can be very



challenging for a machine learning model. Specifically, we found that our regression-oriented machine learning models often suffered from too much bias and far too little variance in the form of making a very small range of predictions concentrated around numbers just a hair above zero. We experimented with alternative tuning methods but found this to be a rather unavoidable consequence of net expected goals as an objective function because of its true distribution.

However, we noticed that this sparsity was far less of an issue with incremental zone time. Incremental zone time is unsurprisingly a right-skewed distribution, but, thanks to the natural distinction between puck possession and goal generation, it is not particularly sparse and is much better suited from a distribution-shape perspective to serve as an objective function for model training. Motivated by this finding, we realized that we could engineer a more stable net expected goals projection from our incremental zone time projection. Specifically, we could use the incremental zone time model to make incremental zone time projections and then translate these incremental zone time projections into net expected goal projections by observing how many expected goals were typically accrued per second of incremental faceoff-attributable offensive zone time. These incremental zone time projections (IZT), along with their translations to net expected goal time projections (NXG), are shown in Table 1 for the 2022-2023 season, shown on the next page.

IZT is measured in seconds per faceoff, and NXG is measured in expected goals per faceoff. For example, on the average even-strength, non-neutral-zone faceoff win (either offensive zone or defensive zone), Anze Kopitar is expected to net his team 5.03 seconds of incremental zone time and 0.022 expected goals more than an average center would if deployed in his situations. This holds for both offensive- and defensive- zone faceoffs because the projection is formed both from the incremental value Kopitar is expected to induce for his team off of an offensive zone win and the incremental value Kopitar is expected to seize from the other team after a defensive zone win. Importantly, we have also controlled all projections for situational deployment (offensive zone faceoff or defensive zone faceoff) with the leaderboard reflecting incremental zone time or net expected goals above or below expected given usage. These projections are for games in both the *Evolving Hockey* play-by-play data and the *All Three Zones* data but are out-of-sample relative to the training of the model. We have implemented elements of regression to the mean for projections based on small quantities of faceoffs.

Of note, we have the ability to utilize individualized player-by-player or line-by-line projections of rates of expected goals accrued per second of incremental faceoff-attributable offensive zone time in this translation. However, we intentionally choose to instead use a constant league-average rate for expected goals accrued per second of incremental faceoff-attributable offensive zone time because here in this table we wish to characterize goals gained or lost strictly through performance at the faceoff dot and not through quality of offense played in the aftermath of a faceoff.

**Table 1: Projected Top Incremental Zone Time (IZT) and Net Expected Goals (NXG) Performers for 2022-2023**

Rank	Player	IZT	NXG	Rank	Player	IZT	NXG
1	Anze Kopitar	5.03	0.022	19	Jean-Gabriel Pageau	1.47	0.006
2	Phillip Danault	4.28	0.019	20	Sean Kuraly	1.46	0.006
3	Sebastian Aho	3.92	0.017	21	Dylan Larkin	1.39	0.006
4	Nick Bonino	3.85	0.017	22	Jack Roslovic	1.39	0.006
5	Joe Pavelski	3.37	0.015	23	Bo Horvat	1.36	0.006
6	Tomas Nosek	3.29	0.015	24	Kyle Palmieri	1.33	0.006
7	Connor Dewar	3.09	0.014	25	Pierre-Luc Dubois	1.23	0.005
8	Jesper Kotkaniemi	3.06	0.014	26	Tyler Seguin	1.20	0.005
9	Derek Stepan	2.93	0.013	27	Aleksander Barkov	1.10	0.005
10	Nick Paul	2.86	0.013	28	Sheldon Dries	1.01	0.004
11	Jeff Carter	2.80	0.012	29	Casey Cizikas	0.83	0.004
12	Charlie Coyle	2.71	0.012	30	Christian Dvorak	0.78	0.003
13	Frederick Gaudreau	2.51	0.011	31	David Kampf	0.60	0.003
14	Joel Eriksson Ek	2.35	0.010	32	Yanni Gourde	0.51	0.002
15	Adam Lowry	2.21	0.010	33	Ryan Strome	0.40	0.002
16	Sean Monahan	1.59	0.007	34	Ryan Johansen	0.39	0.002
17	Martin Necas	1.57	0.007	35	Jordan Staal	0.33	0.001
18	Robert Thomas	1.54	0.007	36	Kirby Dach	0.32	0.001

Table 1 sheds light on several trends in faceoff performance worth exploring. First, faceoff performance in terms of our metrics drops off sharply as we move down the leaderboard: the top performers in incremental zone time and net expected goals are projected to be significantly more elite than even other anticipated high performers who grade out in the top ten to twenty. Second, there is a wide range of centers checkering the list of projected top performers: a variety of top-line centers, bottom-six centers, and centers who play up and down the lineup. We consider this to be a vote of confidence for our models, suggesting that they successfully balance the dual mandate of considering expected goal generation after a faceoff in shaping its value while not simply suggesting that players who provide or are surrounded by elite offense are top faceoff performers. This goes back to effective strategies utilized to control for deployment by linemates and game situations.

An interesting result here is the Los Angeles Kings' top two centers taking spots one and two in our projected rankings for last season. Dom Luszczyszyn and Shayna Goldman of *The Athletic* projected the 2022-2023 Kings to be a 93-point team with a 54% chance at making the playoffs [8]. They surprised, finishing the season with 104 points and a playoff berth. It would be overconfident to say that faceoffs propelled the Kings to the playoffs and could fully explain the gap between preseason expectation and end reality, especially given the Kings' dependence on many maturing young players. However, it is reasonable to posit that supreme faceoff performance may be at least partly responsible for the turnaround.

Delving deeper than initial projected performance to showcase the potential for faceoffs to serve as an effective mechanism for goal creation, we explore how beneficial winning various numbers of



additional faceoffs can be, again using our top performers as examples. How many more non-neutral-zone faceoffs it is reasonable to think a team or player could win per game is up for debate and remains an open question. We show the goal creation implications for a variety of amounts of additional faceoff wins in Table 2. Again, using Kopitar for example, we would project the Kings to net an additional 1.82 goals in offense over the course of a full season beyond his existing projected output from Table 1 if he started winning one more faceoff per game. We would expect 3.65 incremental goals in offense over the course of a full season if he started winning two more faceoffs per game (and so on). These figures assume that the player in question plays a full 82-game season.

**Table 2: Projected Seasonal Incremental Net Expected Goals by Count of Incremental Faceoff Wins and Player per Game**

Player	+1	+2	+3	+4	+5	+6
Anze Kopitar	1.82	3.65	5.47	7.29	9.12	10.94
Phillip Danault	1.55	3.10	4.66	6.21	7.76	9.31
Sebastian Aho	1.42	2.84	4.26	5.68	7.10	8.53
Nick Bonino	1.40	2.79	4.19	5.58	6.98	8.38
Joe Pavelski	1.22	2.44	3.66	4.88	6.10	7.32
Tomas Nosek	1.19	2.38	3.58	4.77	5.96	7.15
Connor Dewar	1.12	2.24	3.36	4.48	5.59	6.71
Jesper Kotkaniemi	1.11	2.22	3.32	4.43	5.54	6.65
Derek Stepan	1.06	2.12	3.18	4.24	5.31	6.37
Nick Paul	1.04	2.07	3.11	4.15	5.18	6.22
Jeff Carter	1.02	2.03	3.05	4.06	5.08	6.09
Charlie Coyle	0.98	1.97	2.95	3.93	4.91	5.90
Frederick Gaudreau	0.91	1.82	2.72	3.63	4.54	5.45
Joel Eriksson Ek	0.85	1.70	2.55	3.40	4.25	5.10
Adam Lowry	0.80	1.60	2.40	3.20	4.00	4.80
Sean Monahan	0.58	1.16	1.73	2.31	2.89	3.47
Martin Necas	0.57	1.14	1.71	2.28	2.85	3.42
Robert Thomas	0.56	1.11	1.67	2.23	2.79	3.34
Jean-Gabriel Pageau	0.53	1.06	1.59	2.12	2.65	3.18
Sean Kuraly	0.53	1.06	1.59	2.11	2.64	3.17
Dylan Larkin	0.50	1.00	1.51	2.01	2.51	3.01
Jack Roslovic	0.50	1.00	1.51	2.01	2.51	3.01
Bo Horvat	0.49	0.99	1.48	1.97	2.47	2.96
Kyle Palmieri	0.48	0.96	1.44	1.92	2.41	2.89
Pierre-Luc Dubois	0.45	0.89	1.34	1.79	2.23	2.68
Tyler Seguin	0.44	0.87	1.31	1.74	2.18	2.61
Aleksander Barkov	0.40	0.80	1.20	1.60	2.00	2.40
Sheldon Dries	0.37	0.73	1.10	1.47	1.84	2.20
Casey Cizikas	0.30	0.60	0.90	1.20	1.51	1.81
Christian Dvorak	0.28	0.57	0.85	1.13	1.41	1.70
David Kampf	0.22	0.43	0.65	0.86	1.08	1.29
Yanni Gourde	0.18	0.37	0.55	0.73	0.92	1.10
Ryan Strome	0.15	0.29	0.44	0.58	0.73	0.87
Ryan Johansen	0.14	0.28	0.42	0.56	0.70	0.84
Jordan Staal	0.12	0.24	0.36	0.48	0.60	0.72
Kirby Dach	0.12	0.23	0.35	0.46	0.58	0.70

In relation to Table 2, it is worth emphasizing that the reflected upside needs to be tempered by an understanding that professional NHL centers are almost certainly maximizing their faceoff win percentages already and that it is likely unreasonable to argue that there is an adjustment or increase in practice that could yield a meaningful increase. Moreover, the reflected numbers assume a linear benefit for each incremental non-neutral-zone faceoff win. However, we still include the table to demonstrate how even incremental faceoff impacts can have powerful effects. On a more actionable note, one way to in effect achieve an incremental increase in faceoffs won is for coaches to either play standout faceoff performers more often or specifically substitute them in more frequently during pre-faceoff line changes after stoppages in play.

## 4. Applications

### 4.1. Integrating Faceoff Performance into Classic Stats

With our projections of goals gained or lost at the faceoff dot for each center, we have the opportunity to integrate faceoff performance into traditional box score stats beyond the limited information provided by faceoff win percentage. We first offer a positive or negative measure of goals gained or lost through projected faceoff performance. We then show a modified sample table where goal totals are now complemented by net faceoff impact in goals. Obviously not all faceoff performance directly affects these scoring totals (it is likely distributed across nearly every statistic in the box score since nearly all depend on puck possession) but visualizing the impact of faceoff performance alongside traditional box-score stats serves as a useful way to intuitively demonstrate the power of faceoffs in influencing high-level conceptions of player talents. We do this in Table 3 (shown on the next page), which makes use of actual 2022-2023 season data from *Hockey Reference* for the selected players <sup>[9]</sup>. Note that the player rankings are different than before because the NXG from Faceoffs column in the table considers how many non-neutral-zone faceoffs the player won and their deployment in the 2021-2022 season as context applied in relation to their projected net expected goals per non-neutral-zone faceoff won for the 2022-2023 season.

Strictly speaking in terms of colloquial heuristics and general hockey fan rhetoric here, faceoffs have the ability to make a coveted 30-goal center in effect be more like a 20- to 25-goal scorer. In reverse, a 15-goal bottom-six center can suddenly approach 20- to 25-goals in effect (often in less ice time). We aim to introduce these faceoff-driven considerations into the hockey lingo and make faceoffs—and their specific impact in terms of goals—a larger part of general hockey discussions.

### 4.2. Special Teams

One of the most interesting implications of our research is that faceoffs could yield a modern reconstruction of power plays. Specifically, we find that the extensive time it takes a team on a power play to successfully re-enter the offensive zone after an offensive zone faceoff loss stands out sharply in relation to the foregone highly valuable faceoff-attributable, power-play offensive zone time that could have been had if a faceoff win were achieved.

This realization fuels a theory that faceoffs may indeed be more valuable on power plays (and accordingly also more valuable on penalty kills because a defensive zone faceoff win on a penalty kill steals would-be valuable offense from a power play). The idea that faceoffs are important on power plays and penalty kills is not necessarily new in its own right, but we consider whether

power plays could be made more effective if faceoff specialists—even faceoff specialists who are less offensively gifted than traditional power play participants—were given power play time to be a worthy topic of future research.

**Table 3: Visualizing Faceoff Impact in relation to Traditional Box-Score Stats**

Player	Unadjusted Traditional Stats			Faceoff Impact
	Goals	Points	+/-	NXG from Faceoffs
Anze Kopitar	28	74	20	12.03
Phillip Danault	18	54	-8	8.03
Sebastian Aho	36	67	8	6.45
Nick Bonino	10	19	-5	5.48
Jeff Carter	13	29	-16	5.41
Joel Eriksson Ek	23	61	4	4.95
Charlie Coyle	16	45	29	4.63
Joe Pavelski	28	77	42	4.00
Adam Lowry	13	36	4	3.64
Frederick Gaudreau	19	38	10	3.48
Tomas Nosek	7	18	9	3.43
Bo Horvat	38	70	-1	3.37
Jean-Gabriel Pageau	13	40	-2	3.24
Nick Paul	17	32	11	3.22
Robert Thomas	18	65	-8	2.97
Sean Kuraly	11	20	-28	2.64
Aleksander Barkov	23	78	10	2.39
Dylan Larkin	32	79	-7	2.38
Sean Monahan	6	17	-5	2.11
Pierre-Luc Dubois	27	63	5	2.09
Derek Stepan	5	11	8	2.04
Tyler Seguin	21	50	3	1.69
Jack Roslovic	11	44	-14	1.67
Christian Dvorak	10	28	-12	1.36
Jesper Kotkaniemi	18	43	10	1.28
Casey Cizikas	6	21	0	1.13
David Kampf	7	27	6	1.06
Ryan Johansen	12	28	-13	0.80
Connor Dewar	6	18	-5	0.71
Jordan Staal	17	34	7	0.64
Yanni Gourde	14	48	23	0.61
Ryan Strome	15	41	-30	0.43
Kirby Dach	14	38	-2	0.21
Martin Necas	28	71	5	0.19
Kyle Palmieri	16	33	13	0.12
Sheldon Dries	11	17	-9	0.01

### 4.3. Offensive Defenseman Overhype?

We further extend our theory that standout faceoff performers may be undervalued and underutilized on power plays to speculate whether offensive defensemen may pose a market inefficiency. Specifically, the top-20 defensemen in points last season accrued on average 37% of their points on the power play and had a median average annual value of \$7.925 million. It's unlikely a standout faceoff performer could match the full offensive impact of an elite offensive defenseman on the power play, but the gap between the annual cost of elite offensive defensemen and many faceoff specialists who often play for salaries near the league minimum (about \$0.75 million) is striking. We leave it as an open question whether faceoff standouts, especially those with passable offensive ability, could efficiently replace some elite offensive defensemen on power plays at a fraction of the cost in a way that delivers significant surplus value.

## 5. Conclusion

Big picture, we concur with and complement the Czuzoj-Shulman et al conclusion that all faceoff wins are not created equal and find estimates slightly more bullish than Schuckers et al regarding the value of the average faceoff using our updated methodologies. We push back against the common pure conclusion that faceoffs are unimportant but concur that their importance comes down to opinions on just how many more faceoffs a team can incrementally win within the scope of reasonable personnel changes, situational usage adjustments, and coaching. We qualify existing arguments to argue that each season features an existent but very small subset of players who net their teams enough expected goals through performance at the faceoff dot to conceivably win meaningfully more games. Our research suggests that the core question is not necessarily whether faceoffs can be important but rather whether it is possible to predict and affordably acquire the small minority of players that will be standout faceoff performers. We suggest that elite faceoff performers can eclipse over five goals per season in projected effect and that, with strong personnel decisions, faceoffs can win meaningful amounts of games over the course of a full season.

## References

- [1] Maxwell, Scott. "How important are faceoffs in hockey?" *Leafs Nation*, 23 November 2021, <https://theleafsnation.com/news/the-leafs-first-line-left-wing-spot-still-seems-to-leave-a-lot-to-be-desired>.
- [2] Sporer, Evan. "Illustrated Review: Searching for a link between face-off wins and longterm success." *Sports Illustrated*, 2 March 2017, <https://www.si.com/nhl/2017/03/03/illustrated-review-importance-nhl-faceoff>.
- [3] Luszczyszyn, Dom. "Why Faceoffs Aren't as Important as They're Made Out to Be." *The Hockey News*, 17 December 2015, <https://thehockeynews.com/news/why-faceoffs-arent-as-important-as-theyre-made-out-to-be>.
- [4] Czuzoj-Shulman, Nick et al. "Winning Isn't Everything: A contextual analysis of hockey face-offs." *MIT Sloan Sports Analytics Conference*, 1 March 2019, <https://www.sloansportsconference.com/research-papers/winning-isnt-everything-a-contextual-analysis-of-hockey-face-offs>.
- [5] Schuckers, M., Pasquali T., and Curro J. "An Analysis of NHL Faceoffs." *Statistical Sports Consulting*, <https://www.statsportsconsulting.com/wp-content/uploads/FaceoffAnalysis12-12.pdf>.
- [6] "Evolving-Hockey." *Evolving Hockey*, <https://evolving-hockey.com/>. Accessed 30 August 2023.
- [7] Sznajder, Corey. "All Three Zones Project," *All Three Zones*, <https://www.allthreezones.com/>. Accessed 23 February 2023.
- [8] Luszczyszyn, Dom and Goldman, Shayna. "Los Angeles Kings 2022-23 season preview: Playoff chances, point projections, roster rankings." *The Athletic*, <https://theathletic.com/3632142/2022/09/27/kings-2022-2023-season-preview/>. Accessed 30 November 2023.
- [9] "2022-23 NHL Skater Statistics." *Hockey Reference*, [https://www.hockey-reference.com/leagues/NHL\\_2023\\_skaters.html](https://www.hockey-reference.com/leagues/NHL_2023_skaters.html). Accessed 30 November 2023.