The Impact of Puck Possession and Location on Ice Hockey Strategy

Andrew C. Thomas*

*Harvard University, athomas@stat.harvard.edu

The Impact of Puck Possession and Location on Ice Hockey Strategy*

Andrew C. Thomas

Abstract

I create a state space within the game of ice hockey by noting which team has possession, and in what location of the rink the puck is located. This space is used to model the game as a semi-Markov process, as data from a series of games in 2004-2005 NCAA play suggest that the system cannot be modeled as a continuous time Markov process. The model is then used to determine the average number of goals scored by a team as a function of the starting state. These scoring probabilities are used to demonstrate the effectiveness of several commonly used tactics.

Introduction

Strategy in sports is routinely described in terms of offence and defence: scoring points for your team, and preventing the opposition from scoring. Separating and isolating the two is critical to determining strategic determinations. In games like ice hockey, offence and defence become most difficult to separate, as control of the play is often difficult to determine. It is therefore challenging to understand the distinct effects on offence and defence as they relate to particular strategies. Ice hockey, arguably the fastest field game, is an excellent candidate for analysis.

Lindsey [1] categorized the expected number of runs scored in baseball given the number and placement of men on base and the number of outs in an inning. He then compiled a 3-by-8 table of expected runs scored in an inning given each of 24 game situations. Bukiet et al. [2] generalized this model by considering certain well-recorded outcomes - base hits, strikeouts, and so forth - and using this information to estimate a transition matrix.

For association football, Hirotsu and Wright [3] modelled the progress of play as a continuous time Markov process with four states: when each team is in possession of the ball, and when each team scores a goal.

Ice hockey has been relatively unexplored compared to other sports; analyses have included the performance of teams on the power play [4] and the best timing to substitute a goaltender for an extra skater in a bid to tie the score late in a game [5]. Amateur research has yielded other information, such as the relationship between the probability distribution of scoring a goal based on the shooter's distance from the net.

The goals of this paper are as follows: to demonstrate that a set of states can adequately describe the flow of a hockey game, so that the Markov transition property holds; to determine whether the game can be modelled as a continuous time Markov process, and if not, to determine what model will be appropriate; and to determine the (time-dependent) probability of scoring or allowing a goal given that, at a particular moment, the game is in a particular state. Given these probabilities, I analyze several strategic decisions common to the game and determine their effectiveness in scoring or allowing goals.

Of the limits imposed by this model, easily the most constraining is that the resolution of our game in motion is described by a mere 19 states - less than baseball's 25, and far less than, as a contrived example, the 120 "states" one could obtain by permuting 5 players between 5 skating positions. While further research, and more detailed game data, could resolve these issues, this model is adequate to answer basic questions of strategy: given a choice of one or more states, which will yield the highest probability of scoring a goal, the lowest probability of conceding a goal, or both? And what will have more impact on a team's ability to score goals: whether they control the puck, or whether they are close to the opposing team's goal?

1   State Space

In order to keep track of the progress of play, I label the teams by A and B (whether or not either team has an advantage at home. Each team can have possession of the puck at any point, and is not considered to surrender it until the other team has made puck contact. This choice also means that the completion of a pass between location zones marks the

change of state, rather than the beginning.

Considering location, the blue lines (see Figure 1) divide the rink into three zones - a neutral zone, and two defensive zones containing the goals. The offside rule in hockey states that a team cannot enter their offensive zone before the puck, or before a player carrying the puck.

Unlike other sports, when play stops due to a minor infraction, penalty, or goal, a faceoff is used to resume the game. An on-ice official conducts a faceoff by dropping the puck between representatives of each team at a location determined by the appropriate rule. The vast majority of faceoffs take place on a faceoff dot (as seen on Figure 2), and so we can group all faceoffs into five zones.

I will also consider goals by each team to be their own states.

These 13 states would appear sufficient to describe the structure of the game. However, it is clear by inspection of the data that the Markov property - the probability of transition to a new state being dependent on only the current state - does not hold for certain subsets of the data. To remedy this, I separate the following three distinct transitions, as viewed on Figure 3:

- Retreat, when a team with neutral possession takes the puck into their defensive zone (the counterpart will be referred to as "Pursuit");

- Defensive Takeaway (Offensive Giveaway)[1], when a team with offensive possession yields the puck to the defending team (possibly forfeiting a scoring chance);

- Offensive Takeaway (Defensive Giveaway), when a team with defensive possession yields the puck to the attacking team (possibly giving the attacking team a scoring chance.

All together, the following state space is constructed, in which transitions between states follow the Markov property (while not considering the time duration of these states.) I label each state with a number from the perspective of Team A, and indicate that each state has a counterpart for the perspective of Team B.

---

[1] For lack of better terms, I consider a giveaway and a takeaway to mean the same process, wherein possession of the puck is transferred between teams, while defensive/offensive refers to the team doing the giving or taking. While it may be argued that these terms mean feats of different skill levels (low and high respectively) the model can't tell the difference anyway.
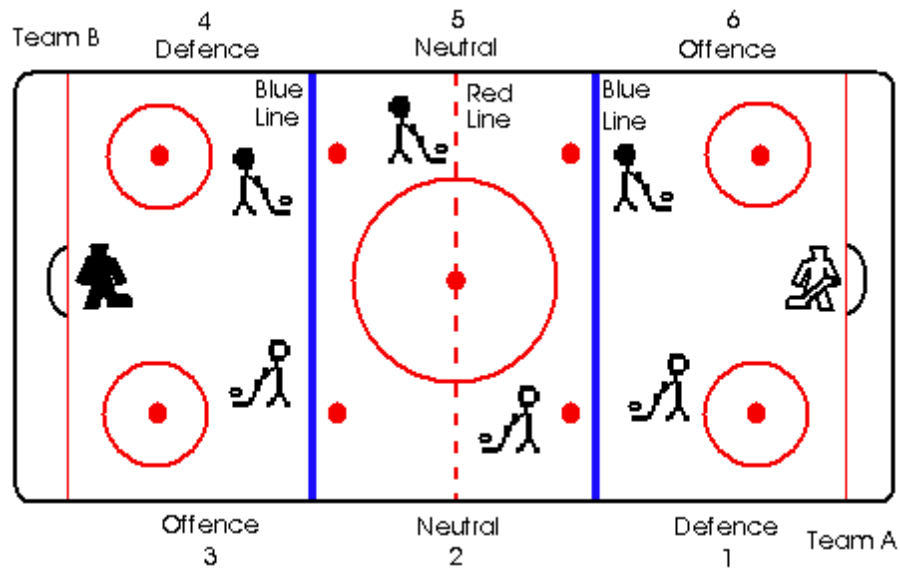
Figure 1: Illustrations of states 1 through 6, describing defensive, neutral and offensive possession for each team A and B.

| State i	State $c(i)$	Team with puck	Location
1	4	A	Defensive possession
2	5	A	Neutral possession
3	6	A	Offensive possession
4	1	B	Defensive possession
5	2	B	Neutral possession
6	3	B	Offensive possession
7	11	Faceoff	A defensive zone
8	10	Faceoff	A blue line
9	9	Faceoff	Center red line
10	8	Faceoff	B blue line
11	7	Faceoff	B defensive zone
12	13	Goal	A goal (B zone)
13	12	Goal	B goal (A zone)
14	17	B	Offensive Takeaway
15	18	A	Retreat
16	19	B	Defensive Takeaway
17	14	A	Offensive Takeaway
18	15	B	Retreat
19	16	A	Defensive Takeaway

Figure 2: Illustrations of states 7 through 13, describing faceoffs and goals for each team A and B.


Figure 3: Illustrations of states 14 through 19, describing several strategic outcomes. Removing these from the basic possession states allows the time-independent model of the system to be Markovian.

Changes In Manpower

For egregious rule violations, players are removed from the ice without replacement for a period of time (typically 2 or 5 minutes.) There are significant changes in strategy by the team with the advantage in manpower (which enjoys a "power play") and the team with the disadvantage (which is shorthanded.) In response, I separate these events into two groups; when the teams are of even strength, and when one team has a manpower edge (though which team that has this edge is also tracked.) This paper will only deal with events which take place while teams are at even strength.[2]

2   Record-Keeping and Data

Records kept by professional and collegiate associations have historically been limited to a small selection of discrete events; goals, shots on goal, and penalties (which can increase the likelihood of scoring.) Recently, the National Hockey League implemented a system which tallies several new events, including every player's presence or absence on the ice throughout the game, locations for shots and stoppages in play, and faceoff win/loss information. [3]

But aside from the verbal information in the radio play-by-play, few records are kept which track the actual progress of play while it happens, and none exist in the public domain. Since radio broadcasts lack precision timing, measurement systems currently in use do not rigorously track puck and player movement.

Due to its ease of access and a high quality of play, I use a set of 18 games involving the Harvard Men's Varsity Hockey team during the 2004-2005 season. Away games were obtained on videotape and examined after the fact.[4] During these 18 games, there were 677 minutes of even-strength play, or roughly 11 full hours/games worth. All teams scored a total of 41 even strength goals, for an average of 1.2 even strength goals per 20 minutes of play.

Because events were collected in real time, standard scorekeeping tools (a pad of paper and a pencil) are impractical. Very high-tech solutions, using GPS technology or optical tracking, have been explored in the private sector, but these are beyond the scope of this level of analysis. I therefore developed a simple data-gathering device for use with a laptop computer.

Table 1 gives a small sample of events recorded in sequence. It contains all information about the time duration of states, the transition made at the end of each duration, the time of the game when each transition was made, and the score of the game at the time.

A Note on Symmetry

Each state has a symmetric counterpart. For example, if team A has the puck in the neutral zone (state 2), team B will be in state 5. In general, if team A is in state i, team B will be in

---

[2]The effect of a penalty call on even-strength play - an increased probability of scoring by the nonoffending team before play stops - has not been taken into consideration, simply because there are too many parameters unaccounted for in the scoring data to attempt a reasonable estimate.

[3]The system once recorded giveaways and takeaways, but the NHL ceased this several years ago for various reasons - possibly given the difficulty in distinguishing the two.

[4]Thanks are given to the Harvard Men's Hockey staff, in particular video coach Joe Heydenburg, for allowing me access to video tapes of their away games.

Table 1: A sample of events and transitions recorded during the course of the season. Time intervals are recorded in seconds. See text for a further explanation. Note that the numbering scheme is representative of our actual state space and is assembled after the fact.

State i	State i+1	Δt	$t_{in-game}$	Score(A)	Score(B)
4	16	3.1	46	0	0
16	17	3.4	49.4	0	0
17	5	7.3	56.7	0	0
5	6	1.9	58.6	0	0
6	13	2.2	60.8	0	0
13	5	0	60.8	0	1
5	6	6.3	67.1	0	1
6	19	0.9	68	0	1

state $c(i)$ (or, verbally, switch offence/defence, possession/non-possession, retreat/pursuit and giveaway/takeaway for the opposite description.)

In particular, I wish to address the issue of balance between teams A and B. Ideally I would like these two teams to be evenly matched; as an investigatory tool, I can consider each event twice by also finding and noting its symmetric counterpart. For example, here is the counterpart of a sample transition record is shown in Table **??**.

	State i	State i+1	Δt	$t_{in-game}$	Score(A)	Score(B)
Original	4	5	7.3	56.7	0	2
Counterpart	1	2	7.3	56.7	2	0

If I were to consider Harvard as Team A, and Harvard's opponents as Team B, across all games, the results would not be symmetric. Harvard tended to outclass many of their opponents, as they finished second in their conference of 12 teams and outscored their opponents in even-strength situations 26-15. Therefore, I consider Team A and Team B to be an amalgam of all events by both teams using the counterpart method.

Continuous Markov or Semi-Markov Process?

The object of constructing a model is to determine the probabilities of scoring or allowing goals under particular conditions. One such condition is the passage of a particular amount of time.

The simplest model for a continuous time, discrete state process such as hockey would be a Markov process. For $X(t)$, the state of the game at time t, this would require that the game be memoryless - that the transition probabilities to other states would be independent of how long this state had been held. This requires that a state holding time be exponentially distributed.

Figure 4 shows the distribution of the holding time for various states conditioned on the next state to be reached. It is clear that these states are not exponentially distributed; moreover, the transition probability conditioned on time is not constant, as can be seen in Figure 5. A continuous time Markov process cannot be used to model this system.

Consider instead a discrete state, continuous time **semi-Markov** process. In this case, given that the system begins in state $I = i$ and moves to state $J = j$ after time T, that $T|i, j$
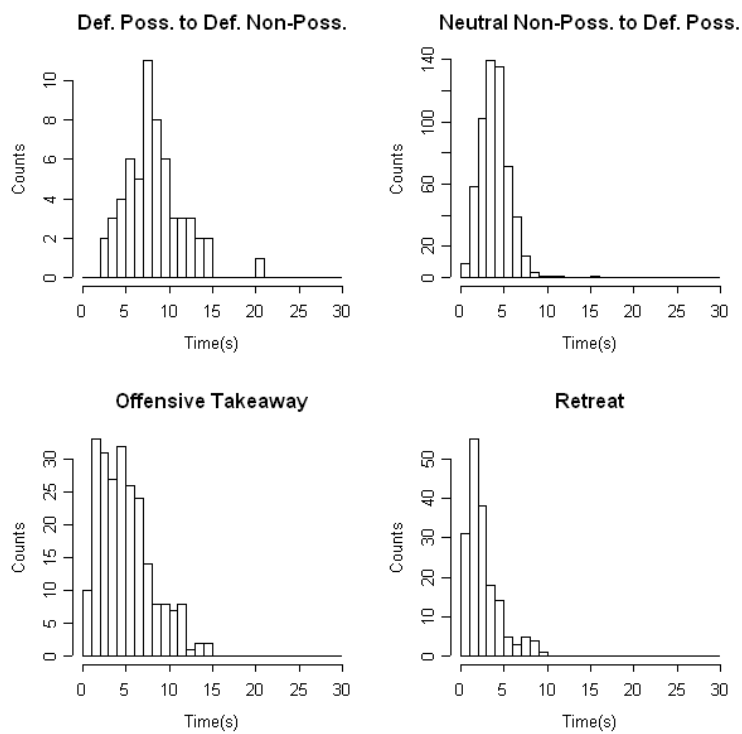
Figure 4: State holding times, given the final state. Since their shape is not exponential, this strongly suggests that the system is not Markovian in continuous time.
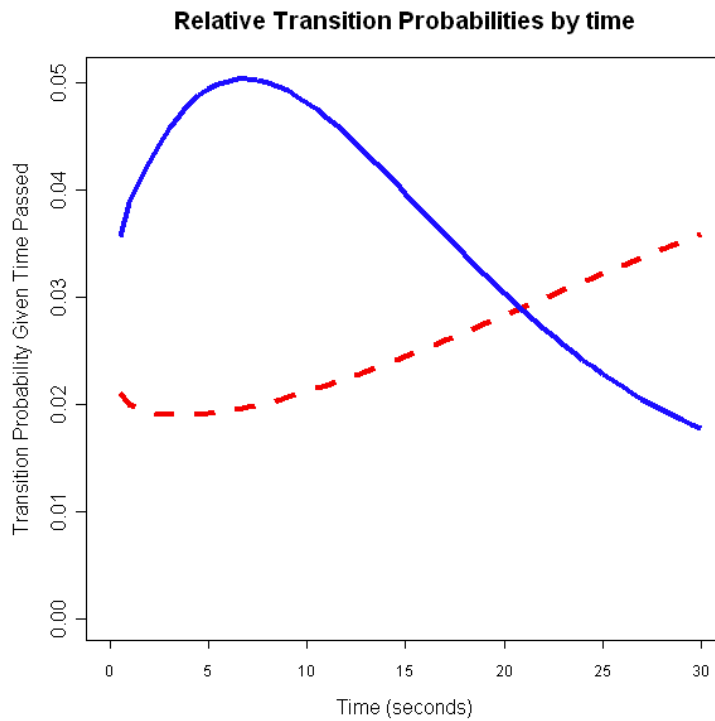
Figure 5: Comparing the relative transition probabilities from offensive possession to neutral possession (solid) and scoring a goal (dotted) as a function of time elapsed in offensive possession.

is drawn from a distribution $Y(i,j)$ with mean $\mu_{i,j}$ [6]. As a quick inspection, transitions involving greater distance tend to have higher mean times and longer left tails. [5]

I considered several different distributions to model the system (see Appendix for further information); of these, the two parameter Gamma distribution was chosen both due to its robustness in fitting the model, and due to its similarity in form to the exponential distribution. Parameter estimates for each distribution $T|i,j$ were determined by the method of moments.

Compiling The Transition Matrix

Since each transition must end up in one state exclusive of all others (and of the starting state), a multinomial model is used to estimate transition probabilities independent of time. Because this model has a conjugate prior distribution of the Dirichlet form, and because subsequent operations may prove to be mathematically complicated and unsuitable for frequentist methods, I choose to use Bayesian inference with a multinomial/Dirichlet model. This approach will also give reasonable confidence intervals for each calculated parameter, some of which will prove significant.

Let D be the data gathered throughout the season, given whatever conditions we wish to impose. (In most cases, the only condition is that the teams are at even strength.) The following steps are used to determine $P(J|I,D)$, the transition probabilities independent of time:

- The sample distribution, $P(D|I,J)$, is multinomial.

- The prior distribution, $P(J|I)$, is Dirichlet (the generalized Beta), where the parameters are 0 for states representing goals and 1 for instances with all other events which can occur under the model.[6] I reduce the dimensionality of the model by removing all "impossible" dimensions, such as the scoring of goals from any state other than possession in the offensive zone and transitions between pairs of faceoff states.

- The posterior distribution, $P(J|I,D)$, is simply the vector sum of the sample and prior coefficients, by conjugacy between the multinomial and Dirichlet distributions. Any zeroes that appear in this row indicate impossible transitions and thus represent a reduction in dimensionality.

Each $P(J|I = i, D)$ represents the i^{th} row of the transition matrix $\mathbf{P}$. We determine our desired quantities by drawing all rows simultaneously and composing $\mathbf{P}$.

Calculation Methods for Desired Quantities

For the sake of this analysis, it suffices to use discrete time intervals for our calculations. Since the estimated measurement error is roughly 0.3 seconds in either direction, an interval on that order of magnitude will be reasonable.

---

[5]This suggests that a transition may be considered as the addition of a "waiting time", which may in fact be memoryless, and a "travel time" which is bounded. I was unable to verify this model; the Gamma distribution, which may be considered a sum of exponential random variables, may contain this very property.

[6]Since the Dirichlet distribution cannot technically take a parameter value of zero, this is an improper prior. However, dimensional reduction ensures that the posterior distribution is indeed proper.

Given this condition, and a known starting state, dynamic programming can be used to determine $P(J|I, T_{elapsed})$ from a joint probability distribution $P(J, T_{holding}|I)$ to solve transition probabilities for a discrete time, discrete state semi-Markov process, noting that multiple transitions may occur within $T_{elapsed}$. Since the probability $P(J, T_{holding}|I)$ is known, we can solve a recursive set of equations $P(J, T_k) = \sum_{n=1}^{t} \sum_{i=1}^{19} P(J, T_k|I = i)P(I = i, T_{k-n})$.

3   Analysis

System Behaviour With Respect to Time

The analysis is conducted on the premise that the system begins in each state I just after a transition has been made.

In the first 5 seconds, the probability of scoring is virtually zero except for two starting states - those leading directly to goals, offensive possession and defensive turnover (see Figure 6.) This confirms a fact we knew to be true by construction: in the extreme short term, possession of the puck in the offensive zone is necessary to score goals.

At 20 seconds, the expected number of transitions grows to such an extent that scoring a goal beginning in any state is now feasible. At 40 seconds, the scoring rates beginning in each state are nearly identical, suggesting that the system has mixed, and the memory of the original starting state is disregarded (unless, of course, a goal was scored.)

Selected scoring probabilities as a function of time are given in the appendix.

Strategic Comparison

When analyzing any type of strategy, the key question that must be answered is, given a choice of alternatives, one choice will produce a superior result than all others, either with certainty or with high probability. While one outcome at a time is observed - and, overwhelmingly, the outcome is "no goal" - the probabilistic interpretation will yield a measurable answer to rough-scale questions of strategy.

First, the exact quantity of interest must be determined. As can be seen from Figure 6, the scoring probabilities rise at roughly the same rate after 40 seconds. It can also be seen that this represents the maximum difference between any two states.

Indeed, since the variances in these estimates after mixing has been achieved, I will use points of minimum variance - scoring probabilities at 40 seconds - as the desired outcomes.

When in trouble, dump the puck or retreat?

When a team is tired and requires a change of players, two maneuvers are commonly used: the puck is surrendered to their opponents and close to their goal (commonly known as the "dump-in"), or the puck is brought back into the team's own defensive zone and control is maintained (or the "retreat"). On average, which strategy is superior?

Graphically, the solution can be seen from Figure 7, assuming that all strategies are successfully executed (which, at this level of play, is reasonable.) The scoring probability for "retreat" is higher than its counterpart, "pursuit", and the scoring probability for "dump-in", offensive non-possession, is higher than that for defensive possession. However, a team

Expected Goals Scored, 5 Seconds
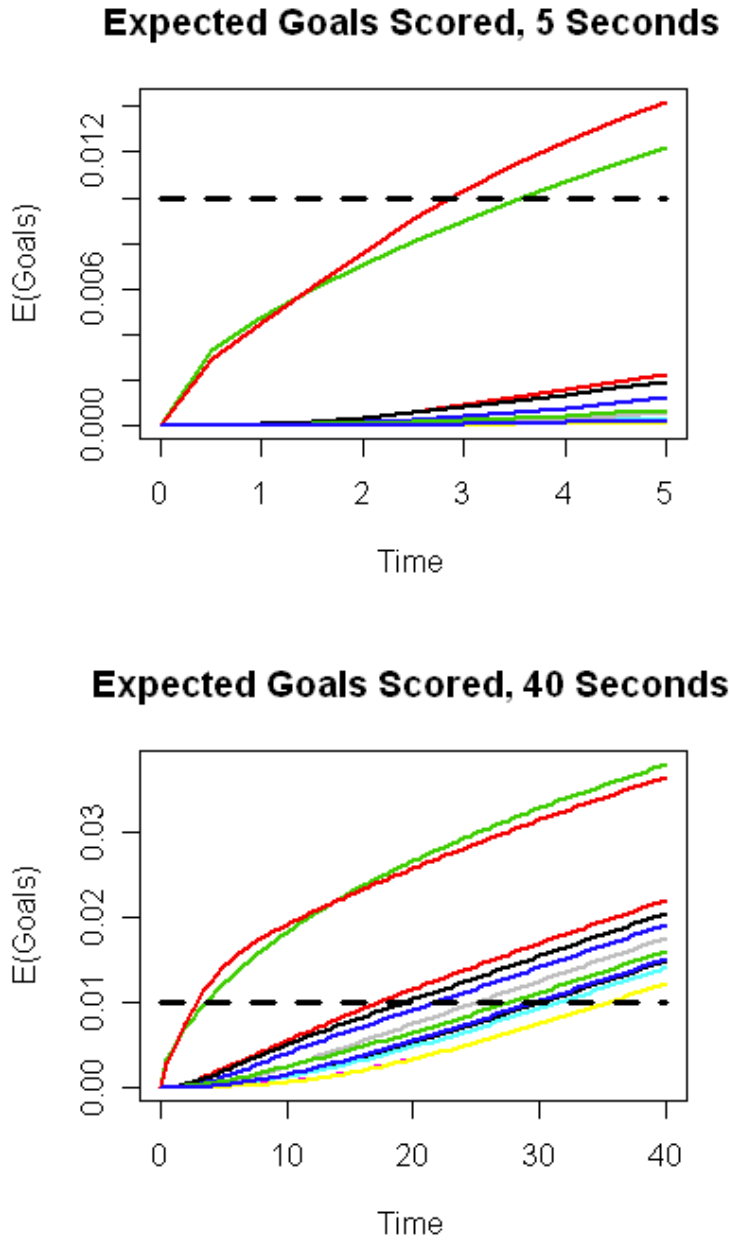

Expected Goals Scored, 40 Seconds

Figure 6: Comparing the scoring probabilities beginning in each state as a function of time. In the top diagram, two states - Offensive Possession and Offensive Takeaway - clearly dominate in the first five seconds. In the bottom diagram, the probabilities all rise uniformly after 40 seconds, suggesting the system has mixed.
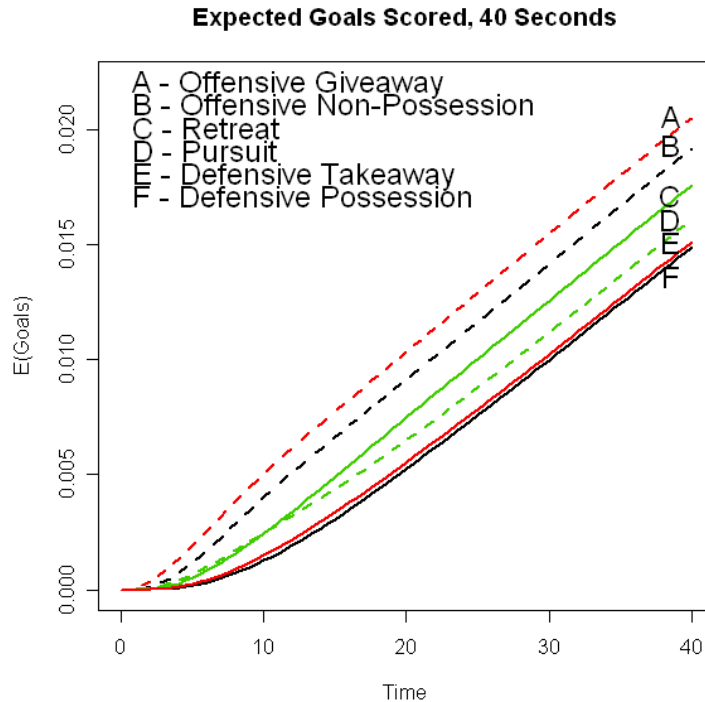
Figure 7: Comparing the scoring probabilities beginning in several states as a function of time.

employing the dump-in strategy has a greater probability of scoring, and a lesser probability of being scored upon, than a team retreating.

When on offense, dump and chase or carry in?

When a team is in the neutral zone, two strategies common to enable the scoring of goals are to send the puck into the offensive zone, then chase after it (also known as the "dump and chase") or to attempt to carry the puck in directly.

This problem is more difficult to identify than in the previous case because of a rather large piece of missing data: the intent of the puckhandling player is not encoded in the model, nor is his attempt. The simplest estimate of these outcomes will be taken by conditioning on those moves that each strategy may produce: the former forbids a neutral-zone giveaway and blue-line faceoff, the latter an offensive-zone giveaway. By simulation, the probabilities of scoring can then be determined for each team. We have two desired quantities: offensive bonus (increased probability of scoring) and defensive bonus (decreased probability of allowing a goal), as well as the sum of the two for overall effect.

As can be seen in Figure 8, the offensive bonus provided by "carry-in" over "dump and chase" is about equal to the defensive penalty, so that the net effect on a team's scoring
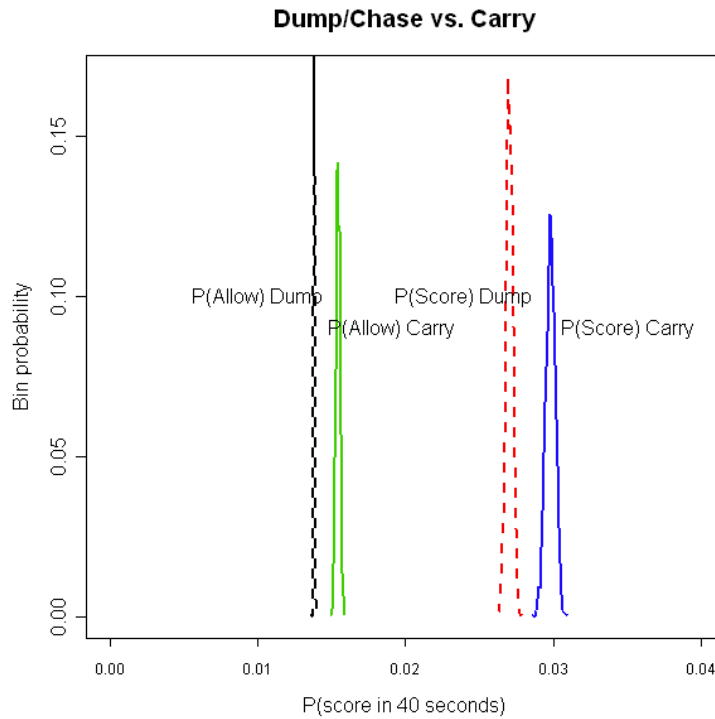
Figure 8: Comparing "dump and chase" vs. "carry in" strategies. The former yields a defensive bonus, the latter an offensive bonus; however, neither strategy is preferable in terms of total scoring advantage.

is zero. However, it is clear that "dump and chase" is a superior defensive strategy, as "carry-in" is superior for the offence.

With defensive possession, clear the puck or press on?

Much the same problem as in the last section: when a team has defensive possession, should they just attempt to remove the puck from their zone ("clearing" the puck) or should they make every attempt to maintain possession ("pressing" the attack)? In this case, the former disallows a defensive giveaway; the latter only allows such a giveaway or a transition to the neutral zone with possession intact.

Since we have two states to consider, defensive possession and defensive takeaway, I consider each case separately in Figure 9. Note that in these cases, allowing a goal is more likely than scoring one. Once again, the total scoring advantage is roughly the same; once again, there is a strategy that increases the likelihood of scoring at the expense of prevention, and one that decreases vulnerability at the expense of scoring.
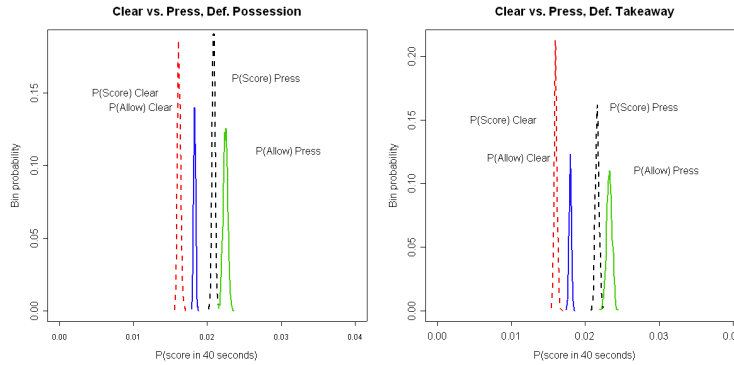
Figure 9: Comparing "clear" vs. "press" strategies while on defence. Neither strategy is generally superior to the other, but "clear" is better defensively while "press" is better offensively.

The value of a faceoff

I have identified five areas where faceoffs tend to occur. The model allows us to identify the advantage of winning a faceoff by comparing the scoring probabilities of each team conditioned on a faceoff win or loss.

Zone	A Wins			B Wins			Victory Bonus
	P(A)	P(B)	P(B)-P(A)	P(A)	P(B)	P(B)-P(A)	
A Def.	0.0149	0.0191	0.0042	0.0123	0.0378	0.0255	0.0213
A Blue	0.0163	0.0181	0.0018	0.0140	0.0228	0.0088	0.0070
Neutral	0.0218	0.0143	-0.0075	0.0143	0.0218	0.0075	0.0150

The short-term gain for faceoffs depends on the starting point. Defensive-zone faceoffs are worth the most, since a loss may more quickly result in a goal; blue line faceoffs are worth the least because the attacking team gains either possession or location but not both. This demonstrates that the faceoff is as important as any other aspect of the game, though far less important when held at a blue line.

4   Conclusions and Future Developments

This model, rough as it may be, is capable of producing strategic determinations by estimating the differences in expected goals scored under multiple outcomes. In particular, noting that giving up puck control in order to gain a territorial advantage is beneficial both offensively and defensively may explain the prominence of location-based defense throughout professional hockey.[7]

However, all these conclusions rest on the assumption that Harvard and their opponents play with a roughly similar level of ability (a legitimate assumption, given the talent base

---

[7]New rules in the National Hockey League appear to have produced an increase in scoring, but their full impact has yet to be fully measured.

from which the schools may draw.) I therefore propose tracking an entire intraleague season within a high-quality league, such as the twelve-team ECAC. Since each team plays each opponent twice, once at home and once away, and since the schedule is determined by outside factors and is therefore essentially randomized, an entire season has the properties of a well-designed experiment.

Other possible improvements include:

- Finer-scale location. While this approach may not be appropriate for dividing the playing surface into, say, 100 units rather than three, being able to determine the likelihood of scoring or allowing a goal starting with offensive possession may prove to show at what point a team should effectively stop shooting, or how much ground to surrender, in order to prevent the opposition from taking the puck away and scoring themselves. Currently available data sets may hold the necessary information to calculate this.

- The "goal value" of taking a penalty to prevent an opponent's goal. This may depend on a similar data set as required for the previous problem.

- Tracking line changes. Unlike in basketball or baseball, where substitutions are relatively uncommon events, an entire team of skaters is typically replaced in less than a minute, and substitutions are often made while play continues. Being able to quantify changes in behavior as a function of shift fatigue would prove interesting, as it might suggest an ideal shift length through a game. This would require tracking shift changes as they happen. This is especially important given the inability of the model to accurately predict offence after a retreat, since the suspected cause of increased goal scoring is catching the other team unprepared during a line change.

- Comparative data from other leagues. Since levels of ability and rules can differ between leagues, the use of this method can illuminate differences in scoring probabilities, among other things. It can also be used to investigate whether certain rule changes might be beneficial or detrimental to scoring (though whether or not more scoring leads to a more exciting game is best left to fan debate.)

- A better recording system. This data collection was taken entirely by hand, and prone to human error. Systems have been developed to track the position of the players, and the puck, electronically for the use of organizations such as the NCAA and National Hockey League. They are entirely in the private sector and dependent on technology licensing agreements, making it impractical for lower-tech operations.

- Four-on-four and empty-net play. I have only gone as far as to say what the balance or imbalance in manpower is. But the behavior of teams playing four-on-four is of great interest, as some believe it makes the game more exciting. Empty net play - determining when to pull the goaltender in favour of an extra attacker - is already a frequent topic of study in hockey statistics. Not enough data exists to pursue these questions using this approach at this time.

References

[1] G. Lindsey. An Investigation of Strategies in Baseball. Operations Research, 11(4). [1963]

[2] B. Bukiet, E.R. Harold, J.L. Palacios. A Markov Chain Approach to Baseball. Operations Research, 45(1). [1997]

[3] N. Hirotsu, M. Wright. Using a Markov process model of an association football match to determine the optimal timing of substitution and tactical decisions. Journal of the Operational Research Society, 53(1). [2002]

[4] C.M. Anderson-Cook, T. Thornton. Measuring Hockey's Special Teams Efficiency. Chance, 11(4). [1998]

[5] S. Berry. My Triple Crown. Chance, 13(3). [2000]

[6] E. Pekoz, S.M. Ross. Estimating the Mean Cover Time for a Semi-Markov Process Via Simulation. Probability in the Engineering and Informational Sciences, 11, pp. 267-271. [1997]

[7] Howard M. Taylor, Samuel Karlin. An Introduction to Stochastic Modeling, 3rd edition. [1998]

A Scoring Probabilities

Table 2: A table of scoring probabilities with respect to a starting state and time elapsed.

Time	5 seconds		20 seconds		40 seconds	
State	P(A Scores)	P(B Scores)	P(A Scores)	P(B Scores)	P(A Scores)	P(B Scores)
A: Def. Poss.	0.000185	0.00121	0.00525	0.00914	0.0149	0.0192
	(3.73e-05)	(0.00027)	(0.000961)	(0.00180)	(0.00251)	(0.00329)
A: Neut. Poss.	0.00221	0.000273	0.0116	0.00477	0.0220	0.0141
	(0.000437)	(5.32e-05)	(0.00206)	(0.000857)	(0.00365)	(0.00238)
A: Off. Poss.	0.0122	8.14e-05	0.0267	0.00335	0.0380	0.0122
	(0.00249)	(2.80e-05)	(0.00491)	(0.000579)	(0.00646)	(0.00204)
A: Def. Give.	7.95e-05	0.0142	0.00334	0.0258	0.0122	0.0366
	(2.70e-05)	(0.00417)	(0.000616)	(0.00580)	(0.00209)	(0.00719)
A: Retreat	0.00054	(0.000691)	0.00748	0.00649	0.0176	0.0161
	(0.000126)	(0.000159)	(0.00146)	(0.00119)	(0.00301)	(0.00266)
A: Off. Give.	0.00194	0.000255	0.0103	0.00554	0.0205	0.0151
	(0.000403)	(4.98e-05)	(0.00197)	(0.000987)	(0.00349)	(0.00254)