A Utility Model for the Entry Draft

Brandon Alcox, P. Poupart, P. Curry, K. Larson, C. Moretti
contact email: bkalcox@uwaterloo.ca

April 13, 2017

Abstract

This paper proposes a model to predict how teams might act in the NHL entry draft. We define utilities to quantitatively state how much each team will prefer each player based on scouting information, playing style, and team depth at each position. We use drafting history to approximate team play style to give each team a unique profile of preferences. Using these utility profiles we conduct a series of draft simulations tailored to specific applications or research questions. One clear application of this work is in the creation of autonomous mock drafts, but the greater value is in using it to study key questions in draft strategy. One such question is when is it beneficial to perform a "strategic pick" by selecting a player other than the most preferred one in the hope that he will be available later. This model also enables a study of how perfectly rational teams engage in pick trading. The trading model proposed results in many pick trades occurring with a high proportion of trades being "mutually beneficial" by causing an increase in the utilities of both teams involved by the end of the draft.

1 Introduction

In many major sporting leagues, the yearly rookie draft is an exciting opportunity for teams to gather and determine the future direction of their franchises. From immense pools of talent, front office staff must decide which individuals will best fit the systems and teams they are designing. Naturally there is a major competitive factor involved since whenever a player is drafted by a team, he becomes unavailable to the other teams. Hence, the ability to reason about the actions of other teams is a critical skill when composing a drafting strategy.

From a mathematical perspective, a draft is a fairly simple event. There is a set of players P, a set of teams T, and a draft order O of picks owned by teams in T. Then, when the first pick $o_1 \in O$ is made with the selection of $p \in P$, the draft can recurse with one less pick in O and one less player in P.

The real challenge lies in how the owner of o_1 comes to select player p, and how reasoning about the potential actions of other teams might alter a team's own choices. This paper tackles the first challenge by defining utility functions over teams' player preferences, allowing each decision to be determined numerically. Using these utilities, actions concerning strategic picking and pick trading will also be closely studied by observing how the derived teams behave and interact with one another during draft simulations.

Most of our experiments deal with the 2016 NHL Entry Draft as a source, using the top 181 prospects and the first 100 picks. Some of the player data that is not included suffers from small sample size and could be overvalued in one or two scouting reports, which would cause some late round players to be drafted very early.

1.1 Related Work

Prior research has investigated certain forms of utilities for specific pick numbers, the most notable of which are the pick value graphs, which assign utilities based on expected games played over a career [12]. Using the utility model defined by Schuckers, the problem of pick trading could be considered a resource allocation problem, where work in [2, 11] would become relevant.

The mathematical model for a draft resembles the well-studied game in AI research known as "picking sequences" where agents select objects in turn from a pool of items [3], without perfect information in the case of a draft. Work by Bouveret and Lang [3] as well as Budish and Cantillon [5] investigates how manipulations affect the picking sequences game. Using draft definitions similar to the one proposed earlier, Straffin [13, 4] shows that teams acting strategically with perfect information concerning ordinal preferences may suffer from the Prisoner's Dilemma. Kohler and Chandrasekaran [7] compare different orderings for picking strategies, which ties into the trade modeling done in this paper.

Another paper considers bidding strategies in fantasy auction drafts where each team has a budget that it may spend [1].

2 Defining Utilities

Utility functions for each team $t \in T$ over all prospects $p \in P$ are meant to quantify how much t would like to draft p. As such, it is important to determine the factors influencing these preferences. The raw skills of p, indicated by a vector s_p, must be the driving factors behind these values. Skill assessments are obtained from scouting records from each prospect's draft year. These records include values across 10 categories, but we used linear combinations of the categories to define six "tools": skating, shooting, passing, defensive, grit, and hockey sense. Of course these skills are not of equal importance, which explains why different teams place differing weights on each skill. Some teams are known for fast-paced styles and others tend towards hard-nosed defense-first play. Thus t should have a set of normalized weights w_t over the skills indicating how they are valued. To acknowledge the effects of team depth on drafting, we use a small scaling factor d_t proportional to how deep teams are at forward and defense in comparison to the 12 forward to 6 defensemen ratio found in most NHL lineups. One additional factor is a random number $r_{t,p} \sim Gaussian(1, 0.1)$, which encompasses other reasons that may influence a team's preferences over players. For example, a team's scouts may have seen games when a player played better or worse than his average. Then the utility gain U_t from drafting p is:

$$U_t(p) = r_{t,p} d_t (s_p \cdot w_t)$$

For the utility of a team over an entire draft where t drafts prospects $\{p_1, ..., p_n\}$ we say

$$U_t(\{p_1, ..., p_n\}) = \sum_{i=1}^{n} U_t(p_i)$$

which is a simplifying assumption that ignores complement effects, though tactics to overcome this issue are discussed in Section 5.

2.1 Setting the Weights

We use past drafting history to determine how teams place value on each of the six defined tools. Specifically, if a team drafts a player p_1, while player p_2 (of the same position) is still available then it is likely that $U_t(p_1) > U_t(p_2)$. If we say $c_{p_1,p_2} \geq 0$ is an additive factor that combines the effects of r_{t,p_1} and r_{t,p_2}, then

$$w_t \cdot (s_{p_1} - s_{p_2}) + c_{p_1,p_2} > 0$$

Noting that this inequality is linear in the elements of w_t and c_{p_1,p_2} we can construct an optimization problem that meets the constraints formed by each draft selection and the next five players at their position to be drafted, while minimizing the total utility required to be from random effect $\left(\sum c_{p_i,p_j}\right)$. For a more concrete example consider the following two players with the following skills:

	Skating	**Shot**	**Passing**	**Defensive**	**Grit**	**Sense**
MM	4.01	3.66	4.25	3.52	3.08	4.04
PZ	3.83	3.97	3.94	3.46	3.73	3.74
Difference	0.18	-0.30	0.31	0.06	-0.65	0.30

Table 1: Comparing skill sets of MM and PZ

If a team drafts MM while PZ is still available, one set of weights might be $(0.3, 0.1, 0.2, 0.3, 0, 0.1)$, and if PZ was drafted first, a possible set of weights would be $(0.1, 0.3, 0.1, 0.3, 0.2, 0)$ both without need for utility from c values.

Leading into the 2017 NHL Entry draft there is an interesting question of how to set the weights for the Vegas Golden Knights as there is no draft history to draw from. Perhaps moving to a management-centric approach as opposed to team-centric would be beneficial here as we could draw from George McPhee's history with the Vancouver Canucks and Washington Capitals to set the weights. Another option would be to observe skill sets of all players available in the expansion draft and create constraints as described above but based on which players are selected for the expansion draft and which were passed over.

2.2 Expected Utility

The most interesting questions to tackle with this utilitarian approach require teams to be able to reason about the future consequences of actions they or others might take. For this reason, we must introduce the ability for teams to calculate expected utility for themselves or estimate the value for others considering the resulting draft that would occur if they took various actions. The expected utility for team t over pick order o will be denoted $EU_t(o)$. We use a Monte Carlo approach where for the 29 teams aside from the one making this calculation the values of w_t and the r values are drawn at random from their believed distributions (weights are normalized after sampling) and a draft is simulated over the result of the action being considered. The average utility that the simulations give for team t is the expected utility of t for the action. In our simulations for expected utility we limit the agents to be myopic in the sense that they will assume all future selections will be made naively and without future pick trades occurring in order to keep runtime manageable. If a simulation of 1000 naive drafts each with 99 picks takes one second then 1000 drafts simulations each requiring 1000 subdrafts of 98 naive picks would take 16.5 minutes, and so on for 68 more levels.

3 Naive vs Strategic Drafting

An important question in draft strategy is whether it is ever beneficial to select a player other than your favourite, hoping that he will be available for selection in your next pick. This is known as "strategic drafting" and it allows for the opportunity to select your favourite two players at the risk that you may only be left with your second favourite when you could have had a player you preferred instead if you had made the naive choice. In terms of the draft simulation, this is equivalent to picking a player that results in higher expected utility that is not the same as the player with maximal utility value.

When running our simulations on the first 100 picks of the 2016 NHL draft, the action that nearly always results in higher expected utility is to draft naively, with strategic picks occurring on 4.2% of the 70 picks where strategic drafting is possible (on a team's last pick, the naive action is trivially optimal). It is expected that this is due to the fact that for most picks the next pick owned by the team is 30 selections later. Coupled with the inherent correlation between utility functions of teams this means that if a player is one team's favourite then he is likely good enough to be drafted by another team before the team's next pick.

This raises an interesting question of how close picks need to be for strategic drafting to become a legitimate option. To investigate this, we created drafts with n many teams where picks were evenly distributed, so that each team makes a selection once every n picks. Then by varying the value of n, we see the effect of pick density on the frequency of strategic picks.

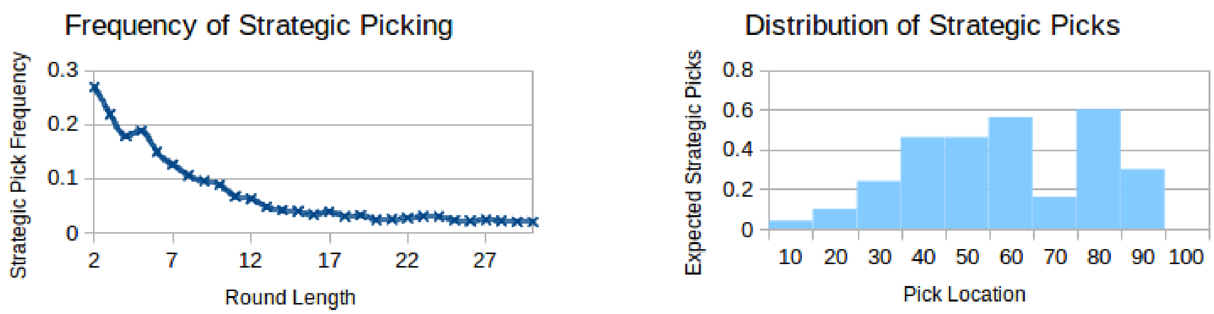

Figure 1: Results of Strategic Drafting Experiments. Left: Frequency of strategic picks decreasing as picks become sparser. Right: Histogram showing the distribution of strategic picking events on 2016 Entry Draft order.

As Figure 1 shows, strategic picking does become far more common when picks are closer together. This data accounts for the fact that final picks for teams are always naive. Also shown is that strategic picking is more common in the middle portion of the draft rather than early when there are clearcut choices or late when teams begin to run out of picks. The picks in the 61 to 70 range were owned by teams whose next picks were relatively far away compared to other stretches, which accounts for notably few instances of strategic picks.

4 Pick Trading

In leagues such as the NHL where pick trading is allowed, a team may be tempted to "trade up" for the current pick in exchange for a package of picks later in the draft if there is a player available that they particularly like. The team with the current pick may "trade down" to gain more or higher quality selections later in the draft if it receives an offer it thinks benefits itself. We enabled teams in our simulation to interact with each other and trade picks to study how rational teams acting in equilibrium would trade picks compared to the cases in real life. We used a simplified negotiation model where, before each selection is made, the 29 teams that do not own the pick may offer a trade for that pick consisting solely of picks to come in the remainder of the draft. These offers may contain multiple picks transferred each way, such as picks in the first and third round in exchange for two picks in the second round. The current owner may choose the offer it believes increases its own utility the most, or keep the draft order the same if no offer is better than making a selection at the time. We do not include existing players or prospects in the trading model as this would require the comparison in value of these assets to picks.

4.1 Creating an Offer

A pick trade between t_1 and t_2 can be thought of as redistributing their picks in such a way that if t_1 originally owned the next pick then t_2 must be the new owner. If the teams own n_1 and n_2 many picks respectively then there are $2^{n_1+n_2-1}$ ways to do this redistribution.

For each possible offer which would lead to order o that t_1 could send to t_2, we consider the expected utility of sending that offer. Specifically this would be the probability the offer is accepted by t_2 multiplied by the expected utility increase for t_1 in comparison to the base ordering o'

$$EU_{t_1}(\text{offering } o) = Pr[t_2 \text{ accepts } o](EU_{t_1}(o) - EU_{t_1}(o'))$$

The probability term is calculated as the number of times in a Monte Carlo simulation (using the method from Section 2.2) where the order o would give t_2 higher utility than o' and each of the best offers of the other 28 teams. This policy will result in a Nash equilibrium as each team will be taking their best possible action relative to everyone else.

Once each team identifies the offer action that would yield the highest expected utility, the original pick owner simply selects the order that would give it the highest expected utility, including possibly selecting o' and keeping the pick.

4.2 Experiments and Results

We ran a series of 30 drafts in which teams were allowed to trade using the methods described above. On each instance of a trade, we evaluated its quality by running a subdraft on the draft ordering before and after the trade. We compared the change in utility caused by the trade for the teams involved, marking the trade as beneficial if the change was greater than zero. For a point of reference during the draft weekend of the 2016 NHL entry draft, there were 8 trades involving at least one of the first 100 picks, only one of which involved only picks (the remaining 7 also included players or prospects), and in 2015 there were 9 pick-for-pick trade ups during the first 100 picks of the draft [8, 9]. On average our model produced 23.2 pick-for-pick trades, each of which used only picks in the top 100.

	Average Utility Gain	**% of Trades Beneficial**
Trade Up	1.26	79.5
Trade Down	1.59	81.5

Table 2: Utility Gains from Trades

Our results shown in Table 2 make some intuitive sense. Specifically, using our trading model leads to an increase in utility and the resulting trades are beneficial with high probability. There are also notably larger benefits for the team trading down than the one trading up. This makes sense considering that when trading down a team will pick its favourite drafting order out of the 30 possibilities, while the team trading up must balance their gains against their trading partner's or suffer low probability of acceptance. For a trade to occur both teams are expecting the result to be beneficial, but there can be a difference between expectation over 1000 draft simulations and what happens in the single real draft where another team could happen to take a player you expected to be available later, which is why we find that about 20% of the time teams may lose a small amount of utility as a result of a trade.

One final result is that when we compare total utility gained by teams over the course of a draft with trading compared to one without trading there is a total increase of 59.4 units, which is greater than the average utility value of drafting Auston Matthews (40.8 units). This calculation of social utility includes the changes to utility for all teams, not just the two involved in any given trade.

5 Future Work

Some work could be done to join our utility model with other similar works. For example, while our utilities can be shown to decrease over the course of the draft in a rate similar to the pick values shown by [12], it may be worthwhile to alter our utilities to fit the utility curve found by Schuckers as closely as possible.

More sophisticated utility models could be considered where team utilities are dynamic to account for substitution and complement effects, since common sense suggests that a team will not seek out a team of identical players. Work with similar dynamic targets in the ordinal preferences instead of utilities has been studied by [6] and could be of use for this line of work. We could also institute some concept of enforcing a balanced team through the utility functions by penalizing teams that do not meet certain thresholds of skill across all categories. An example rule could be that in the total pool of skill of players owned by the team, each tool must account for at least 10% of the pool. This could lead to interesting results for strategic drafting as it might be worthwhile to draft players supplying low utility early if they would help achieve the threshold values to avoid penalization.

Similar to the approach used in [10], agents could be able to move away from equilibrium strategies in the context of interacting with human agents. Human input and reinforcement learning could be used to determine what is a good pick trade to offer to a human peer and how that might differ from the equilibrium trade offer.

Agents could be given higher-level reasoning capabilities to determine what trade offers and pick selections their peers might make using past actions in the draft as data. The risk with this line of work is that it may become beneficial for agents to take actions that are not optimal early in the draft to "lie" to their peers.

6 Conclusions

We proposed a way to look at draft day decisions from a numerical perspective taking into account what makes each team's preferences unique from the rest. Not only is this an exciting new way for fellow analysts to consider why certain teams make the selections they did, but it also allows us to use computers to do the reasoning for us. This allows for study of countless questions revolving around the draft that were previously out of reach.

One question that we tackled in our research was "When is it wise for a team to draft strategically?" We showed that for the true draft where picks are very far apart an instance of a strategic pick should be very rare, though it can be the correct action more often when picks become closer together. However even when picks are only two places apart strategic drafting only occurs at around 25% of the time, so maybe taking the "Best Player Available" is generally the best policy. This does assume that all teams are drafting with the similar utility functions, so if one team thinks they may have found a draft inefficiency which breaks the correlation between their utilities to everybody else then it is still an open question as to whether they should draft strategically or not.

We also looked at cases where pick trading is beneficial to both teams involved. With our trade policy in a Nash equilibrium, we found far more trades than were performed in the actual 2016 Entry Draft, and showed that with high probability they would lead to mutual utility gain. This would certainly add another level of excitement to the draft for the average fan trying to keep up with who owns which picks as the draft order would constantly be changing, but would also leave the individual teams and the entire league far happier by the end of the event.

The experiments run in this paper only scratch the surface of the possible research that our methods could lead to, and we look forward to seeing how others might use our model to help the hockey world better understand the draft.

Acknowledgements

The authors wish to thank HockeyTech Inc for providing the scouting data that made this work possible, and to the OCE VIP II grant for helping to fund this project.

References

[1] Aris Anagnostopoulos, Ruggiero Cavallo, Stefano Leonardi, and Maxim Sviridenko. Bidding strategies for fantasy-sports auctions. In *International Conference on Web and Internet Economics*, pages 102–115. Springer, 2016.

[2] Haris Aziz, Páter Biró, Jérôme Lang, Julien Lesca, and Jérôme Monnot. Optimal reallocation under additive and ordinal preferences. In *Proceedings of the 2016 International Conference on Autonomous Agents & Multiagent Systems*, pages 402–410. International Foundation for Autonomous Agents and Multiagent Systems, 2016.

[3] Sylvain Bouveret and Jérôme Lang. Manipulating picking sequences. In *ECAI*, pages 141–146, 2014.

[4] Steven J Brams and Philip D Straffin. Prisoners' dilemama and professional sports drafts. *The American Mathematical Monthly*, 86(2):80–88, 1979.

[5] Eric Budish and Estelle Cantillon. The multi-unit assignment problem: Theory and evidence from course allocation at harvard. *The American economic review*, 102(5):2237–2271, 2012.

[6] Hadi Hosseini, Kate Larson, and Robin Cohen. Matching with dynamic ordinal preferences. In *AAAI*, pages 936–943, 2015.

[7] David A Kohler and R Chandrasekaran. A class of sequential games. *Operations Research*, 19(2):270–277, 1971.

[8] NHL Public Relations. Trades completed during the 2015 nhl draft weekend. https://www.nhl.com/news/trades-completed-during-2015-nhl-draft-weekend/c-772519, 2015.

[9] NHL.com. 2016-17 nhl trade tracker. https://www.nhl.com/news/2016-17-nhl-trades/c-281003458, 2017.

[10] Noam Peled, Sarit Kraus, et al. A study of computational and human strategies in revelation games. *Autonomous Agents and Multi-Agent Systems*, 29(1):73–97, 2015.

[11] Arvind Sathi et al. Constraint-directed negotiation of resource reallocations. 1989.

[12] Michael E Schuckers. What's an nhl draft pick worth? a value pick chart for the national hockey league. 2011.

[13] Philip D Straffin. *Game theory and strategy*, volume 36. MAA, 1993.