Emergent League Parity in a Multi-Agent Simulation of the National Basketball Association

Authors Redacted

Abstract

The National Basketball Association (NBA) seeks to gain and retain as many fans as possible. We propose that a reasonable way to numerically measure the attractiveness of the league to fans is by examining the league's parity. Thus, the NBA seeks to maximize league parity in order to attract as many fans as possible. By simulating how various regulations on the league's free agency affect the parity of the league as a whole, we believe that the NBA can begin to understand how various policy changes may affect parity and thus fan attraction to the league. In addition, we seek to explain why the various league policies affect parity through analysis of our simulations of free agency.

Introduction

The NBA desires to achieve a certain level of league parity in order to retain fandom and maintain revenues in all markets. The NBA believes that it can establish and maintain league parity by enforcing certain rules related to offseason free agency. In general, NBA Players and Teams have self-centered goals. NBA Players are concerned with maximizing the amount of money they can earn and/or achieving success with their team, and NBA Teams seek to attain the best Players and become the best Team. By imposing regulations related to total Team spending on Players, the NBA can attain some control over the outcomes of the offseason free agency period. This paper presents a multi-agent system (MAS) that simulates the relationship between Players and Teams within the boundaries of adjustable league policies. The desired emergent behavior for this MAS is to achieve and maintain a high level of league parity. The parity will be achieved via autonomous agents seeking to maximize their utility. This is a desirable outcome for the NBA because it will help keep fans of all Teams interested and engaged. NBA league parity will be measured by computing a Power Index for each Team at the beginning of each iteration (after all offseason free agency activity has finished) and tracking how that Power Index changes over time. The specifics of parity computation are discussed in the Desired Emergent Behavior section. The NBA seeks to maximize this parity value.

Each agent type will operate under two conditions:

1. An agent will not be able to communicate directly with any other agent of the same type.
2. Agents will strive to maximize their respective utility (see Simulation Design section for more information on agent utility).

There are multiple properties that can be changed in order to affect NBA league parity. These are discussed fully in the Environment Design section. In order to conduct experiments to analyze the emergence of desired NBA parity, we have addressed several hypotheses that will serve as the basis for our experimentation. These hypotheses are:

- **Hypothesis 1:** Decreasing the salary cap will increase league parity. This will make it more difficult for Teams to retain multiple highly skilled Players, and as a result few Teams will be significantly better than the rest.
- **Hypothesis 2:** Increasing the Team Preference Factor, which is the weight Players place on Team Prestige (as opposed to contract size), will monotonically

Commented [LS1]: Hmm .. not appropriate here. This means that you already know what your emergent behavior is … thus no point in carrying out the research.

decrease league parity. Players with high skill values will be more likely to cluster on the best Teams.
- **Hypothesis 3:** Increasing the Max Contract Size will increase parity in the league. This is because a higher Max Contract Size makes it more difficult to retain good Players, since other Teams will be more likely to offer bigger contracts.
- **Hypothesis 4:** We believe that parity will be minimized when 50% of the Teams are willing to go over the salary cap, and will be maximized when either all or none of the Teams will exceed the salary cap. When half of the Teams are willing to exceed the salary cap, these Teams will be able to offer more money to the best Players.
- **Hypothesis 5:** Teams that sign Players to longer deals will have a higher Power Index, since they will be more likely to keep good players for longer. This will result in lower parity, since the Teams willing to sign longer contracts will remain the same year-to-year.

This report will focus on the MAS design, as well as providing analysis of experimentation relating to the hypotheses listed above.

Related Work

In [1], the author proposes using the knapsack model to optimize a Team's approach to signing new free agents each offseason. The author does explore similar concepts presented in this paper; however, he does not examine free agency as it relates to the Player's utility function. In addition, the author does not explore how a Team's approach to signing free agents might affect overall league parity as discussed in this paper.

In [2], Rockerbie attempts to explain the relatively higher parity in the NBA compared to other North American sports leagues. He proposes that the high volume of scoring opportunities in basketball could be the cause of the relatively higher parity in the NBA. The paper, however, does not address our core experimentation goal of attaining parity.

Other literature has examined more general concepts in the economics of sporting leagues and player trading [3]; however, we believe that our research in the field is unique in the fact that it models utility functions for both the NBA Player and Team. Furthermore, we seek to explore how these dynamics affect the parity of the league as a whole and how parity will respond to changes in the free agency policy.

Simulation Design

The system will be comprised of two agent types: NBA Players and NBA Teams. Each agent class will be primarily concerned with obtaining its highest individual utility. NBA Players will attempt to maximize their annual contract size. NBA Teams will attempt to maximize their overall team Power Index.

NBA Player Design

A NBA Player can either be currently contracted by a Team or available to be signed by Teams. Players will be available if
- They were on a Team but their contract expired
- They entered the system through the Draft (discussed in Environment Design)

Players have a property that defines a finite number of years for which they are available to play in the league. Once a Player has been in the league for its allotted time, a Team can no longer contract the Player and the Player will be retired from the system.

The goal of a Player that has not yet been signed to a Team is to sign a new contract with a desirable team for as much money as possible. The Player's Team Preference Factor (TPF) represents the relative weight of these two priorities. A TPF of zero means the Player only considers the dollar amount of an offer, and a TPF of one signifies that the Player only considers the skill of the Team that the offer is from.

Players also have a skill value that is based on Player Efficiency Rating (PER), a metric commonly used to evaluate real NBA players. This skill value is used by both the Player and Teams sending offers to the

Commented [LS2]: Nice!

Commented [LS3]: Nice

Player when ~~the~~ computing the values of offers and in determining whether an offer is accepted or rejected.

Player Decision Function

The decision function for a Player looks at the offers that the Player has received in the past day, decides which one is best, decides if it is a "good enough" offer, and then either accepts the best offer or rejects all offers and waits for the next day's offers. Each offer is evaluated using the following formulas:

$$value = \frac{offer\ value - \min(salary)}{\max(salary) - \min(salary)}$$

$$team = \frac{offering\ team's\ PI}{\max(PI) - \min(PI)}$$

$$offer\ value = (1 - TPF) * value + TPF * team$$

The maximum possible Power Index for a Team is approximately thirty-five (if a team somehow had all of the best possible players) and the minimum possible Power Index is zero (if all players on the team have a skill level of zero, which is the minimum skill level). The value and team components are therefore both scaled to be between zero and one. The Player selects the offer with the highest value according to this equation.

Once the decision function has found the best offer, it determines if it is an acceptable offer. The factors that contribute to this decision are the Player's value and desperation factor.

Player Value

The monetary value of a Player depends entirely on its skill level. In order to compute this value, NBA player salaries were scraped from ESPN.com, paired with the matching player skill, and used to create a least squares regression model [6]. The best fitting function had an R^2 of 0.989, which indicates a strong model. The equation is shown below, along with a graph of the function and the points used to fit the model:

$$player\ value = 12671051.1 - 3003452.8 * PER + 217568.5 * PER^2 - 3483.8 * PER^3$$

Figure 1: Player Value

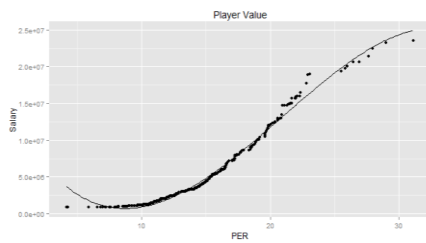

This function is then used to find a Player's monetary value given its skill, which is used in both the Player and Team decision functions.

A few modifications were made to the function output by the regression model. First, the value of the function is limited inside the range of PER values that produced its local optima (8.735 and 32.899). For skill levels outside of this range, the function was adjusted to a straight line that approached the minimum individual salary (for PERs less than 8.735) or the maximum individual salary (for PERs greater than 32.899). Limits were also placed so that Players are never valued at values above the maximum individual salary or below the minimum individual salary.

Desperation Factor

The other factor that the Player decision function takes into account when deciding whether or not to accept the best offer is the desperation factor. This is a Player-specific attribute that models growing desperation as the offseason passes by while the Player remains unsigned. The desperation factor can be expressed as:

$$desperation = 0.993^d$$

where *d* is the number of days into the offseason. This results in a multiplier that rapidly decays to 0.495 by the end of an offseason if the Player has not yet signed with a Team. At the end of each offseason, the desperation factor is reset to 1 for all Players.

Commented [LS7]: Each figure must be referred to at least once in the text, and that reference should come before the figure appears in the text.

Commented [LS4]: Hmm … both these equations use global information: max(salary), min(salary), max(PI), min(PI)… not that good.

Commented [LS5]: How's this computed?

Commented [LS6]: Nice work.

Commented [LS8]: How did you get 0.993?

Commented [LS9]: What is the length of an offseason? Not clear.

The final decision by a Player on whether to accept or reject the best offer is based on the idea of a reservation price [8]. The desperation factor multiplied by the Player's calculated value forms this Player's reservation price. When deciding to accept or reject the best offer, the Player compares the value of the offer to its reservation price and will simply accept the offer if it is greater than or equal to its reservation price, and reject otherwise.

The Player decision function also does not allow Players to accept contracts that last longer than the number of years the Player has remaining in the league. If the decision function has decided to accept an offer where this is the case, the length of the contract is cut down to the number of years that the Player has remaining in the league.

NBA Team Design

A NBA Team will be a collection of Players. Each year, a Team is primarily concerned with maximizing its Power Index relative to the environmental constraints discussed in the Environment Design section below.

Power Index

The Power Index for a NBA Team is ~~calculated is~~ calculated by taking a weighted sum of the skills of its Players. The player ratings will be ordered such that the highest rating on the team receives the highest weight with each successive rating receiving a lower weight. The weight associated with each rank will be tied to minutes-played statistics from the real NBA. Each Team's total minutes per game sums to 240 minutes (5 players on the court for 48 minutes each). On average, the Player that plays the most minutes for an NBA team plays 26.2 minutes per game, so the weight given to the rating of a Team's best player is $\frac{26.2}{240} = 0.109$. The rest of these weights can be found in Table 1.

Table 1: Player Rank Weights

Player	Average Minutes	Weight
1	26.2	.109
2	24.9	.104
3	23.5	.098
4	22.2	.092
5	20.9	.087
6	18.6	.078
7	17.1	.071
8	15.4	.064
9	14.0	.059
10	13.0	.054
11	11.2	.047
12	10.0	.041
13	8.6	.036
14	7.7	.032
15	6.7	.028

Teams attempt to maximize this Power Index are subject to a few constraints. A Team cannot have more than fifteen players on its roster, which is consistent with NBA regular season regulations. Teams are allowed to have fewer than fifteen players; they will simply receive a contribution of zero to their Power Index from that roster spot. Another constraint faced by Teams is the salary cap. The sum of the annual salaries of all the Players on a Team's roster must be less than this salary cap. The 2014-2015 NBA cap (and the default value used in the simulation) is $63,065,000 [5]. Since the NBA allows teams to exceed this cap and choose to pay a luxury tax, we will allow for this in our simulation. Some Teams will randomly be designated as willing to exceed the salary cap, and the designated Teams will be assigned a random amount that they are willing to exceed the cap by. The salary cap,

Commented [LS10]: Why not also representing this as an equation?

probability of a Team being willing to exceed the cap, and maximum amount a willing Team will go over the cap by are all parameters that can be controlled in the simulation.

Each tick in the simulation represents a day in the offseason. Each Team has the option to offer one contract to one Player each day. Within the constraints of the current roster size and salary cap, a Team can offer a salary and contract length to any Player not currently signed to a Team. At the end of each day, Players with contract offers must accept one offer, or decline all of their offers and remain unsigned.

Team Decision Function

The decision function for a Team determines to which Player the Team will send an offer ~~to~~, and what dollar amount to offer that Player. Every Team performs this every day in the offseason (Teams with full rosters simply choose to not send out an offer). The decision function begins by looking at all available players and determining how much the addition of each Player would increase the Team's Power Index. It then uses the ratio of that amount to the Player's value to determine the "power ratio" for that Player to the Team. The Team selects the Player with the highest "power ratio" to send an offer to.

To determine the dollar amount for an offer, the Team takes into consideration how much money it has left to spend (until it hits its salary cap) and the number of roster spots it still has left to fill. Let $P = \{p_1, p_2, \ldots, p_m\}$ be the set of available Players, in decreasing order of impact they would have on the Team's Power Index (so adding p_1 would cause the Team's Power Index to increase by the largest amount). Let *r* be the dollar amount that the Team has available to spend (*r* is the Team's salary cap, including however much it is willing to exceed the base salary cap by, minus the sum of the contracts for Players currently on its roster). Let *n* be the number of spots remaining for a team (meaning there are 15 - *n* Players on the Team's current roster). The Team will offer the selected Player the amount determined by the following formula:

$$\max\left(\min\left(r * \frac{p_1}{\sum_{i=1}^{n} p_i}, \max(size)\right), \min(size)\right)$$

This formula guarantees that offers will be within the individual contract size restrictions for the NBA, but it remains a possibility that the offer is far too low for the Player. To prevent offers from being wasted, the decision function checks if the offer meets an estimated reserve price for the targeted Player. This reserve price is estimated by taking the value of the Player, and multiplying it by a discount factor, which by default is a random number between 0.8 and 1 that represents how risky the Team is. If the computed offer amount is greater than or equal to the estimated reserve price, then the offer remains under consideration. If it is not, the decision function decides that the current Player is too expensive to make an offer to, removes the Player from P, and repeats this process with the new set of available Players.

The last step in the decision function makes sure that the Team is planning ahead and remains far enough below the salary cap to be able to sign fifteen Players. To do so, the decision function finds the value of the offer that it would give to the fifteenth Player that it would sign, and makes sure that the value is above the NBA's minimum salary. This value is found by the following formula:

$$offer\ value = r * \frac{p_n}{\sum_{i=1}^{n} p_i}$$

If this value is less than the NBA's minimum salary, this indicates that the Team would be spending too much of its remaining budget on the current target Player. If this is the case, the target Player is then removed from the list of available Players, and the process is repeated. If the list of available Players ever becomes empty, the Team does not make an offer on the current day. If an acceptable target Player and offer amount is selected by the decision function, the Team sends an offer to the selected Player.

Commented [LS12]: Proportional, okay.

Commented [LS13]: Why this?

Commented [LS14]: Why 0.8? why 1? Ad hoc, without sufficient justification.

Commented [LS11]: Would be better to also express this as an equation.

Environment Design

Agent Initialization

The simulation is initialized with 30 Teams corresponding to the 30 teams in the NBA. The initial Players in the environment are based on real NBA players. The PER values and team assignments of NBA players were scraped from ESPN.com and used to assign Players to simulated Teams that match the real skill levels of the NBA players on the actual 2014-2015 NBA rosters.

Player Entry/Exit

To emulate the NBA draft, which serves as the entry method for Player agents into the environment, we add ninety randomly generated Players without assigned Teams to the pool of available agents. To populate the attributes of these randomly generated Players, Players are randomly sampled from the initial set of Players created based on the 2014-2015 NBA rosters, and their attributes are copied to create new Players, allowing us to maintain a distribution of Player skills equivalent to that of the real NBA.

The one attribute of Players that is not copied when creating new Players to add to the environment is the number of years until retirement of that Player. In order to maintain an equilibrium number of Players in the league, this value is selected from an exponential distribution, since exponential distributions are commonly used to model time until death. To determine the parameter for this exponential distribution, we found the actual NBA distribution of years of service for 2014-2015 players by scraping data from ESPN.com. From this distribution, we observed the number of active years of service in the NBA and found that the average years of service is 5.59. We decided to round this number to six and therefore model the distribution of years to retirement as an exponential distribution with a λ of ⅙. This allowed us to maintain a distribution with approximately 450 players in the league each year.

To confirm that our distribution correctly modeled the NBA, we compared the distribution of years of service of our agents in the final year of our simulation to the distribution of years of service among real 2014-2015 NBA players. The following plot in Figure 2 compares the distributions:

Figure 2: Distribution of Years of Service

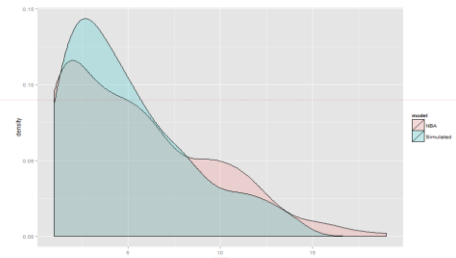

As you can see, our generated distribution closely models the actual distribution of years in the NBA from our scraped data.

Parameters

The NBA has certain rules and regulations regarding free agency, and the goal of this research was to determine the effect changes in these rules would have on league parity. The system was built with customizable parameters that allow for the adjustment of these rules. The different parameters are listed below, with default values and any justification provided.

- Contract Adjustment **{0}:** By default, contracts are randomly assigned lengths of three, four, or five years (each with equal probability). This parameter lets the simulation adjust those lengths. For example, a contract adjustment of -1 would cause contracts to be randomly assigned a length of two, three, or four years.
- Max Individual Contract Size **{$20,644,400}:** The maximum amount that can be paid to a particular athlete each year as part of a contract. The default value is the 2014-2015 NBA maximum [4].
- Max Team Salary Cap Overage **{0.2}:** The maximum proportion that a team is willing to exceed the salary cap by.
- Min Individual Contract Size **{$507,336}:** The minimum amount that can be paid to a particular athlete each year as part of a

Commented [LS15]: Good.

Commented [LS17]: Well. The area under the curve is not negligible. So, I wouldn't say "closely models".

Commented [LS16]: Thoughtful.

Commented [LS18]: Justification?

contract. The default value is the 2014-2015 NBA minimum [4].

- Percentage of Teams willing to Exceed Salary Cap **{0.1}:** The percentage of Teams who are willing to pay the penalties associated with exceeding the salary cap.
- Proportion of Risky Teams **{1}:** The percentage of Teams that will give Players offers that are lower than their value.
- Risk Percentage (for Risky Teams) **{0.8}:** The lowest value that a risky Team will offer a Player (as a proportion of the Player's value).
- Salary Cap **{$63,065,000}:** The maximum amount that a Team can pay to its entire roster each year. The default value is the 2014-2015 NBA Salary Cap [5].
- Team Preference Factor **{0.5}:** The relative weight given to the Power Index of the offering Team versus the offer amount when a Player is evaluating offers.

Desired Emergent Behavior

The desired emergent behavior for this multi-agent system is achieving and maintaining a high level of league parity. This is desirable for the NBA because it will keep fans of all Teams interested. Parity is computed as follows:

$$v_i = \sum_{t\epsilon T}\left(\sum_{d=1}^{i-1}\left(\left|I_{t_i} - I_{t_{i-d}}\right| * w_d\right)\right)$$

where I_{t_i} is the Power Index of Team *t* in year *i*, T is the set of Teams, and w_d is the weight given to that entry. We define the weights as follows: $w_1 = w_2 = 0.1, w_3 = w_4 = w_5 = 0.2, w_6 = w_7 = 0.1, and\ w_i = 0\ \forall\ i > 7$. These weights were chosen to value changes in the three to five year span more than recent or distant changes. For calculating parity in years before the eighth year, the unused weights are be dropped, and the remaining ones are be recalculated proportionately. For example, to calculate the parity in year 4, the weights would be: $w_1 = w_2 = 0.25\ and\ w_3 = 0.5$.

Repast

Our multi-agent system was created using Repast Simphony. We used the Repast modeling system to define our agents as Java classes. Repast allowed us to visualize various aspects of the state of agents in the multi-agent simulation, which was useful during the development and tuning of the system. One such visualization was the average number of agents signed to a team (shown below).

Figure 3: Average Number of Agents Signed to Teams

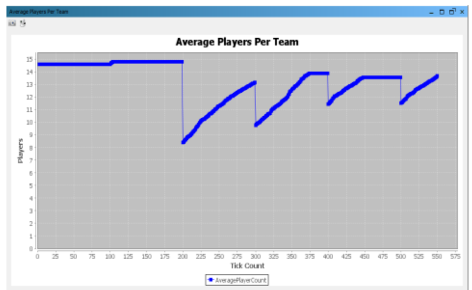

This helped us ensure that Players are correctly being signed to Teams. We also used text sinks to log data on the parity of the league in csv files. This data was then loaded into the R programming environment for our final analyses.

Experimentation

Experiment Setup

In each experiment, the default values found in the Parameters section above were used with the exception of the variable being manipulated. Each experiment was run with seed number 60,718,094.

Results

Experiment 1

The results of Experiment 1 can be seen in Figure 4. For this experiment, we ran the simulation 250 times with salary caps ranging from $40 million to $120 million and plotted the resulting median league parity. We then fit a linear model to the data, resulting in a negative coefficient (downward trend). In the linear regression, the value of the coefficient

Commented [LS19]: I have the same comment before. Why 0.8?

Commented [LS20]: What's the rationale behind this design?

Commented [LS21]: grammar

Commented [LS22]: I don't get this. Why 0.5 for W_3?

of salary cap indicates that the model is statistically significant ($p < 0.05$).

We believe that the data from this experiment supports our hypothesis that decreasing the salary cap will increase league parity. We also believe that this outcome is intuitive; lowering the salary cap makes it more difficult for teams to retain the best players (whose contract offerings will be relatively high) when contracts run out, meaning that more teams have opportunities to sign the most valuable players.

Figure 4: Salary Cap vs. Parity

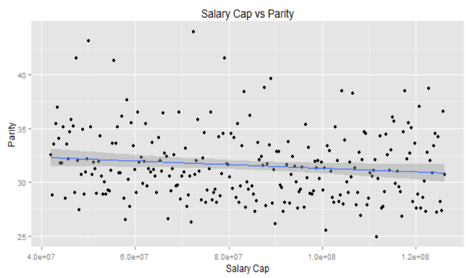

Experiment 2

In this experiment, we investigated the impact of Team Preference Factor on league parity. The results of the experiment are available in Figure 5. The increase in players' TPF was varied from 0 to 1 by increments of 0.005. While there is an upward trend to the linear regression model below, we found the results to be inconclusive. The p-value of the regression model was significantly greater than 0.05 (0.43). In addition, the standard error was greater than the linear coefficient. These conditions indicate that the results are not statistically significant.

We believed that a relatively higher TPF across Players would cause the Players with the highest PERs to cluster to the best Teams, resulting in an overall low parity. However, the data indicates that this is not the case. Our team proposes a couple explanations for this: a large percentage of the Players would necessarily have to receive multiple offers for their TPF to have an effect on which offer they accept. However, we have reason to believe that this may not be this case; if most Players get at most a single offer per round, or their offers are significantly different in price, then varying their TPF would not change the outcome. Additionally, we believe that changing Players' TPFs would not affect league parity because of the low Player turnover within Teams year-to-year. At the end of each season, the average Team loses up to three Players to retirement or contracts ending. This relatively low amount of turnover, and therefore relatively small amount of cap space freed up by the Players leaving, may not be sufficient to contract the best Players that would seek to play with the best Teams.

Figure 5: Team Preference Factor vs. Parity

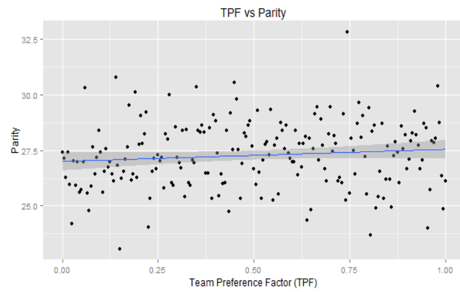

Experiment 3

Next, our team sought to determine the effect raising the maximum contract size that a team could offer a Player would have on league parity. The results from this experiment can be seen in Figure 6. Each box in the plot comprises 50 runs of an experiment with the given maximum Player contract size. It is interesting to note that $20.6 million is the actual Player contract cap mandated by the NBA during the 2014-2015 season. The horizontal line in the middle of each box denotes the median parity for each dataset. Varying the max Player cap does not have a significant effect (if any) on league parity. Our team believed that a larger maximum Player contract size would increase parity because it would make it more difficult for teams to retain good Players (as it would become easier for other Teams to "poach" the Players by starting a bidding war with the newly higher contract cap). However, the data does not

Commented [LS23]: Good.

Commented [LS25]: Why not check out the data and see whether any of these two are true??

Commented [LS26]: Why this observation?

Commented [LS27]: Interesting.

Commented [LS24]: Interesting.

support this hypothesis and we believe the results to be inconclusive.

Our team believes that this outcome is at least partially due to a simplifying assumption made in our player model: contract offers are based on Player values, which were modeled off of the real NBA during the 2014-2015 season. Raising the maximum Player cap enables Teams to offer larger bids to individual Players; however, it does not change how Players value themselves. Therefore, Teams have the option to bid more but will not because Players' values have not increased.

Figure 6: Player Cap vs. Parity

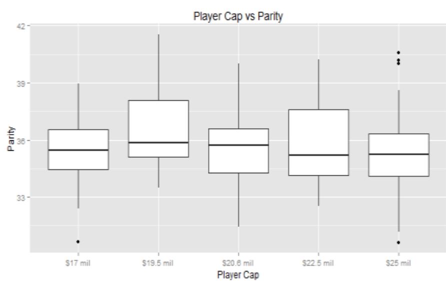

Experiment 4

In this experiment, we explored the effect of having Teams willing to exceed the salary cap. The results from this experiment can be seen in Figure 7. Our hypothesis that parity would be minimized when 50% of Teams were willing to exceed the salary cap and maximized when all or none of the Teams would exceed the salary cap was the opposite of what was observed during this experiment. Our initial justification was that when only half of the Teams were willing to exceed the salary cap there would be a large disparity between those Teams that exceed the cap and those who did not.

We believe that there are a few reasons for the difference in our hypothesis and what was actually observed. The first possible explanation is that the Teams that are willing to exceed the salary cap are not necessarily the best Teams (which teams that are willing to exceed the cap is random); this means that lesser Teams could outspend better Teams, therefore improving their Power Index and increasing parity. Additionally, we believe that roster turnover year to year may not be large enough to see the spending above the salary cap have a large effect on league parity. Teams are limited to 15 Players and only lose two or three per year on average. It is possible that the maximum roster size was the limiting factor when determining roster makeup rather than salary cap (teams have sufficient funds to fill their 15-man rosters with good Players).

Figure 7: Percentage of Teams Over Salary Cap vs. Parity

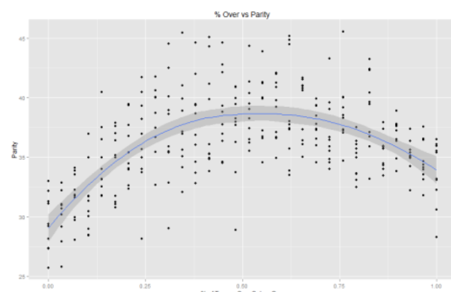

Experiment 5

In the fifth experiment, we investigate the impact of contract length on league parity. By default, Teams randomly offer contracts of length 3 to 5 years to Players. We investigated modifying this range by two fewer, one fewer, one more, and two more years. The results of the experiment are available in Figure 8. We believed that increasing contract length would decrease parity because the best Players will be signed to Teams for longer, decreasing the variance in Power Indices year to year and therefore decreasing league parity. However, the results of our experiment do not completely support our hypothesis.

We believe that there are a couple possible explanations for these results. With 1 to 3 year contracts (-2 in Figure 5), the league experiences a very high turnover rate. This prevents Teams from locking down the best Players for long periods of times, increasing parity. As the contract length

Commented [LS29]: Hmm …. Insightful but not completely convincing ….

Commented [LS28]: But how would the teams know how each player values himself? I don't see the logical flow in your justification here.

Commented [LS30]: Hmmm .. if best players stay with a team longer, then it is more likely that the team dominates longer.

Commented [LS31]: Yes, exactly my thought.

increases from -2 to 0, our hypothesis holds and parity decreases. However, when the contract lengths increase, parity begins to rise. We believe this unanticipated result is due to Teams having relatively low turnover rates; instead of losing two or three Players a year on average, most Teams only lose one or two. This results in most Teams having only a single roster spot to fill, freeing them up to use any remaining cap space on the best Players available, rather than having to divide it between multiple Players more economically. In summary, both very low and very high Player turnover rates among Teams result in higher league parity.

Figure 8: Player Contract Length vs. Parity

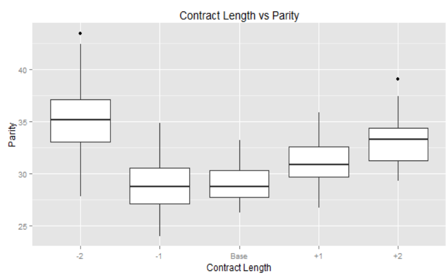

Discussion

Two experiments yielded inconclusive results. We believe that the slow average turnover rate of two to three Players a year was a factor in the inconclusive results and also decreased the effect of the changes introduced in the other three experiments relative to our expectations. When new policies affect 20 percent or less of a Team in a given year, they likely will not have a large effect on parity for a short period of time. While the slow turnover rate was an unexpected factor in our experiments, we would leave it unchanged for future work. In the real NBA, losing two or three Players a year is the norm and thus our simulation accurately models the real NBA in this aspect.

The dynamics involved even in a simplified system like the one we designed are clearly more complex than they might initially seem. Intuitive hypotheses proved to be incorrect or inconclusive. Though these differences are explainable through additional analysis, the causes of the differences were not initially clear.

Emergent Behavior

Our simulation exemplifies the emergent behavior that we as system designers desired at the outset of this paper: league parity (a global value) is affected by completely autonomous agents acting to maximize their utility. What's more, Teams and Players act in ways that model the real NBA without explicit enforcement by the system. For example, NBA teams have roughly a minimum of 13 players [6] and a maximum of 15 players [7]. The system enforces the maximum but not the minimum; however, teams rarely fall below 12 players on their roster during the offseason.

Real-world Application

This simulation applies to the real world in meaningful ways because it was derived from a real-world problem using real data. Based on our experimentation, we can recommend a few changes that the NBA could implement in order to help maximize parity. The first recommendation would be to lower the salary cap. In Experiment 1 we determined that our hypothesis was supported by data suggesting that reducing the salary cap increases league parity. In addition, we would recommend that the NBA encourage roughly half of its teams to spend beyond the salary cap. In Experiment 4 we found that league parity is maximized when 50 percent of teams spend beyond the cap. Finally, we would recommend that the NBA attempt to attain either a relatively high or relatively low turnover rate. We saw either end of the contract length spectrum yield high parity. However, in order to maximize its Power Index, NBA Teams seek to sign players to longer contracts. This would conflict with the recommendation offered to the NBA.

Hindsight

In our current model of a Player, a Player enters the simulation with a skill rating generated from a distribution. Throughout the career of the Player in the simulation, their skill does not increase or decrease. In the real NBA, skill increases based on experience. We believe that modifying Player skill

Commented [LS33]: Insightful.

Commented [LS32]: Nice!

Commented [LS34]: Cool.

Commented [LS35]: Pretty neat! Meaningful!

based on their years in the league and previous success (the history of the power indices of their previous Teams) would be a reasonable way to model Player skill over time. In addition to more accurately modeling the NBA free agency, we believe that this tweak could lead to conclusive results in Experiment 3, as Players' values would change over time and thus a greater Player contract cap could become effective.

An interesting data set that we did not collect is data on offers sent out to Players. This data set would have given us insights such as offer success rate, average number of offers per Player per day, the values of offers, and many other interesting offer-related metrics. Given this information, we could have found more detailed explanations as to why certain experiments resulted in unexpected outcomes. For example, data on the number of offers received by a player each day, as well as data on the range of values for those offers, would likely have shed some light on why Experiment 2 resulted in TPF having an inconclusive effect on parity.

Future Work

The research we conducted provided some interesting and useful insights into the dynamics of free agency in the NBA. However, there were simplifying adjustments made and some factors not included in our modeling that could have important implications for league parity.

For contract simplicity, our system does not consider the previous Team a Player was signed to as a factor when deciding between offers. In the NBA, however, it can be argued that there is generally a significant chance that a player will re-sign with his current team when its contract expires. However, this is not always an assumption that can be made, and would vary greatly between players. Due to this complexity, and the fact that it is a fairly intangible factor, it would require additional research and modeling before it could be successfully implemented.

Another contract limitation of our system is the fact that the decision function for a Team finds the amount it will pay given how a Player would impact its Power Index. This misses the potential for Teams to enter bidding wars on a particular Player, since the initial offer is the upper limit of what a Team would be willing to pay.

Our simulation incorporates a yearly draft to repopulate the system with new Players as other agents retire and leave the league. This is currently implemented as an influx of Players to the Free Agency pool. In the NBA, however, teams will directly draft new players each year. Assigning particular generated Players to Teams based off of their record could successfully incorporate a more realistic draft system. In the NBA, teams have a greater chance of getting an early pick if they had a poor record. This could also allow for the incorporation of tanking. This is where a particular team knows it does not have a good chance of doing well in a particular year and intentionally loses games. This leads to a worse record and increased probability of attaining a high draft pick. This adds significant complexity to the system, but would also allow these concepts and their repercussions to be better explored.

Lastly, our experiments focused primarily on the alteration of one variable (NBA league rule) at a time. Some of our conclusions suggest, however, that altering multiple parameters at the same time could potentially yield more positive results. This would require careful analysis and planning, however, because changing more than one parameter could have cascading effects on different aspects of the final result. The scope of these considerations is outside the focus of this paper.

Commented [LS36]: Insightful!

Commented [LS37]: Yes, insightful!

References

1. Sam Kirchner. Optimizing an NBA team's approach to free agency using a multiple-choice knapsack model. MIT Sloan Sports Analytics Conference. 2011.
2. Duane W Rockerbie. Exploring inter-league parity in North America: the NBA anomaly. Munich Personal RePEc Archive. 2012.
3. John Vrooman. The economics of American sports leagues. Scottish Journal of Political Economy. 2000.
4. Larry Coon. NBA salary cap FAQ. cbafaq.com. 2015.
5. 2014-15 NBA player contracts. basketball-reference.com. 2015.
6. ESPN. NBA Teams. ESPN.com. 2015.
7. National Basketball Association. Collective Bargaining Agreement. 2011.
8. Amos Azaria, Yonatan Aumann, and Sarit Kraus (2014). Automated agents for reward determination for human work in crowdsourcing applications, A~~a~~utonomous A~~a~~gents and M~~m~~ultiagent S~~s~~ystems, 28(6):934-955.