The Automated General Manager: An Unbiased, Backtested Algorithmic System for Drafts, Trades, and Free Agency that Outperforms Human Front Offices

Philip Z. Maymin
Vantage Sports and
University of Bridgeport
Bridgeport, CT 06604
+1 203 576 2027
philip@maymin.com

ABSTRACT

I introduce an automated system and interactive tools for NBA teams to better decide who to draft, trade for, and sign as free agents. This automated general manager can serve as an expert-system replacement, a complement to a team's front office, or as a benchmark against actual performance. Backtested over ten years, the automated GM outperforms every single team, and by substantial margins. From draft decisions alone, the average team lost about $130,000,000 worth of on-court productivity relative to what they could have had with the automated GM; this shortfall represents about a quarter of the average franchise value. Historically the automated GM's choices would have produced about twice as much as the human choices actually did: approximately one extra win per year per draft pick. The system is calibrated using an innovative extension of traditional machine learning methods, applied to a uniquely broad historical database incorporating both quantitative and qualitative evaluations, in a way that avoids survivorship bias, and for a variety of performance metrics; it is thus robust, comprehensive, realistic, and does not overfit information from the future. I provide all of the interactive tools supporting this paper, including backtesting results, projections, scenario analysis, and more, online, for free, at nbagm.pm.

CCS Concepts

• Computing methodologies→Machine learning→Machine learning approaches→Neural networks

• Computing methodologies→Machine learning→Machine learning approaches→Classification and regression trees

• Applied computing→Law, social and behavioral sciences→ Economics; Psychology

Keywords

Machine Learning; Draft; Basketball; Nba; Algorithm.

1. INTRODUCTION

NBA basketball operations front offices have numerous responsibilities including “coach selection, scouting, [salary cap management and] contract negotiations and perhaps most importantly, player personnel decisions.” (Wong and Deubert, 2011). This paper essentially automates the player personnel decisions. Such an ambitious project was impossible until recently, but due to advances in statistical techniques, cloud computation, high quality data, and an integrated, backtested, visual design, it is now one more thing, alongside chess and Jeopardy!, that artificial intelligence does better than human intelligence.

I provide an automated, algorithmic system to handle all of the decisions of which players to draft, trade, or sign. The system does not overfit information from the future. It does not only project players with future professional experience, but rather all top prospects, so it does not have lookahead bias. It is flexible enough to be customized to a team’s preferred on-court NBA performance metric. It outperforms historical human decision making in backtesting. It can be used either to entirely automate such decision making, or to provide a default evaluation framework. As such, this objective, testable, systematic standard can help bring discipline and accountability to these issues of central importance in team success. The framework can also be used to evaluate alternative systematic strategies historically. In short, the automated general manager can provide a benchmark for owners to use in evaluating their own front offices as well as a tool for the front offices to simplify most decisions, allowing them to add value in the more difficult situations.

The three parts of the system are each dealt with separately in the following sections: the draft, free agency, and trades.

One issue in backtesting is ensuring a comprehensive past history of available alternatives. The draft is the cleanest application because it is clear what other prospects were available instead; thus, for drafts, both a fully automated system and interactive tools are provided.

Trades are the most difficult application because there is no data on what teams would have been willing to trade players for in the past; thus, for trades, interactive tools allowing live evaluation of actual trade proposals are the most useful.

Free agency is in the middle because it is a workable but far from perfect assumption that a free agent who later signed a particular deal would have agreed to a more lucrative deal had it been offered to him—indeed, the larger the contract, the less plausible this assumption becomes; thus, for free agency, an automated system for lower salaried players plus interactive tools are available.

Finally, tools involving combinations of these three areas are also provided. The entire system and all of the tools are available for free at nbagm.pm.

Far from being just another input, advanced analytics can be an additional revenue source in a team’s business model. As far as I am aware, no team currently employs an approach like this; the first one to do so will likely reap the vast majority of the available gains.

I predict systems like this will eventually become standard among the top NBA teams.

2. DATA AND METHODS

I project historical NCAA college basketball performance to subsequent NBA performance for prospects using an extension to machine learning techniques that prevents snooping bias without limiting the data available. I find that the projections would have improved the drafting decisions of every team: over the past ten years, teams forfeited an average of $130,000,000 in lost productivity relative to the recommendations of the automated GM, corresponding to about one extra win per season per draft pick. The automated choices had double the productivity of the historical human choices. I provide team-by-team comparisons, summaries of lost profit, sensitivity evaluations, and other interactive tools.

NCAA and NBA data was scraped from public websites including realgm.com and basketball-reference.com, combine measurements and RSCI from draftexpress.com, salary data from Spotrac, and mock drafts from ESPN, nbadraft.net, draftexpress.com, mynbadraft.com, and CNN/SI. NBA Wins Produced numbers are courtesy of David Berri of wagesofwins.com and, for the 2013-2014 and 2014-2015 seasons, from boxscoregeeks.com.

A win is worth about $1.65 million (c.f. Silver, 2014) because about $2 billion is spent in total on player salaries annually and there are only 1,230 wins available per season. (These numbers will increase in subsequent years with the influx of new revenue to the league. However, for purposes of comparing to the past decade, the $1.65 million value is appropriate.)

I use the random forests machine learning algorithm of Breiman (2001). Specifically, I used the `Predict[]` function of Mathematica v10 with the option `Method->"RandomForest"`. This automatically chooses reasonable parameters: in our case it chose `"TreeNumber"->200` and `"LeafSize"->5`. The results are also similar with neural networks.

The draft model used in this paper projects professional performance from both on-court and off-court prospect information using an extension to machine learning described below. The data used for projections include both the usual and advanced quantitative on-court college performance information, as well as proxies for qualitative off-court prospect information, such as consensus rankings after high school, appearance in various mock drafts, awards won, combine measurements, and more.

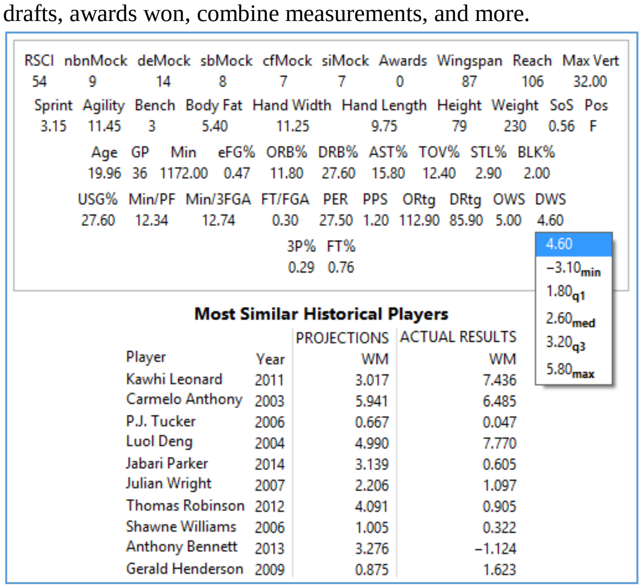

RSCI	nbnMock	deMock	sbMock	cfMock	siMock	Awards	Wingspan	Reach	Max Vert
54	9	14	8	7	7	0	87	106	32.00

Sprint	Agility	Bench	Body Fat	Hand Width	Hand Length	Height	Weight	SoS	Pos
3.15	11.45	3	5.40	11.25	9.75	79	230	0.56	F

Age	GP	Min	eFG%	ORB%	DRB%	AST%	TOV%	STL%	BLK%
19.96	36	1172.00	0.47	11.80	27.60	15.80	12.40	2.90	2.00

USG%	Min/PF	Min/3FGA	FT/FGA	PER	PPS	ORtg	DRtg	OWS	DWS
27.60	12.34	12.74	0.30	27.50	1.20	112.90	85.90	5.00	4.60

3P%	FT%
0.29	0.76

4.60
-3.10_{min}
1.80_{q1}
2.60_{med}
3.20_{q3}
5.80_{max}

Most Similar Historical Players

		PROJECTIONS	ACTUAL RESULTS
Player	Year	WM	WM
Kawhi Leonard	2011	3.017	7.436
Carmelo Anthony	2003	5.941	6.485
P.J. Tucker	2006	0.667	0.047
Luol Deng	2004	4.990	7.770
Jabari Parker	2014	3.139	0.605
Julian Wright	2007	2.206	1.097
Thomas Robinson	2012	4.091	0.905
Shawne Williams	2006	1.005	0.322
Anthony Bennett	2013	3.276	−1.124
Gerald Henderson	2009	0.875	1.623

Figure 1: Prospect evaluation tool

Figure 1 shows a one-off evaluation tool for exploring the sensitivity of a player's ratings. It is for Kawhi Leonard, who was picked #15 in 2011 but which the automated GM would have drafted as high as #5. It shows all of the inputs used, the most similar historical players, and a dropdown allowing the user to see how Leonard's rating would change if his collegiate DWS were changed. Similar dropdowns are available for every input, allowing a human GM to evaluate sensitivities as well as to override the data if new information becomes available.

The extension to the machine learning (ML) approach that allows for the maximum amount of data to be used for each player is to re-run the ML algorithm for every player, excluding that player from the corpus. Thus, rather than running just one ML equation on a single training subset, we re-run for each player. This assumes that the relationship between college and pro performance is independent of time. This assumption is warranted for the data analyzed here: the correlation between the predictions of the overall ML and those of an overall ML including the player's draft season is 0.99. In other words, including the season a player is drafted as part of the input basically does not matter.

The downside to this approach is time: one ML run takes about 4 seconds on a single core of an Intel i7 CPU at 2.69GHz. Thus, a typical level 2 machine learning draft model could be completed in 4 seconds. But this deeper analysis requires 30 minutes. Further, that half hour is for each choice of metric for player production evaluation. For robustness and flexibility, I use 11 separate metrics of future production for NBA prospects, each of which calculates the average over the player's first three seasons in the NBA: minutes played (MP), wins produced (WP) and wins produced per 48 minutes (WP48), win shares (WS) and win shares per 48 minutes (WS48), offensive win shares (OWS), defensive win shares (DWS), Hollinger's player efficiency ranking (PER) and the associated estimated wins added (EWA), and a simple consensus wins metric (WM)—the average of WP, WS, and EWA—as well as the wins metric per 48 minutes (WM48). See Berri (2014), Sports Reference (2014), and Hollinger (2009) for relevant definitions. Thus, the total computing time on a single core is about 6 hours.

3. RESULTS AND DISCUSSION

3.1 Draft

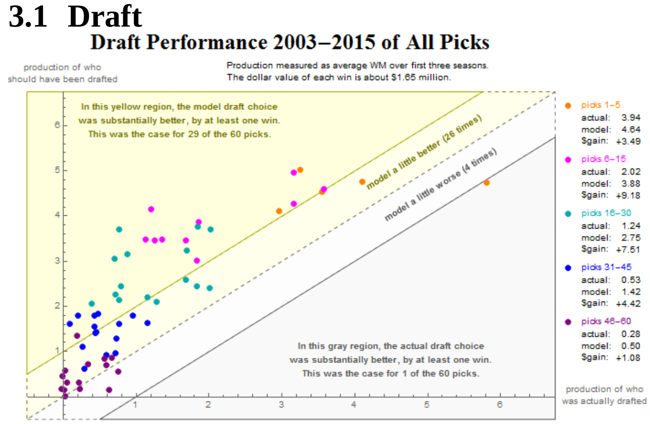

Figure 2: Comparison of automated GM model draft choice vs. actual historical draft choice

Figure 2 shows how the automated GM's choices performed (y-axis) relative to the actual draft choices made by teams (x-axis). Performance is measured with the consensus wins metric WM. The automated GM would have generated far more production than the actual draft choices for all picks, except for the first overall pick (solely because LeBron James was not available as an option to the system). For example, for picks 6 through 15, the actual choices missed out on an extra 1.86 wins per year, or about $9.18 million in value over three years.

These results are conservative because the model only evaluates NCAA players while actual choices often drafted foreign or high school players. Despite this informational handicap, the system still substantially outperformed human decisions.

On a team-by-team basis, every single team lost out on substantial on-court productivity relative to what they could have had if the automated GM had been available to them. Table 1 lists the lost profits per team; the average was $132 million and the largest was nearly $300 million by Minnesota. Estimated team values are by Forbes as of January 2014 from Badenhausen et al. (2014) (prior to the recent sales of the Clippers and Bucks at higher valuations). The average team lost out on profits equal to 23 percent of its 2014 estimated market value; Minnesota lost out on a whopping two-thirds. The model choice would have produced about one more win per draft pick per year, essentially double the human choice.

Note that the lost profits or win gains in Table 1 cannot be aggregated. This is not a general equilibrium model aiming to make all teams more optimal but rather a partial equilibrium model quantifying the benefit available to whichever single team does use the automated GM. The lost profits are perhaps best viewed as a convenient translation from wins into dollars.

Note also that the draft strategy does not take into account optionality or the cost of acquiring additional later draft picks. Incorporating this optionality improves the results even further.

Table 1: Team-by-team draft performance and lost profits relative to the automated GM

	Team	Total Picks	Actual Choice Avg. Production	Model Choice Avg. Production	Model – Actual Avg. Production	↓$ Gains from Model (total lost profits)	Team Value Forbes (2014)	Portion of Team Value
1	MIN	33	0.69	2.51	1.83	$298,145,869	$430,000,000	69.34%
2	UTH	31	1.15	2.81	1.66	$254,059,117	$525,000,000	48.39%
3	GOS	25	1.22	3.18	1.96	$242,655,844	$750,000,000	32.35%
4	POR	33	1.33	2.73	1.40	$229,237,247	$587,000,000	39.05%
5	OKC	35	1.31	2.46	1.14	$198,061,593	$590,000,000	33.57%
6	TOR	23	1.22	2.77	1.55	$176,038,302	$520,000,000	33.85%
7	MEM	27	1.00	2.30	1.29	$172,921,324	$453,000,000	38.17%
8	CHA	20	1.61	3.22	1.61	$159,310,473	$410,000,000	38.86%
9	BRK	26	1.00	2.23	1.23	$158,257,137	$780,000,000	20.29%
10	DAL	18	0.74	2.44	1.70	$151,330,360	$765,000,000	19.78%
11	LAC	21	1.14	2.55	1.41	$146,792,234	$575,000,000	25.53%
12	NYK	28	1.30	2.33	1.03	$142,790,963	$1,400,000,000	10.20%
13	PHX	21	0.37	1.74	1.36	$141,818,473	$565,000,000	25.10%
14	WAS	21	1.00	2.35	1.35	$140,061,716	$485,000,000	28.88%
15	LAL	26	0.49	1.52	1.03	$133,107,848	$1,350,000,000	9.86%
16	PHL	34	1.37	2.06	0.70	$117,674,421	$469,000,000	25.09%
17	ATL	27	1.25	2.08	0.83	$110,820,200	$425,000,000	26.08%
18	IND	19	1.37	2.52	1.15	$108,296,306	$475,000,000	22.80%
19	DET	32	1.13	1.79	0.66	$103,809,392	$450,000,000	23.07%
20	CLE	29	1.35	2.06	0.71	$102,396,676	$515,000,000	19.88%
21	ORL	26	1.20	1.92	0.73	$93,554,737	$560,000,000	16.71%
22	BOS	33	1.12	1.66	0.54	$89,009,119	$875,000,000	10.17%
23	MIL	25	1.39	2.06	0.67	$82,750,582	$405,000,000	20.43%
24	MIA	20	1.01	1.84	0.83	$82,344,869	$770,000,000	10.69%
25	SAC	22	1.39	2.06	0.67	$73,208,652	$550,000,000	13.31%
26	SAN	27	0.71	1.26	0.55	$73,132,757	$660,000,000	11.08%
27	HOU	29	1.12	1.58	0.46	$66,056,615	$775,000,000	8.52%
28	CHI	25	2.16	2.55	0.39	$48,352,132	$1,000,000,000	4.84%
29	NOP	21	2.12	2.55	0.43	$45,106,627	$420,000,000	10.74%
30	DEN	20	1.21	1.32	0.11	$10,754,119	$495,000,000	2.17%

3.2 Free Agency

The first step in free agency, and player performance generally, is to look at the difference between the dollar value of their on-court production and their annual salary. This would explains who outperformed their contract in the past.

The next step is to forecast future performance to evaluate appropriate salary amounts and contract length. This is done with a machine learning model on past historical data: forecasting future wins made given current age, height, weight, years of service, and their maximum and most recent minutes played, EWA, WP, and similar advanced box score statistics. Projections are done up to four years out.

Figure 3 shows these projections. Not shown here but computed separately are forecasts of likely market salaries based on past contracts. The automated GM could combine that, or other indications of market range, to determine which free agents to pursue most heavily.

An underappreciated consideration is that franchises must underpay players in order to win. If a team pays every player on its roster exactly what they produce on the court, and if their total salary is in line with the average team salary, then that team will likely have a 0.500 record.

For example, Anthony Davis was the second most productive player in the NBA last season (barely behind MVP Stephen Curry), producing approximately $30 million of on-court value, and the New Orleans Pelicans just signed him to a deal that will pay him an average salary of $29 million per year. Is this a good deal?

On the one hand, the conventional thinking is that if you can lock up a bona fide superstar, you do it. On the other hand, if he is producing approximately what he is being paid, then the Pelicans will have to find value in the rest of their roster. But on the third hand, and this is the important hand, the dollar value of wins in the league will increase substantially in the coming years, so that his production in dollars will far outstrip his salary. This is an excellent deal for New Orleans.

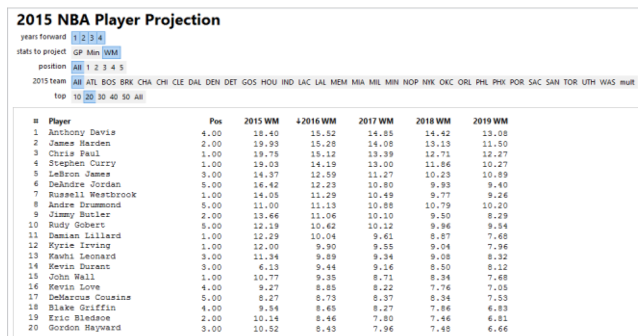

2015 NBA Player Projection

years forward 1 2 3 4
stats to project GP Min WM
position All 1 2 3 4 5
2015 team All ATL BOS BRK CHA CHI CLE DAL DEN DET GOS HOU IND LAC LAL MEM MIA MIL MIN NOP NYK OKC ORL PHL PHX POR SAC SAN TOR UTH WAS mult
top 10 20 30 40 50 All

#	Player	Pos	2015 WM	↓2016 WM	2017 WM	2018 WM	2019 WM
1	Anthony Davis	4.00	18.40	15.52	14.85	14.42	13.08
2	James Harden	2.00	19.93	15.28	14.08	13.13	11.50
3	Chris Paul	1.00	19.75	15.12	13.39	12.71	12.27
4	Stephen Curry	1.00	19.03	14.19	13.00	11.86	10.27
5	LeBron James	3.00	14.37	12.59	11.27	10.23	10.89
6	DeAndre Jordan	5.00	16.42	12.23	10.80	9.93	9.40
7	Russell Westbrook	1.00	14.05	11.29	10.49	9.77	9.26
8	Andre Drummond	5.00	11.00	11.13	10.88	10.79	10.20
9	Jimmy Butler	2.00	13.66	11.06	10.10	9.50	8.29
10	Rudy Gobert	5.00	12.19	10.62	10.12	9.96	9.54
11	Damian Lillard	1.00	12.29	10.04	9.61	8.87	7.68
12	Kyrie Irving	1.00	12.00	9.90	9.55	9.04	7.96
13	Kawhi Leonard	3.00	11.34	9.89	9.34	9.08	8.32
14	Kevin Durant	3.00	6.13	9.44	9.16	8.50	8.12
15	John Wall	1.00	10.77	9.35	8.71	8.34	7.68
16	Kevin Love	4.00	9.27	8.85	8.22	7.76	7.05
17	DeMarcus Cousins	5.00	8.27	8.73	8.37	8.34	7.53
18	Blake Griffin	4.00	9.54	8.65	8.27	7.86	6.83
19	Eric Bledsoe	2.00	10.14	8.46	7.80	7.46	6.81
20	Gordon Hayward	3.00	10.52	8.43	7.96	7.48	6.66

Figure 3: The top 10 NBA players with the highest forecast wins made for the 2015-2016 season

3.3 Trades

While there is a history of consummated trades between NBA teams, there is no comprehensive public history of offered but unaccepted trades, or, even more, uncontemplated trades that would have been agreed to (see Maymin, Maymin, and Shen (2013) for examples of possible mutually beneficial trades).

Further, the complete evaluation of a trade needs to include: the value of future draft picks, including options and protection on those picks; salary cap impact both in the short and long term; and considerations on future market salaries and extension possibilities for acquired assets. Thus, trades cannot be automated. However, tools incorporating backtested results can help a human GM quickly and accurately determine the near-term basketball value of a proposed trade: will the team be better off over the next few seasons?

Player on-court productivity can be projected in the same way as free agents in the previous section, and the totals compared for the outgoing vs. incoming packages, allowing the GM to decide if the difference in actual and projected salary offsets the difference in projected productivity. However, the evaluation of draft picks needs to be addressed.

We have used the wins made (WM) metric for evaluating the production of players, but how stable is it? In general, how stable are metrics like minutes played or field goal percentage? We can look at the correlation within players across years: the autocorrelation of WM is 0.76. This is roughly in line with the other win metrics, and is higher than, for example, the correlation of minutes played, which is only 0.68.

The disadvantage of WM is that it is not derived from a single theoretical framework; the advantage is that, as an average and a "consensus" measure, it is less prone to extreme valuation. Of course, it is a simple extension to use any metric a team prefers.

We can use the wins made metric and the trade tools incorporating its forecasts to evaluate, for example, the three-team Kevin Love trade consummated in the summer of 2014, as described for example in Zillgitt (2014).

Cleveland is the overall winner in terms of pure on-court basketball value, even though they only received one of the six players involved in the trades, and zero draft picks. But that player is Kevin Love, which makes all the difference (at least in forecast; his actual untimely injury was unpredictable).

Minnesota was forecast to lose relative to what they would have had if they had been able to retain Kevin Love, but what was not included in these trade evaluations, but every human GM knew, was that there was virtually no way to retain Love beyond one lame duck season. To be sure, Bennett's forecasts are based off a historically terrible rookie year; if he were to perform as an average top overall pick, then Minnesota would have fully replaced Love's lost production.

Philadelphia, meanwhile, was forecast to lose more games the following season as a result of the trade, but to do better in later years than they would have had they retained Young. Note that the future "2015" pick traded from Cleveland ultimately to Philadelphia was a top-10 protected Miami pick; indeed, Miami ended up with the tenth selection and kept its pick, thus the pick to be transferred will actually be for 2016, c.f. Rapaport (2014).

4. CONCLUSION

In this paper I introduced an extension to traditional machine learning methods that can have broader applicability outside of sports. The idea is straightforward although the implementation can be time-consuming: for each data point, run a complete machine learning fit using the entire data set except that single data point. This method requires the assumption that the relationship between a given player's college and professional production is relatively independent as a function of time. In other words, if a player had been born and drafted a few years earlier or later, had the same pre-draft inputs, and was drafted onto the same team with the same teammates and against the same opponents, his professional career should have been roughly similar. For the sample period and the context of this paper, this assumption is reasonable.

A traditional machine learning application uses early data solely as a training set, therefore backtesting over that same period is not proper, and it also does not update itself from later results from the test set, therefore it is noisier. The method used here overcomes both of those obstacles and allows for a proper and reliable backtest across the entire sample period without overfitting from future data.

I find that the average team could have earned $130 million in additional on-court production had they used the draft component of the automated GM system presented here. These lost profits average about 23 percent of a team's total franchise market value. I also present free agency and trade components to help a team make more disciplined algorithmic decisions in those areas as well. The interactive tools are available for free on nbagm.pm.

As the first wave of analytics focused on standalone statistical models for drafting and player valuation, the next wave will now likely focus on the comprehensive backtesting and interactive evaluation of entire algorithmic strategies. This will require building on the successes of the first wave with new systems and tools such as the ones presented here.

Future extensions for the draft component could incorporate more prospects including both high school and foreign players, textual sentiment analysis of scouting reports if sufficient historical data is available, and alternative measures of future production such as wins metrics based on optical or other fine-grained data. The free agent component can be extended to the development league, including a predictive model of which characteristics of players suggest they will be successful if called up. The trade component can attempt to fit a model of decision making onto the transactions of every general manager or team, and the results of those fitted models can be used to predict future trade interest.

5. REFERENCES

Badenhausen, Kurt, et al., (January 22, 2014). "NBA Team Values", Forbes, retrieved on December 10, 2014 from http://www.forbes.com/nba-valuations/list.

Berri, Dave (2011). "Wins Produced comes back better and stronger!", WagesOfWins.com, retrieved on December 9, 2014 from http://wagesofwins.com/2011/12/11/wins-produced-comes-back-better-and-stronger/.

Breiman, Leo (2001). "Random Forests," *Machine Learning* 45:1, 5-32.

Hollinger, John (2009). "PER Diem: Value Added", ESPN, retrieved on December 9, 2014, http://sports.espn.go.com/nba/columns/story?columnist=hollinger_john&page=PERDiem-090325.

Maymin, Allan, Philip Maymin, and Eugene Shen (2013), "NBA Chemistry: Positive and Negative Synergies in Basketball," International Journal of Computer Science in Sport, vol. 12, no. 2, pp. 4-23.

Rappaport, Max (2014). "Sixers Acquire 2015 First-round Pick, Two Players", Sixers.com, retrieved on December 11, 2014 from http://www.nba.com/sixers/news/140823-trade-story.

Silver, Nate (May 20, 2014), "How Much Is Winning the (NBA Draft) Lottery Really Worth?," FiveThirtyEight, retrieved on December 9, 2014 from http://fivethirtyeight.com/features/how-much-is-winning-the-nba-draft-lottery-really-worth.

Sports Reference LLC, "Calculating Win Shares," Basketball-Reference.com, retrieved on December 9, 2014 from http://www.basketball-reference.com/about/ws.html.

Wong, Glenn M. and Chris Deubert (2011), "National Basketball Association General Managers: An Analysis of the Responsibilities, Qualifications and Characteristics," Jeffrey S. Moorad Sports Law Journal, vol. 18, no. 1, pp. 213-266.

Zillgitt, Jeff (2014). "Three-team Kevin Love trade finally becomes official", USA Today, August 23, 2014, retrieved on December 10, 2014 from http:// usatoday.com/story/sports/nba/2014/08/23/14442179