Walker Szurek

Stats 50: The Mathematics of Sports

Project Abstract

2 December 2014

Assessing NBA General Manager Performance

The purpose of this project is to determine a methodology for assessing general manager performance in the National Basketball Association. Ultimately, this project seeks to answer two questions. First, do general managers have a meaningful impact on their teams with free agent signings and trades? Second, are some general managers better than others? Basketball presents an interesting case study for front office evaluation because it is sport in which individual performance can be difficult to predict. Lacking the clearly defined sabermetrics of baseball, basketball GMs are thus required to effectively assess talent, fit, and overall player value when building a team. This project is valuable because very little analysis exists which quantitatively evaluates general managers. Quick research revealed only qualitative General Manager evaluations such as sensationalist power rankings and non-analytical gut-feeling based lists. At first glance, it could be tempting to oversimplify the analysis of a general manager's performance by evaluating a GM based solely on the team's championships, wins, or even year over year change in wins. While these metrics, especially the change in wins from year to year, are clearly indicative of GM performance, I do not believe they tell the entire story. Championships, wins, and team improvement depend on more than GM performance: coaching and natural player development are large contributors to those achievements. The goal of this research is to determine a way to isolate the effect of a GM's actions on the incremental improvement of his team. In this research I evaluated free agent signings and trades to determine the effectiveness of a GM on improving his team.

The first step towards evaluating general managers was determining a way to quantify team improvement through roster management. In order to do this, I turned to the statistic player efficiency rating, or PER, created by basketball statistician and current vice president of basketball operations for the Memphis Grizzlies John Hollinger. PER attempts to combine all of a player's contributions on a basketball court into a single, per-minute statistic. PER is calculated by combining all the positive contributions a player makes on a basketball court such as made field goals, assists, rebounds, and blocks and then subtracting the negative contributions,

turnovers and missed shots, and finally dividing the total by the player's minutes played[1]. PER has been criticized for overvaluing the contributions of players who play a small number of minutes. For example, in the 2014-15 season, Dwight Powell, a rookie power forward out of Stanford, has only played 7 minutes and sports an astronomical a PER of 33.2: higher than LeBron James' current PER of 24.2. Dwight Powell is not a better basketball player than LeBron James. In order to eliminate the bias of PER towards players who play fewer minutes, I multiplied each player's PER by the number of minutes that player played, giving me a stat I will call Player Value, or PV. PV is a useful stat in the context of evaluating a GM because common sense tells us that a team's success will not depend the average PER of each player on the team because some players play much more than others. Team performance will depend on the total difference between positive and negative contributions a team has over the course of a season. I exported 97 team seasons[2] of PER data and calculated each team's total player value by summing the PER*MP for each player for each season. I regressed the results and found that Player Value is strongly correlated with a team's wins (Figure 1). I chose to regress PV against wins because wins are the ultimate metric which determines team quality at the end of a season. A statistical analysis showed that team wins and PV have a correlation coefficient of 0.66, and a t-test of the correlation coefficient showed that the results have a statistical significance of >99%. Therefore, the relationship between a team's total player value and the number of regular season games that team will win is very strong, and PV is thus a valid metric for evaluating the quality of an NBA team.

With a baseline metric for evaluating team quality, the next step was to determine the parameters for my research. The current average tenure of an NBA GM is 5 years, and the current median tenure is 2.5 years. Therefore, in order to obtain complete data for analysis I extended my timeframe for analysis back to 2000 and included any GM with a tenure of 5 or more seasons with the same team in my candidate pool for analysis. There were 30 candidates for analysis. In order to gather the data for each GM, I defined the time from the end of one NBA finals until the next year's NBA finals as a GM's season. This definition is different than the NBA's definition, which defines a season as July 1st of one year to June 30th of the next. It is important that I adjusted the definition of a season otherwise trades occurring after the NBA

1" More detailed information on PER can be found at http://www.basketball-reference.com/about/per.html

2" A team season is the data for one team over one season, e.g. the 2013-2014 Chicago Bulls season

finals, such as draft day trades, which take place in late June, would be counted towards the wrong season. Next, I used the transaction lists from Basketball-Reference.com, which include the transaction type (free agent signing or trade), the date, and the players involved. For each transaction I calculated the PV for each player joining the team of the GM being evaluated. I separated each player's PV into the PV accumulated over the player's first season with the team and the total PV the player generated during his time with the GM's team to give myself multiple avenues for analysis. I chose to only look at incoming players to the team because trying to calculate each transaction's ΔPV would have introduced a very large amount of uncertainty to my analysis. I would have had to arbitrarily chosen whether to evaluate an outbound player's PV based on his time with the GM's team or with his next team after the transaction. This process would have brought factors such as coaching and system into play. Because a GM's job is to evaluate how a player will perform on his team and not how an outbound player will perform after the transaction with another team and based on my later results, I felt that evaluating only the incoming player's PV adequately captured the quality of the transaction. I collected data up to the most recent complete NBA season. With this methodology in place, I collected 97 team seasons of transaction based PV data which made up the tenures of 13 general managers.

The next step in my research was to determine how to use PV to evaluate a general manager's performance. The first thing I looked at was how the average player value a general manager added via trades and free agency each year (Player Value Added, or PVA) compared to the average wins of the general manager's team over his tenure (Figure 2). For this analysis I used the total player value a transaction produce for the team, not just the PV produced in the first year. The results did not show a meaningful relationship between average PVA and average wins. The two values had a correlation coefficient of -0.32 with a statistical significance of <75%. Therefore, the average player value a GM produces per year via free agent signings and trades does not have a strong relationship with team success. There could be numerous factors contributing to this result. First, many of the best NBA players are acquired via the draft, not trades or free agency; therefore, the totally quality of a team relies on much more than the GM's non-draft actions. Second, PVA does not account for the quality of a team's coaching. Elements of team success such as defensive performance and offensive scheme are not captured by PER or PV, yet they are crucial to a good team and are often credited to the quality of coaching a team

receives. Therefore, PVA is not a good metric of performance by which to measure a GM's performance.

Next, I regressed the one year PV a general manager produced against the teams change in wins relative to the previous year. The purpose of this analysis is to determine if a team's incremental change in wins versus losses is tied to the player value they add for that season only. However, a graph of the change in wins relative to the previous year vs the one year player value added showed no significant relationship over the 97 team seasons analyzed (Figure 3). The change in wins and one year PVA generated by a GM had a correlation of 0.09 and a significance of <75%. One major problem evident in relating the change in wins to the player value added for that season is that teams are all of different qualities. Adding a large amount of one year player value may dramatically change a bad team, but a good team may stay relatively similar. Thus, it was necessary to take the next step and connect a GM's PVA to the overall performance and incremental improvement of his team.

In order to calculate the change in overall team quality due to a GM's actions, I compared the relative value of the one year PVA generated by a GM to the value it would have had on the previous year's team. If the player value added by a GM accounts for a higher percentage of the previous year's total player value than the current year's total player value, the team has improved due to the trades and free agents signings made by the GM. In order to quantify this relationship I calculated a statistic I am calling the change in player value percentage, or ΔPV%. It is the percentage of a team's PV in year n-1 (the previous season) created by the one year PVA minus the percentage of a team's PV in year n created by the one year PVA. The formula for ΔPV% is:

Looking at a graph of the 97 seasons analyzed as a part of this project, it is clear the ΔPV% is strongly correlated to the incremental change in wins from the previous season (Figure 4). The change in wins and the change in player value percent had a correlation coefficient of 0.62 with a statistical significance of >99%. Therefore, the change in player value relative to the team's total value due to a GM's actions is a strong indicator or the change in said team's wins. Thus, we can conclude that the actions of a general manager are very important in regards to his team's success.

Finally, since the last goal of this project was to determine if certain GMs are better than others, I aggregated the data and averaged the ΔPV% and ΔWins over the course of each GM's tenure and found a statistically significant relationship (Figure 5). We can thus say that over time some GM's are better than others at improving their teams through trades. Aggregated ΔPV% and ΔWins have a correlation coefficient of 0.64 with a statistical significance of >98%. Therefore, evaluating the ΔPV% a GM produces is a meaningful statistical indicator of GM ability and a GM's effect on team improvement. Table 1 shows the values of ΔPV% and ΔWins for the 13 general managers analyzed in this research:

GM	ΔWins	ΔPV%
Bryan Colangelo (Raptors 2006-2012)	1.00	1.15
Gar Forman (Bulls 2009-P.)	1.40	1.04
Danny Ainge (Celtics 2003-P.)	-1.73	0.46
John Paxson (Bulls 2003-2008)	1.83	0.15
Billy Knight (Hawks 2003-2007)	0.40	0.10
Pat Riley (Heat 1995-P.)	0.14	0.03
Joe Dumars (Pistons 2000-2013)	-0.93	-0.09
Billy King (76ers 1998-2006)	-2.00	-0.17
Isiah Thomas (Knicks 2003-2007)	-2.80	-0.25
John Hammond (Bucks 2008-P.)	-1.83	-0.52
David Kahn (Pacers 1995-2004)	-2.40	-0.61
Otis Smith (Magic 2005-2011)	0.14	-0.63
Larry Harris (Bucks 2003-2007)	-3.20	-0.63

The names on this list go a long way towards qualitatively validating the results of this research. The top performing general managers in terms of ΔPV% are names we consider to be good general managers such as Danny Ainge and Pat Riley, while famous GM busts such as Otis Smith and Isiah Thomas are in the bottom half of this list. Therefore, based on the quantitative and qualitative results of this research, I believe this project adequately shows that general managers do have a meaningful impact on the improvement of their team's year over year and that some general managers are in fact statistically better than others beyond merely analyzing wins, losses, and championships. This research and the statistic ΔPV% could be a useful tool for owners and fans to evaluate the performance of their teams and there could be meaningful potential for further analysis using Player Value to analyze coaching, player evaluation, and even informing general manager decisions.