The Pot Calling the Kettle Black
Are NBA Statistical Models More Irrational than "Irrational" Decision-Makers?

December 2007

David Lewin
Macalester College and Cleveland Cavaliers

Dan T. Rosenbaum
University of North Carolina at Greensboro and Cleveland Cavaliers
rosenbaum@uncg.edu

Abstract:
Recent research suggests that "statistics and analysis" typically lead to better decisions than "intuition and human intellect" in diverse areas such as choosing which students to admit to college and assessing mortality risks among cancer patients. Sports economics, with its rich data and abundant decisions to analyze has provided a fertile laboratory for studies of the efficiency of decision-making. In fact, researchers of the National Basketball Association (NBA) have used statistical models of player productivity to make strong claims about the "rationality" of NBA decision-makers. Yet these statistical models have rarely been subjected to any rigorous examination of their ability to forecast the future. We examine how well several player productivity metrics, including (a) John Hollinger's Player Efficiency Rating, (b) *Wages of Wins* Wins Produced, and (c) the NBA Efficiency metric, do in predicting future team wins and future player productivity (the latter as measured by plus/minus statistics). In addition to a comprehensive examination of the player productivity metrics used by NBA statistical analysts, this paper is the first academic presentation of plus/minus statistics. Our findings provide a counterweight to much of prevailing literature and suggest that models that assume simplistic NBA decision-making often outperform more sophisticated statistical models.

Acknowledgments:
We would like to thank the APBRmetrics community for years of commentary and critique of the ideas contained in this paper. Many of the ideas presented in this paper first appeared in the APBRmetrics discussion board. We thank Guy Molyneux, Stephen Iliardi, Eli Vitus, and seminar participants at the NESSIS conference and at the Council of Economic Advisers for their helpful comments. We also would like to make it clear that this paper represents the views of the authors alone and not the organizations that the authors are associated with.

The Pot Calling the Kettle Black
Are NBA Statistical Models More Irrational than "Irrational" Decision-Makers?

1. Introduction

A recurring theme spawned by the growth of statistical analysis in sports decision-making is the tension between "statistics and analysis" and "intuition and human instinct." It permeates Michael Lewis's *Moneyball* which describes the use of statistical analysis by the low-payroll, yet highly successful Oakland A's, as well as Malcolm Gladwell's observation in his *Wages of Wins* review that in "complex situations, like basketball, the limitations of 'seeing' become enormous." David Leonhardt joins the chorus both in his sports and economics writings in the *New York Times*.

> "Academic research, however, is pretty much on the side of statistics. Whether diagnosing patients or evaluating job candidates, human beings vastly overestimate their ability to make judgments, research shows. Numbers and analysis almost always make people better. 'There have been hundreds of papers on subjects from picking students for a school to predicting the survival of cancer patients,' said Richard Thaler, a University of Chicago economist who uses sports examples in his class on decision-making. When a computer model is given the same information as an expert, the model almost always comes out on top, Thaler said." (Leonhardt, 2005)

This last sentence, however, begs the question, because it is rare that the "computer model" ever has the "same information" as the scout or coach or doctor. The more typical tradeoff is whether the benefits of more data points (often collected and analyzed in a more systematic manner) outweigh the costs of a simplified model that ignores important aspects of reality. Thaler argues above that in most circumstances the answer appears to be yes – an argument that is supported in baseball due to the success of Billy Beane and other sabermetricians.

But in this paper we analyze basketball and see that the answer is not quite so clear. Models put forth by leading academics and statistical analysts often perform quite poorly relative to simple approximations of NBA decision-makers. Stanley (2005) asserts that this difference between baseball and basketball may be due to the fact that "baseball is mostly about a small number of repetitive hand/eye coordination tasks, while basketball involves constant maximizing interaction between optimizing actors on the court." Tabulating statistics may very well be the best way to form predictions about the "repetitive hand/eye coordinating tasks" of baseball, but applying those same techniques to the game of basketball which "involves constant maximizing interaction between optimizing actors" may leave out too much. The costs of a simplified model may be too high.

The most vociferous argument against this possibility comes from Berri et al. (2006) in the *Wages of Wins* and its related blog, "our story of the overrated and underrated indicates that the NBA may have a problem evaluating talent. The overrated players can all score, and most of these players have also scored major paydays. . . All of this suggests that people making decisions in the NBA are not as "rational" as economists tend to expect." (p. 199).

Rosenbaum (2005b) makes a similar argument that "people in the game often claim to know instinctively how to measure intangibles, but salaries suggest otherwise. Teams pay for little more than the glory statistics (points, rebounds and, to a lesser extent, assists). . . Although steals, blocks, shooting percentage and an ability to avoid turnovers are crucial to a team's performance, players proficient in these aspects are rarely rewarded with bigger paychecks."

What is curious about these arguments, in particular given they are made by statistical analysts, is that there is very little evidence that basketball statistical models are better at evaluating players than NBA decision-makers. The most common argument put forth (see both quotes above) is the low correlation between salaries and many of the assumed determinants of wins.[1]

This paper takes a different approach. We lay out simple approximations of NBA decision-making and then compare how these models do versus more sophisticated statistical models in (1) predicting future team wins and (2) explaining and predicting how teams play when a particular player is on the court. Somewhat shockingly, the simple approximations of NBA decision-making often perform better than their more sophisticated counterparts.

2. Statistical Measures of Player Productivity

Below we describe several player productivity measures, ranging from extremely simple to fairly complex. Some, like PER and Wins Produced, are "advanced" metrics that were developed specifically to measure overall player productivity. Others like Minutes per Game, Points per Game, and NBA Efficiency are simpler metrics that are used to approximate NBA decision-making. This section describes each in more detail.

The first metric is Minutes per Game for player *P* (MPG_P).

$$(1)\quad MPG_P = {MIN_P}/{G_P},$$

where MIN_P is total minutes played by player *P* and G_P is total games played by player *P*. The assumption here is that Minutes per Game is a proxy of player productivity *per minute*. In other words, a 40-minute-per-game player not only plays more minutes than a 30-minute-per-game player, but also is more productive per minutes played. Minutes per Game is a very simple measure of player productivity and theoretically is a good estimate of NBA decision-making. Assuming win-maximizing teams, minutes should be allocated in such a way that maximizes the probability of winning; hence minutes should be allocated roughly proportionally to player productivity within a team. However, decision-makers can only allocate minutes within their team, thereby making inter-team evaluations difficult. For this reason an adjustment is made by team that assumes that NBA decision-makers, on average, incorporate (past) team productivity into their assessments of players. Below we discuss these team adjustments in more detail.

The second metric is Points per Game for player *P* (PPG_P), which again is assumed to be a proxy of *per minute* productivity.

[1] The brief salary analysis in this paper suggests that NBA decision-makers do, in fact, consider more than just the points players score in their valuations of players.

$$(2)\ PPG_P = PTS_P / G_P = \left(MIN_P / G_P \right) \times \left(PTS_P / MIN_P \right),$$

where PTS_P is total points scored by player *P*. Points per Game is perhaps the most reported NBA statistic, in particular in the media. According to some, most notably Berri (2006), Points per Game is the metric of choice for NBA decision-makers and awards voters. This metric assumes that NBA decision-makers view points *per game* as a proxy of player productivity *per minute*. At first glance, such an assumption may seem odd, but it should be noted that points per game is equal to minutes per game multiplied by points per minute. Hence, it is minutes per game adjusted by a player's scoring rate. Like Minutes per Game, this metric is team adjusted in order to allow for cross-team comparisons.

The third metric is NBA efficiency for player *P* (EFF_P).

$$(3)\ EFF_P = \left(\frac{48}{MIN_P} \right) \times \left(PTS_P + REB_P + AST_P + STL_P + BLK_P - TO_P - MissFG_P - MissFT_P \right),$$

where REB_P is rebounds by player *P*, AST_P is assists by player *P*, STL_P is steals by player *P*, BLK_P is blocks by player *P*, TO_P is turnovers by player *P*, $MissFG_P$ is missed field goals by player *P*, and $MissFT_P$ is missed free throws by player *P*. NBA Efficiency, a simple linear weights metric, adds up all the good things that a player does on the court and subtracts the bad; all statistics are weighted equally. Rarely referenced in the media, NBA Efficiency is available on NBA.com and was recently reported that Utah Jazz owner Larry Miller utilizes a similar metric (Siler, 2007). Berri et al. (2006) argues that NBA Efficiency does a good job approximating NBA decision-making. Team and position adjustments are added to NBA Efficiency (see discussion below).

The fourth metric is John Hollinger's Player Efficiency Rating (PER) for player *P* (PER_P).[2]

[2] In addition to the statistics already defined, $3PM_P$ is three pointers made by player *P*, AST_T is assists by team *T*, FGM_T is field goals made by team *T*, FGM_P is field goals made by player *P*, FTM_P is free throws made by player *P*, PF_P is personal fouls by player *P*, FTM_L is league-wide free throws made, PF_L is league-wide personal fouls, ORB_P is offensive rebounds by player *P*, AST_L is league-wide assists, FGM_L is league-wide field goals made, FGA_L is league-wide field goals attempted, TO_L is league-wide turnovers, FTA_L is league-wide free throws attempted, ORB_L is league-wide offensive rebounds, DRB_L is league-wide defensive rebounds, TRB_L is league-wide total rebounds, $PACE_L$ is average league-wide possessions per game, and $PACE_T$ is average possessions per game for team *T*.

$$(4)\ \ uPER_P = \left(\frac{1}{MIN_P}\right) \times \left[\begin{array}{l} 3PM_P + \left(2/3\right) \times AST_P + \left(2 - FACTOR_L \times \left(AST_T / FGM_T\right)\right) \times FGM_P \\ + \left(1 - \left(1/6\right) \times \left(AST_T / FGM_T\right)\right) \times FTM_P - \left(FTM_L / PF_L\right) \times PF_P \end{array}\right]$$

$$+ \left(\frac{VOP_L}{MIN_P}\right) \times \left[\begin{array}{l} STL_P - TO_P + DRB\%_L \times \left(BLK_P + ORB_P - MissFG_P\right) \\ + \left(1 - DRB\%_L\right) \times DRB_P \\ -0.44 \times \left(\left(0.44 + 0.56 \times DRB\%_L\right) \times MissFT_P - \left(FTA_L / PF_L\right) \times PF_P\right) \end{array}\right],$$

$$where\ \ FACTOR_L = \left(2/3\right) - 0.25 \times \frac{AST_L \times FTM_L}{\left(FGM_L\right)^2},$$

$$VOP_L = \frac{PTS_L}{FGA_L + TO_L + 0.44 \times FTA_L - ORB_L},\ and$$

$$DRB\%_L = DRB_L / TRB_L\,.$$

$$PER_P = \left(PACE_L / PACE_T\right) \times \left(15 / uPER_L\right) \times uPER_P.$$

PER is a careful and detailed accounting for the contribution that each box score statistic makes in a possession. Popularized by several books (Hollinger, 2003, 2004, and 2005) and ESPN.com, PER is the most widely cited advanced player evaluation metric. PER most likely is used by several NBA teams and Berri (2006b) argues that it is a good approximation of NBA decision-making. In the later evaluations of metrics, team and position adjustments are added to PER (see discussion below).

The fifth metric is Wins Produced for player *P* (WP_P) from Berri et al. (2006), Berri et al. (2007), and Berri (2008).

$$(5)\ \ WP_P = \left(\frac{40}{MIN_P}\right) \times \left[\begin{array}{l} 0.033 \times PTS_P + 0.034 \times \left(TRB_P + STL_P - TO_P - FGA_P\right) \\ +0.022 \times AST_P + 0.021 \times BLK_P - \lambda_T PF_P - 0.01496 \times FTA_P \end{array}\right],$$

where λ_T is a team-specific multiplier for personal fouls. Wins Produced has become the player evaluation metric of choice among academic economists and for that reason we give it additional attention, including describing it in detail in the Appendix.[3] Wins Produced is derived from the relationships between wins, points, and possessions and uses a combination of theory and regression analysis to estimate the values of these weights. Team and position adjustments are

[3] Wins Produced is employed in Price and Wolfers (2007) and in papers that David Berri has co-authored with many leading sports economists, including Stacey Brook, Erick Eschker, Aju Fenn, Rodney Fort, Bernd Frick, Anthony Krautmann, Young Hoon Lee, Michael Leeds, Michael Mondello, and Martin Schmidt.

added to Wins Produced (see discussion below). A simplified version of Wins Produced is *Win Score* for player *P* (WS_P).

$$(6)\ \ WS_P = \left(\frac{40}{MIN_P}\right) \times \begin{bmatrix} PTS_P + TRB_P + STL_P - TO_P - FGA_P \\ +0.5 \times \left(AST_P + BLK_P - PF_P - FTA_P\right) \end{bmatrix}.$$

Win Score is not evaluated in this paper, because it is very highly correlated with Wins Produced, but it motivates another metric that will be evaluated: *Alternate Win Score* for player *P* (AWS_P).

$$(7)\ \ AWS_P = \left(\frac{40}{MIN_P}\right) \times \begin{bmatrix} PTS_P + STL_P - TO_P + 0.5 \times \left(AST_P + BLK_P - PF_P\right) \\ +0.7 \times \left(ORB_P - FGA_P\right) + 0.3 \times \left(DRB_P - FGM_P\right) \\ -0.35 \times FTA_P - 0.15 \times FTM_P \end{bmatrix}.$$

Alternate Win Score is essentially the Win Score from Berri et al. (2006) fit into the framework set out in Kubatko et al. (2007). As discussed in the Appendix, it assumes that the possession production function is the same for both the own team and opponent. Because about 70 percent of missed shots are rebounded by the opposing team, Alternate Win Score assumes that missed shots count 70 percent of a possession and defensive rebounds count 30 percent. A made shot counts 100 percent and offensive rebounds count 70 percent. Team and position adjustments are added to Alternate Win Score (see discussion below).

3. Position and Team Adjustments

Wins Produced incorporates a position adjustment. The position adjustment embodies the assumption that average productivity is the same across positions, an assumption that is difficult to justify theoretically. Theory would only imply that *marginal productivity* is the same across positions. Given the prominence of Wins Produced in the academic literature, we have chosen to add position adjustments to NBA Efficiency, PER, and Alternate Win Score. This allows for an apples-to-apples comparison between these metrics.[4]

Like in Berri et al. (2007) the position adjustments subtract the average value for each metric separately for guards (point guards and shooting guards), small forwards, and big men (power forwards and centers).[5] Given the arbitrariness of position designations (e.g. players often play different positions on offense and defense), well-designed metrics would capture productivity differences by position without resorting to position adjustments. In fact, the need for a large position adjustment could be a symptom of a poorly designed metric. In this light the very high correlations for PER (0.994) and Alternative Win Score (0.980) with and without position

[4] Position adjustments were not added to Minutes per Game or Points per Game in order to keep these metrics simple approximations of NBA decision-making.

[5] Unlike Berri et al. (2007) the positions used are just those given by either www.basketball-reference.com or www.dougstats.com. This implies that position-minutes are not necessarily equal across all teams. These differences are unlikely to affect the results significantly, in particular since the same positions are used for each metric that uses position adjustments.

adjustments are preferable to the low correlations for NBA Efficiency (0.918) and Wins Produced (0.874).[6] On the other hand, it is possible that some positions may be more productive, on average, than others in which case these position adjustments could mistakenly be equalizing productivity across positions.

Team adjustments are added to all of the metrics. Following Berri et al. (2006), the team adjustment forces the player metric to aggregate up by team to team efficiency – the difference between points per possession for own team and opponents. The team adjustment allocates the part of team efficiency that is not allocated through individual-level box score statistics to all players according to their minutes played.[7] Because aggregated box score statistics will differ from metric to metric, the team adjustment will differ from metric to metric. Yet they are all created in the same spirit as Wins Produced, because they all result in the player metrics aggregated by team to sum to team efficiency. This argument is presented in more detail in the Appendix.

Because the team adjustment forces each of the player metrics to add up to team efficiency, how well player metrics aggregated by team explain current team wins tells us nothing more than how well team efficiency explains team wins. Explaining current team wins cannot be used to evaluate player metrics; all player evaluation metrics with team adjustments will explain current team wins equally well (see Appendix for more explanation).

Team adjustments for each of the metrics can be justified using arguments similar to those in Berri et al. (2006), but that still begs the question why a team adjustment should be added to player evaluation metrics that are designed to approximate NBA decision-making, such as Minutes per Game and Points per Game. Table 1 provides such justification by showing that conditional on box score statistics, player salaries increase a highly statistically significant $0.9 million (in 2006-07 dollars) for every one standard deviation increase (about 4.8 points per 100 possessions) in team efficiency. So even after factoring in box score statistics such as points, rebounds, and assists, NBA decision-makers clearly take account of team-level success in assessing the value of players. That suggests that any player evaluation metric that attempts to approximate NBA decision-making should include team efficiency.

[6] These correlations are for 10,582 player-seasons between 1980-81 and 2006-07 and are weighted by minutes.
[7] Also, as described in the Appendix the team adjustment can also be used to subtract off any part of player productivity that is not directly related to team efficiency or possession production, as in Berri et al. (2006).

Table 1
The Effect of Box Score Statistics on Salaries

Statistic	Coefficient	Standard Error	Standard Deviation	Standardized Coefficient
Points per game	0.499	0.122	6.474	3.232
True shooting percentage	-0.130	0.030	5.378	-0.698
Rebounds per game	0.495	0.091	2.854	1.412
Assists per game	0.380	0.114	2.127	0.808
Turnovers per game	-0.598	0.657	0.835	-0.499
Blocks per game	1.379	0.311	0.678	0.935
Steals per game	-1.357	0.441	0.507	-0.688
Personal fouls per game	-0.724	0.288	0.731	-0.529
Team efficiency	0.183	0.038	4.835	0.886
Intercept	7.493	1.567	--	--

Notes: Data are from 7,269 player-seasons from 1990-91 through 2006-07. OLS regression is weighted by salary and minutes played and standard errors are calculated assuming that errors are clustered by player. R^2 is 0.3868. The dependent variable, salary, is expressed in millions of 2006-07 dollars and is normalized by dividing the raw salary by its league-wide mean each year. *True shooting percentage* is 0.5 times points divided by the sum of field goal attempts plus 0.44 times free throw attempts. *Team efficiency* is average team point differential measured per 100 possessions. *Standardized Coefficient* gives the estimated effect of a one standard deviation increase in the variable. See text for details.

Table 1 also suggests that, contrary to the statements made in the introduction, teams pay for more than just points. Points per game is certainly the most important box score statistic, but if the absolute values of the effects of one standard deviation increases for each of the independent variables are added up, the points per game effect is only one third of the total. Points, rebounds, and assists together only add up to only 56 percent, suggesting that nearly half of player salaries are explained by box score statistics other than points, rebounds, and assists. Moreover, even if box score statistics that are not statistically significant (turnovers) or have "wrong" signs (steals and shooting percentage) are excluded, points, rebounds, and assists account for less than two thirds of the total. This suggests that even with a team adjustment, Points per Game and Minutes per Game may leave out much of goes into determining player salaries. Hence, even if the team adjustment in a sense gives NBA decision-makers too much credit in assuming that they accurately value differences in team efficiency across teams, this regression evidence suggests that simple NBA decision-making models such as Minutes per Game or Points per Game leave out key components. In addition, these models ignore all of the non-box score statistics data – such as character, willingness to work hard, etc. – that go into player evaluation, so on net it is probably still the case that these simple models are lower bounds for how well NBA decision-makers actually evaluate players.

These statistically significant wrong signs might suggest that decision-makers are not optimally maximizing wins, but that assumes the model is specified correctly; it may be that the "wrong" signs are due to omitted variables bias. The somewhat low R^2 of 0.39 also might suggest that NBA decision-makers are "irrational." However, the salary caps, luxury taxes, rookie scale contracts, minimum and maximum salaries, and many other provisions in what is likely the most complicated collective bargaining agreement in sports (Coon, 2007) create markets for players that are sometimes competitive, sometimes monopolistic, and sometimes monopsonistic. Therefore, without carefully accounting for those significant barriers to a free market, it is very difficult econometrically to use R^2 to infer anything about the "rationality" of NBA decision-makers. A better approach to assessing the "rationality" of NBA decision-makers is to evaluate how well the proxies for NBA decision-making perform relative to more advanced player evaluation metrics.

Table 2 provides additional evidence that adding a team adjustment is consistent with NBA decision-making. In every case correlations with salary increase once team adjustments are added and this is particularly true with Points per Game.[8]

Table 2
Correlations of Salary with Player Metrics

Player Metric	Without Team Adjustment	With Team Adjustment	In Regression
Minutes per Game	0.464	0.468	0.497
Points per Game	0.536	0.561	0.557
NBA Efficiency	0.473	0.485	0.482
PER	0.497	0.505	0.506
Wins Produced	0.335	0.343	0.355
Alternate Win Score	0.420	0.433	0.431

Notes: Data are from 7,269 player-seasons from 1990-91 through 2006-07 and are weighted by salary and minutes played. The "without team adjustment" correlations do include position adjustments, where appropriate. The "In Regression" correlations are the square root of R^2 in a salary regression onto the player metric and team efficiency. See text for details.

Like with position adjustments, it is possible to calculate correlations of each metric with and without team adjustments. Higher correlations again suggest that the player metric itself does a better job capturing productivity without relying on a somewhat arbitrary team adjustment. PER (0.987), Wins Produced (0.979), Alternate Win Score (0.969), and NBA Efficiency (0.969) all have high correlations with and without team adjustments, while Points per Game (0.940) and especially Minutes per Game (0.730) have much lower correlations.

4. Correlation with Box Score Statistics

Section 2 and the Appendix carefully present the formulas behind each of the player evaluation metrics, but that sheds little light on how the metrics differ empirically. Table 3 addresses this issue by presenting correlations of each of the player evaluation metrics with box score statistics measured per minute and adjusted by position. These correlations illustrate important difference between the various metrics.

[8] Table 2 also presents the correlations of salary with the predicted values from a regression of the salary onto the player metric and team efficiency. When these correlations are much higher than the with team adjustment correlations, it suggests that the team adjustment may be over-correcting or under-correcting for team efficiency. In a few cases the with team adjustment correlation is higher than the regression-based predicted value correlation, which is due to the fact that the team adjustment subtracts off the indirect effect of the player metric on team efficiency.

Table 3
Correlations of Player Metrics with Box Score Statistics

Player Metric	TS%	Per Minute Statistics						Team Eff.
		Shots	REB	AST	TO	STL	BLK	
Minutes per Game (MPG) with team adjustment	0.452	0.345	0.139	0.246	0.060	0.075	0.058	0.679
Points per Game (PPG) with team adjustment	0.466	0.750	0.094	0.198	0.278	0.056	0.032	0.353
NBA Efficiency (EFF) with position and team adjustments	0.591	0.528	0.456	0.382	0.312	0.252	0.233	0.351
Player Efficiency Rating (PER) with position and team adjustments	0.613	0.668	0.282	0.332	0.290	0.225	0.156	0.286
Wins Produced (WP) with position and team adjustments	0.556	0.061	0.652	0.333	0.051	0.295	0.285	0.343
Alternate Win Score (AWS) with position and team adjustments	0.658	0.381	0.370	0.347	0.079	0.285	0.181	0.377

Notes: Data are from 10,582 player-seasons from 1980-81 through 2006-07 and are weighted by minutes played. *Shots* is field goal attempts plus 0.44 times free throw attempts. *REB* is rebounds, *AST* is assists, *TO* is turnvoers, *STL* is steals, and *BLK* is blocks. *Shots*, *REB*, *AST*, *TO*, *STL* and *BLK* are measured per minute. *TS%* is true shooting percentage, which is 0.5 times points divided by shots. *Team Eff.* is team efficiency, which is team point differential measured per possession. See text for details.

The largest difference between the player evaluation metrics is between how they treat players who take a lot of shots and players who get a lot of rebounds (relative to other players who play the same position). Points per Game, Minutes per Game and PER are all much more highly correlated with shots taken (including both field goals and free throws) than with rebounds. The ratio of the correlations for shots taken to rebounds is 8.0 for Points per Game, 2.5 for Minutes per Game, and 2.4 for PER. The ratio is closer to one for NBA Efficiency (1.2) and Alternate Win Score (1.1). Wins Produced is practically uncorrelated with shots taken and very highly correlated with rebounds; the corresponding correlation ratio is 0.1.

Each of the four advanced player evaluation metrics are more highly correlated with shooting efficiency (TS%) than are Minutes per Game and Points per Game, but among the advanced metrics, Wins Produced has the lowest correlation with shooting efficiency. None of the metrics have correlations with turnovers, blocks, steals, or personal fouls more than 0.32. Minutes per Game is much more highly correlated with team efficiency than are the other metrics. Unlike the other metrics, the Minutes per Game metric when aggregated up by team is uncorrelated with team efficiency (prior to adding the team adjustment). For that reason the team adjustment is much more important for Minutes per Game than for the other metrics; it is the only mechanism that allows Minutes per Game metric to value players from better teams more highly.

5. Evaluating the Player Evaluation Metrics using Team Wins

The previous sections present and describe a series of metrics, some of which may approximate NBA decision-making and other advanced metrics that allegedly are better in some way. But how should we go about evaluating these metrics? Oliver (2004) is skeptical about the very prospect of evaluating player metrics.

> "All we can do is create some statistics that *look like* individual wins and losses. There is then no great way to calibrate those statistics, improve them, or verify them. Without those reality checks to indicate what the method does wrong, there is no good way for us to say what form of an individual win-loss record 'works.'" (p. 243)

Berri et al. (2006) is more optimistic and proposes that a metric can be evaluated by how well it explains current team wins when current values of the metric are aggregated up by team. However, as discussed in the previous section and in the Appendix, all of the metrics with team adjustments add up to team efficiency so they do an equally good job of explaining current team wins. So current team wins cannot be used to distinguish between these metrics (for the current season).

This, however, leads to an obvious next step. How well do current values of these metrics explain future team wins? Every time players switch teams or experience significant changes in playing time, each of the player evaluation metrics predict a different change in team wins, creating thousands of mini-experiments that we can evaluate. Such evaluations have some problems (described later) and require assumptions to implement, but it is a reasonable place to start to evaluate a metric. To implement these evaluations, we use the following conventions.

1. Player productivity per minute is the same as in the previous season.
    - Points per game and minutes per game are assumed be predictors of *per minute* productivity.
    - This assumption abstracts away from the issue of how to project how productivity changes over a player's career, but since this is done for all of the metrics, it should not significantly affect the results.
2. All players who played less than 250 minutes in the previous season (including rookies) are assigned their future season values of the various metrics.
    - Predicting the productivity of rookies and low minutes players is beyond the scope of this paper and since all of these metrics are treated the same in this regard, this should not significantly affect the results. Note that this implies that these low minutes players and rookies will have little effect on the relative evaluations of these various metrics.
3. The assigned minutes played are those from the future season.
    - Again, predicting minutes per game is beyond the scope of this paper and since all of these metrics are treated the same in this regard, this should not significantly affect the results.

With these conventions it is straightforward to evaluate various metrics and we do so using data from the 1980-81 through 2006-07 seasons, which is a total of 722 team-seasons.[9] Table 4 presents correlations between player metrics aggregated up to the team level and team wins one season ahead. For example, for all of the players on the Orlando Magic in 2004-05 productivity is predicted using their 2003-04 values.[10] These predictions are then aggregated up and compared to 2004-05 team wins.

---

[9] We thank Justin Kubatko at www.basketball-reference.com and Doug Steele at www.dougstats.com for these data.
[10] Players with less than 250 minutes in 2003-04 (including rookies in 2004-05) are assigned their 2004-05 value of the given metric. Also, players are assigned their actual minutes in 2004-05.

The first column of Table 4 gives the correlations for team wins with the given player metric from the previous season aggregated up to the team level. For the Minutes per Game metric that correlation is 0.823. The key question in Table 4 (and later tables) is comparing whether the correlation for team wins and aggregated Minutes per Game differs from the correlations with other aggregated player metrics. For example, is the difference in correlation between team wins and aggregated Minutes per Game (0.823) different from the correlation between team wins and aggregated Wins Produced (0.803)?

It would be possible to compute standard errors for each of these correlations, but these correlations are highly correlated and so these standard errors would not help us assess whether differences in correlations are statistically significant. [11] To account for this high correlation we use the following method. We regress team wins for team T ($WINS_T$) onto aggregated (and standardized) player metrics 1 and 2 for team T ($aPM1_T$) and ($aPM2_T$).

$$(10)\ WINS_T = \beta_0 + \beta_1 aPM1_T + \beta_2 aPM2_T + \varepsilon.$$

We then run the hypothesis test that $\beta_1 = \beta_2$, which is identical to running the following regression and testing whether $\lambda_2 = 0$.

$$(11)\ WINS_T = \lambda_0 + \lambda_1 aPM1_T + \lambda_2 \left(aPM2_T - aPM1_T\right) + \varepsilon.$$

This re-formulation in (11) sheds light on precisely what this hypothesis test is testing. Does the difference in the two aggregated player metrics have any explanatory power? If so, the two aggregated player metrics are statistically significantly different.[12] If not, they are not. Table 4 gives these differences in correlations and the p-values for these hypothesis tests.[13] For example, the Minutes per Game row and WP column indicate that the correlation difference for team wins with aggregated Minutes per Game minus team wins with aggregated Wins Produced is 0.021 with a p-value 0.030, suggesting that the difference in correlations is statistically significant at the 3% level.

---

[11] Standard errors can be calculated by regressing team wins onto the aggregated player metric (after both variables have been normalized to have a minutes-weighted mean of zero and standard deviation of zero). The standardization of the dependent and independent variables implies that the formula for the coefficient for the aggregated player metric is identical to the formula for the correlation coefficient between the two variables. In addition, the standard error is the appropriate standard error for that correlation. This approach would allow for computing standard errors assuming that errors are clustered by team (observations for a team are not fully independent), which has the effect of increasing standard errors.

[12] This is not the only way in which to test whether two correlations are statistically significantly different from each other, but this regression-based approach is in the spirit of Wooldridge (2002) and allows this hypothesis test to be tested under the assumption that errors are clustered by team, again a more conservative assumption that should lead to fewer correlation differences being statistically significant.

[13] Correlation differences that are statistically significant at the 1% level are in bold and those statistically significant at the 5% level (but not the 1% level) are in italics.

Table 4
Explaining Team Wins One Season Ahead

		Difference in Correlations (P-Value)					
Player Metric	**Correlation**	**MPG**	**PPG**	**EFF**	**PER**	**WP**	**AWS**
Minutes per Game (MPG) with team adjustment	0.823	--	0.006 (0.299)	0.004 (0.681)	0.019 (0.066)	*0.021 (0.030)*	-0.005 (0.502)
Points per Game (PPG) with team adjustment	0.817	-0.006 (0.299)	--	-0.003 (0.722)	0.012 (0.103)	0.014 (0.144)	-0.012 (0.122)
NBA Efficiency (EFF) with team and position adjustments	0.820	-0.004 (0.681)	0.003 (0.722)	--	**0.015 (0.003)**	**0.017 (0.003)**	*-0.009 (0.050)*
Player Efficiency Rating (PER) with team and position adjustments	0.805	-0.019 (0.066)	-0.012 (0.103)	**-0.015 (0.003)**	--	0.002 (0.824)	**-0.024 (0.000)**
Wins Produced (WP) with team and position adjustments	0.803	*-0.021 (0.030)*	-0.014 (0.144)	**-0.017 (0.003)**	-0.002 (0.824)	--	**-0.026 (0.000)**
Alternate Win Score (AWS) with position and team adjustments	0.829	0.005 (0.502)	0.012 (0.122)	*0.009 (0.050)*	**0.024 (0.000)**	**0.026 (0.000)**	--

Notes: Data are for 722 team-seasons from 1980-81 through 2006-07. *Correlation* gives the correlation of team wins with the player metric in the previous season aggegated up to the team level. The *Difference in Correlations* gives the difference in the correlation of the row aggregated player metric with team wins minus the correlation of the column aggregated player metric. The *P-Value* gives the p-value for the hypothesis test that those two correlations are equal. P-values are calculated assuming that errors are clustered by team. See text for details.

Overall in Table 4 Alternate Win Score and Minutes per Game (aggregated up to the team level) have the highest correlation with team wins in the following season. Two advanced player evaluation metrics – PER and Wins Produced – have the lowest correlations. In fact, the correlations of teams wins (one season ahead) with PER and Wins Produced are statistically significantly different than from correlations of team wins with NBA Efficiency and Alternate Win Score (with p-values of 0.003 or lower). The correlation with Wins Produced also is statistically significantly lower than with Minutes per Game at the 3% level.

Tables 5 and 6 repeat this exercise, but look at teams win two seasons and three seasons ahead.[14] More players have switched teams over a two or three season period, and so team adjustments become less of a factor. On the other hand, player productivity changes over the course of a player's career, so looking two or three seasons ahead results in more noise, lower correlations, and fewer statistically significant differences in correlations.

Results looking two and three seasons ahead are surprisingly similar to those looking one season ahead. Once again, Wins Produced and PER perform quite poorly. Interestingly, Alternate Win Score consistently outperforms Wins Produced. Predicting one season ahead, it performs better with a p-value less than 0.001. Two seasons ahead, the p-value is 0.001 and three seasons ahead, it is 0.036. Essentially, the only difference between the two methods is the way they value shots versus rebounds. The Appendix demonstrates a flaw with the theoretical model for Wins Produced and these results indicate that it appears to quite important empirically, as well.

[14] Player metric values for the season ahead (or two seasons ahead) are used when a player has played less than 250 minutes in the given season.

Table 5
Explaining Team Wins Two Seasons Ahead

Player Metric	Correlation	Difference in Correlations (P-Value)					
		MPG	PPG	EFF	PER	WP	AWS
Minutes per Game (MPG) with team adjustment	0.761	--	0.016 (0.126)	-0.008 (0.503)	0.013 (0.233)	0.016 (0.232)	-0.010 (0.280)
Points per Game (PPG) with team adjustment	0.745	-0.016 (0.126)	--	-0.024 (0.077)	-0.003 (0.740)	0.000 (0.993)	*-0.026 (0.025)*
NBA Efficiency (EFF) with team and position adjustments	0.769	0.008 (0.503)	0.024 (0.077)	--	**0.021 (0.001)**	**0.024 (0.001)**	-0.002 (0.784)
Player Efficiency Rating (PER) with team and position adjustments	0.749	-0.013 (0.233)	0.003 (0.740)	**-0.021 (0.001)**	--	0.003 (0.729)	**-0.023 (0.000)**
Wins Produced (WP) with team and position adjustments	0.745	-0.016 (0.232)	0.000 (0.993)	**-0.024 (0.001)**	-0.003 (0.729)	--	**-0.026 (0.001)**
Alternate Win Score (AWS) with position and team adjustments	0.771	0.010 (0.280)	*0.026 (0.025)*	0.002 (0.784)	**0.023 (0.000)**	**0.026 (0.001)**	--

Notes: Data are for 699 team-seasons from 1981-82 through 2006-07. *Correlation* gives the correlation of team wins with the player metric two seasons previous aggegated up to the team level. The *Difference in Correlations* gives the difference in the correlation of the row aggregated player metric with team wins minus the correlation of the column aggregated player metric. The *P-Value* gives the p-value for the hypothesis test that those two correlations are equal. P-values are calculated assuming that errors are clustered by team. See text for details.

Table 6
Explaining Team Wins Three Seasons Ahead

Player Metric	Correlation	Difference in Correlations (P-Value)					
		MPG	PPG	EFF	PER	WP	AWS
Minutes per Game (MPG) with team adjustment	0.725	--	*0.033 (0.011)*	0.013 (0.386)	*0.039 (0.018)*	0.031 (0.059)	0.010 (0.485)
Points per Game (PPG) with team adjustment	0.692	*-0.033 (0.011)*	--	-0.020 (0.207)	0.006 (0.651)	-0.003 (0.901)	-0.024 (0.124)
NBA Efficiency (EFF) with team and position adjustments	0.712	-0.013 (0.386)	0.020 (0.207)	--	**0.026 (0.003)**	0.018 (0.080)	-0.003 (0.696)
Player Efficiency Rating (PER) with team and position adjustments	0.686	*-0.039 (0.018)*	-0.006 (0.651)	**-0.026 (0.003)**	--	-0.008 (0.562)	**-0.029 (0.000)**
Wins Produced (WP) with team and position adjustments	0.695	-0.031 (0.059)	0.003 (0.901)	-0.018 (0.080)	0.008 (0.562)	--	*-0.021 (0.036)*
Alternate Win Score (AWS) with position and team adjustments	0.716	-0.010 (0.485)	0.024 (0.124)	0.003 (0.696)	**0.029 (0.000)**	*0.021 (0.036)*	--

Notes: Data are for 676 team-seasons from 1982-83 through 2006-07. *Correlation* gives the correlation of team wins with the player metric three seasons previous aggegated up to the team level. The *Difference in Correlations* gives the difference in the correlation of the row aggregated player metric with team wins minus the correlation of the column aggregated player metric. The *P-Value* gives the p-value for the hypothesis test that those two correlations are equal. P-values are calculated assuming that errors are clustered by team. See text for details.

Table 7 summarizes the results from Tables 4 through 6. The highest average correlation across the three tables belongs to Alternate Win Score (0.772), followed by Minutes per Game (0.770) and NBA Efficiency (0.767). Points per Game (0.751), Wins Produced (0.748), and PER (0.747) all have much lower average correlations. Because Points per Game is so different than the advanced player metrics, its correlation differences are statistically significant in only one of 15 comparisons, whereas the correlation difference is statistically significant at the 5% or less level 40 percent of the time with Wins Produced and 47 percent of the time with PER. Clearly, PER and Wins Produced poorly predict future team wins.

Table 7
Distribution of How Well Metrics Predict Team Wins

Player Metric	Average Correlation	Positive Correlation Differences P-Values ≤ 0.01	Positive 0.01-0.05	Positive > 0.05	Negative Correlation Differences P-Values > 0.05	Negative 0.01-0.05	Negative ≤ 0.01
Minutes per Game	0.770	0.0%	20.0%	60.0%	20.0%	0.0%	0.0%
Points per Game	0.751	0.0%	0.0%	23.3%	63.3%	13.3%	0.0%
NBA Efficiency	0.767	33.3%	0.0%	33.3%	26.7%	6.7%	0.0%
PER	0.747	0.0%	0.0%	20.0%	33.3%	6.7%	40.0%
Wins Produced	0.748	0.0%	0.0%	16.7%	43.3%	13.3%	26.7%
Alternate Win Score	0.772	33.3%	20.0%	40.0%	6.7%	0.0%	0.0%
Overall	0.759	11.1%	6.7%	32.2%	32.2%	6.7%	11.1%
Expected	--	0.5%	2.0%	47.5%	47.5%	2.0%	0.5%

Notes: The p-value distribution comes from Tables 4, 5, and 6. *Positive (Negative) Correlation Differences* indicate that the correlation with team wins is higher (lower) for that player metric than it is for the comparison.

6. Evaluating the Player Evaluation Metrics using Adjusted Plus/Minus

Without players who switch teams or substantially change their minutes played, evaluating how well player evaluation metrics predict future wins would have no power. Looking two or three seasons ahead is a step in the right direction, but even when there is enough player variation to assess the metrics, using team wins to evaluate a player metric gives credit (or blame) to the player when that player is in the game *and when he is not*. It would be more efficient to evaluate how player metrics predict how the team performs when particular players are in the game and to somehow account for the players the player is playing with and against.

And this precisely what *adjusted plus/minus statistics* do; they measure how the team does when a given player is in the game, holding the effects of the other players (on both teams) constant. Adjusted plus/minus statistics measure all of the contributions a player makes that affect the point differential while they are in the game. Even contributions that do not show up in the box score, such as a good pick, poor help defense, or keeping the floor spread are accounted for. To estimate *adjusted plus/minus statistics*, games are broken down into all of the combinations of players from both teams (over 182,000 combinations between 2002-03 and 2006-07). Then with these data we run the following regression separately by season.[15]

$$(12)\quad DIFF = \beta_0 + \beta_1 X_1 + \beta_2 X_2 + \ldots + \beta_K X_K,\ \textit{where}$$

DIFF = home team points per possession – away team points per possession, $X_i = 1$ if player i is playing at home, = -1 if player i is playing away, = 0 otherwise (i.e. if player i is not playing),

[15] These regressions are weighted by the number of possessions. The clutch/garbage time weighting in Rosenbaum (2004) is not used.

and ε = i.i.d. error term.

The intercept β_0 measures the average home court advantage across all teams, and the slope coefficients (β_1, β_2, . . . , β_K) measure how the *team* point differential changes when a particular player is on the court (relative to the reference players), holding the effects of other players (on both teams) constant. The reference players are all players playing less than 250 minutes in a given season.

What does this "holding other players constant" mean? Strictly speaking, it means comparing how a team does with a particular player in the game (and four teammates and five opponents) relative to how the team does when that player is replaced by a replacement player keeping all of the other players the same. This is what is meant by "holding the other players constant," since this exercise can be repeated with any other combination of other players. Another way to think of these βs is that they measure how the team does when a given player is on the floor adjusting for the other players on the floor. This takes out the effect of a player who is fortunate to always play with Kevin Garnett or unfortunate enough to always being matched with rookies or marginal NBA players.

Table 8 presents correlations and differences in correlations similar to those in Tables 4 through 6, except that now the outcome is a player's adjusted plus/minus (in the current season) and the player metric is the individual-level player metric rather than the aggregated player metric. The sample is limited to players who played 500 minutes or more and is weighted by minutes played. Standard errors are computed assuming that errors are clustered by player.

Correlations are much smaller in Table 8 than in the team wins tables, but the differences in correlations in Table 8 are much greater than in the team wins tables. PER, for example, has a statistically significantly higher correlation with adjusted plus/minus than any of the other player evaluation metrics (with p-values that are 0.005 or lower). Eleven of the fifteen comparisons between player metrics are statistically significant at the 1% level. Six of the fifteen are statistically significant at less than the 0.1% level. The lower correlations are due to the fact that adjusted plus/minus may be "noisy" at the individual-level (discussed in more detail below), but as a barometer it is extremely powerful.[16]

Using adjusted plus/minus PER clearly outperforms all of the other metrics, followed by NBA Efficiency and Alternate Win Score. Points per Game clearly outperforms both Minutes per Game and Wins Produced. Wins Produced is statistically significantly worse than all of the other player metrics other than Minutes per Game at less than the 0.1% significance level.[17]

[16] A "noisy" dependent variable is not a problem in cases like this, where (a) the "noise" is uncorrelated with the independent variables and (b) the "noise" isn't so great that the regression has no power.
[17] Note that an earlier version of Wins Produced used in Berri and Krautmann (2006) has a much lower correlation of 0.0918 with adjusted plus/minus.

Table 8
Explaining Adjusted Plus/Minus

		Difference in Correlations (P-Value)					
Player Metric	**Correlation**	**MPG**	**PPG**	**EFF**	**PER**	**WP**	**AWS**
Minutes per Game (MPG) with team adjustment	0.412	--	**-0.093 (0.000)**	**-0.112 (0.000)**	**-0.130 (0.000)**	0.003 (0.934)	**-0.105 (0.000)**
Points per Game (PPG) with team adjustment	0.505	**0.093 (0.000)**	--	-0.019 (0.186)	**-0.038 (0.004)**	**0.096 (0.000)**	-0.013 (0.268)
NBA Efficiency (EFF) with team and position adjustments	0.524	**0.112 (0.000)**	0.019 (0.186)	--	**-0.019 (0.001)**	**0.114 (0.000)**	0.006 (0.531)
Player Efficiency Rating (PER) with team and position adjustments	0.543	**0.130 (0.000)**	**0.038 (0.004)**	**0.019 (0.001)**	--	**0.133 (0.000)**	**0.025 (0.005)**
Wins Produced (WP) with team and position adjustments	0.409	-0.003 (0.934)	**-0.096 (0.000)**	**-0.114 (0.000)**	**-0.133 (0.000)**	--	**-0.108 (0.000)**
Alternate Win Score (AWS) with position and team adjustments	0.518	**0.105 (0.000)**	0.013 (0.268)	-0.006 (0.531)	**-0.025 (0.005)**	**0.108 (0.000)**	--

Notes: Data are for 1,627 player-seasons from 2002-03 through 2006-07 with at least 500 minutes and are weighted by minutes played. *Correlation* gives the correlation of adjusted plus/minus with the player metric. The *Difference in Correlations* gives the difference in the correlation of the row aggregated player metric with adjusted plus/minus minus the correlation of the column player metric. The *P-Value* gives the p-value for the hypothesis test that those two correlations are equal. P-values are calculated assuming that errors are clustered by player. See text for details.

Table 9 repeats the exercise in Table 8, but looking one season ahead. The results are nearly identical to those in Table 8 with PER clearly outperforming the other metrics; NBA Efficiency, Alternate Win Score, and Points per Game are next in line. Once again, Minutes per Game and Wins Produced perform poorly. Both do a very poor job explaining or predicting how teams play when particular players are in the game.

Table 9
Explaining Adjusted Plus/Minus One Season Ahead

		Difference in Correlations (P-Value)					
Player Metric	**Correlation**	**MPG**	**PPG**	**EFF**	**PER**	**WP**	**AWS**
Minutes per Game (MPG) with team adjustment	0.348	--	**-0.106 (0.000)**	**-0.130 (0.000)**	**-0.146 (0.000)**	-0.006 (0.644)	**-0.109 (0.000)**
Points per Game (PPG) with team adjustment	0.454	**0.106 (0.000)**	--	-0.025 (0.142)	*-0.040 (0.011)*	**0.100 (0.004)**	-0.003 (0.668)
NBA Efficiency (EFF) with team and position adjustments	0.478	**0.130 (0.000)**	0.025 (0.142)	--	*-0.016 (0.023)*	**0.124 (0.000)**	*0.021 (0.028)*
Player Efficiency Rating (PER) with team and position adjustments	0.493	**0.146 (0.000)**	*0.040 (0.011)*	*0.016 (0.023)*	--	**0.140 (0.000)**	**0.037 (0.000)**
Wins Produced (WP) with team and position adjustments	0.354	0.006 (0.644)	**-0.100 (0.004)**	**-0.124 (0.000)**	**-0.140 (0.000)**	--	**-0.103 (0.000)**
Alternate Win Score (AWS) with position and team adjustments	0.457	**0.109 (0.000)**	0.003 (0.668)	*-0.021 (0.028)*	**-0.037 (0.000)**	**0.103 (0.000)**	--

Notes: Data are for 1,306 player-seasons from 2002-03 through 2006-07 with 500 minutes both in the current and past seasons and are weighted by minutes played. *Correlation* gives the correlation of adjusted plus/minus with the player metric in the previous season. The *Difference in Correlations* gives the difference in the correlation of the row aggregated player metric with adjusted plus/minus minus the correlation of the column player metric. The *P-Value* gives the p-value for the hypothesis test that those two correlations are equal. P-values are calculated assuming that errors are clustered by player. See text for details.

Table 10 summarizes the results from Table 8 and 9. Using adjusted plus/minus as a barometer has been criticized because it is noisy and varies a lot from season to season. But if noise was a problem with using it as a barometer, we would find very few differences in correlations that were statistically significant. The results in Table 10 show exactly the opposite; a remarkable two thirds of comparisons are statistically significant at the 1% level or less. This is powerful evidence of its usefulness as a barometer, especially given that box scores are not used to construct adjusted plus/minus and so there is no reason to expect it would be biased towards any particular box score based metric. This evidence strongly suggests that PER does a good job explaining how teams perform when particular players are in the game. Wins Produced does a particularly poor job (as does Minutes per Game), and it is worth noting that in both Tables 8 and 9, it does statistically significantly worse than Points per Game at the 0.4% level or lower.

Table 10
Distribution of How Well Metrics Explain Adjusted Plus/Minus

Player Metric	Average Correlation	Positive Correlation Differences P-Values ≤ 0.01	Positive Correlation Differences P-Values 0.01-0.05	Positive Correlation Differences P-Values > 0.05	Negative Correlation Differences P-Values > 0.05	Negative Correlation Differences P-Values 0.01-0.05	Negative Correlation Differences P-Values ≤ 0.01
Minutes per Game	0.380	0.0%	0.0%	10.0%	10.0%	0.0%	80.0%
Points per Game	0.480	40.0%	0.0%	0.0%	40.0%	10.0%	10.0%
NBA Efficiency	0.501	40.0%	10.0%	30.0%	0.0%	10.0%	10.0%
PER	0.518	80.0%	20.0%	0.0%	0.0%	0.0%	0.0%
Wins Produced	0.382	0.0%	0.0%	10.0%	10.0%	0.0%	80.0%
Alternate Win Score	0.488	40.0%	0.0%	20.0%	10.0%	10.0%	20.0%
Overall	0.458	33.3%	5.0%	11.7%	11.7%	5.0%	33.3%
Expected	--	0.5%	2.0%	47.5%	47.5%	2.0%	0.5%

Notes: The p-value distribution comes from Tables 8 and 9. *Positive (Negative) Correlation Differences* indicate that the correlation with adjusted plus/minus is higher (lower) for that player metric than it is for the comparison. metrics.

In comparing the two sets of results for team wins and adjusted plus/minus, the relative correlations for PER and Minutes per Game differ substantially. PER is the best metric explaining adjusted plus/minus, i.e. how teams play when particular players are in the game, but when aggregated predicts team wins poorly. Minutes per Game predicts team wins well, but does poorly explaining adjusted plus/minus. What explains these differences for the two different outcomes?

The most likely culprits here are (a) assigning future season values to rookies and low minutes player and (b) the greater importance of the team adjustment in explaining team wins. On the first point, limiting the sample to those teams with fewer than 10 percent of minutes being played by rookies/low minutes eliminates more than 90 percent of the correlation difference between Minutes per Game and PER in predicting team wins. On the second point, note that in Table 3 team efficiency has by far the highest correlation with Minutes per Game; alternatively, the lowest correlation is with PER. The team adjustment captures elements of player productivity that are difficult to attribute to specific players, such as team defense. Those elements may be highly correlated with coaching and thus carry over from season to season, even when players change. This may bias using team wins as a barometer towards metrics that place greater weight on team efficiency, such as Minutes per Game. Further evidence of this possible bias is that the metric with the second highest correlation with team efficiency – Alternate Win Score – performs the best predicting team wins, but does not perform as well explaining adjusted plus/minus. This possible bias of the team wins barometer toward metrics highly correlated with team efficiency is unlikely to carry over to the individual player level adjusted plus/minus, where team adjustments don't matter nearly as much. Overall we believe that both sets of results support our hypothesis that advanced metrics can sometimes be worse than approximations of NBA decision-making, but of the two we find the adjusted plus/minus results more convincing.

7. Adjusted Plus/Minus as a Player Evaluation Metric

Given the usefulness of using adjusted plus/minus as a barometer, it is natural to wonder how well it does as a predictor. Table 11 examines how well adjusted plus/minus (measured one year at a time) explains team wins one season, two seasons, and three seasons ahead and how well it explains adjusted plus/minus one season ahead. In short, one-year adjusted plus/minus is terrible as a predictor; it performs far worse than all six of the metrics on all four outcomes – even predicting adjusted plus/minus.

Table 11
Adjusted Plus/Minus as a Predictor

Outcome	Correlation	Standard Error	N	Range of Correlations	
				Minimum	Maximum
Explaining team wins one season ahead	0.534	0.080	119	0.803	0.829
Explaining team wins two seasons ahead	0.401	0.101	90	0.745	0.779
Explaining team wins three seasons ahead	0.263	0.144	60	0.686	0.725
Explaining adjusted plus/minus one season ahead	0.264	0.033	1,055	0.348	0.493

Notes: Data are for team-seasons from 2003-04 through 2006-07 for explaining team wins one season ahead, from 2004-05 through 2006-07 for explaining team wins two seasons ahead, and from 2005-06 through 2006-07 for explaining team wins three seasons ahead. Data are for player-seasons from 2003-04 through 2006-07 with 500 minutes in both current and past seasons for explaining adjusted plus/minus one season ahead. Standard errors are clustered by team when explaining team wins and by player when explaining adjusted plus/minus. *Range of Correlations* gives the range of correlations for the other player metrics from the earlier tables. See text for details.

Despite this poor predictive ability, one-year adjusted plus/minus was one of the first advanced player evaluation metrics to be used in the NBA. The *Indianapolis Star* refers to the one-year adjusted plus/minus metric developed by Wayne Winston and Jeff Sagarin as the "Michael Jordan of statistics" (Oliver, 2004). Leonhardt (2003) reports that Dallas Mavericks owner Mark Cuban "said the ratings had influenced every player signing he had made over the past few years."

The poor predictive ability of one-year adjusted plus/minus is due to its "noisiness," which is discussed in Rosenbaum (2004, 2005a, and 2005b), Lewin (2006), and Iliardi (2007). Teams do not randomize over player combinations playing in the game at any given time, e.g. starters often play together, so econometrically if takes a lot of games to get precise estimates of individual player effects. Using several seasons of data and/or combining adjusted plus/minus with box score statistics can reduce the "noisiness" of adjusted plus/minus, but evaluating those variations of adjusted plus/minus is beyond the scope of this paper.

Because adjusted plus/minus statistics can also be used to rate players, it is sometime argued that this somehow invalidates using adjusted plus/minus statistics as an evaluation method. An analogous argument would be that if someone used team wins to evaluate players, somehow that would invalidate team wins as method of evaluating player metrics. Adjusted plus/minus statistics simply measure how the team does when a given player is in the game and that is precisely what player metrics are supposed to measure. And they do so without using any box score statistics, so they are not systematically biased towards any box score based metric.

8. Conclusion

Does statistical analysis lead to better decisions than intuition and human judgment? Leonhardt (2005) presents examples of decisions involving medical diagnoses, college admissions, and baseball player evaluation where the answer appears to be yes. Hence, it is understandable that such evidence might lead analysts, such as Berri et al. (2006) to be confident that statistical analysis would outperform intuition and human judgment in NBA decision-making, as well.

> "One can play basketball. One can watch basketball. One can both play and watch basketball for a thousand years. If you do not systematically track what the players do, and then uncover the statistical relationship between these actions and wins, you will never know why teams win and why teams lose. Staring at these players play is not a method that will ever yield the answers that the proper analysis of statistics will yield. And this is true if stare for one day, or as we said, if you stare for a thousand years." (Berri et al. 2006, p. 215)

Academics rarely use such harsh language, and the results in this paper suggest why. Statistical analysis is not always better than intuition and human judgment. We find that the metric of choice for academics – the Wins Produced from Berri et al. (2006) – does a very poor job predicting team wins and explaining (or predicting) adjusted plus/minus, i.e. how a team plays when a given player is in the game.

The evidence is overwhelming. In our main analysis we compare six different metrics. Two of the metrics – Minutes per Game and Points per Game – are proxies of NBA decision-making. Salary regressions indicate that NBA decision-makers consider more than Minutes per Game and Points per Game in their valuations of players. In particular, NBA decision-makers appear to take into account the success of the team when setting salaries, so team adjustments are added to these metrics. Even with the team adjustment, we expected that these proxies would perform far worse than any of the advanced metrics, because they clearly leave out other determinants of salaries. But that did not turn out to be the case. Wins Produced performed worse than both of these metrics, especially Points per Game.

We also compare Wins Produced to two other advanced metrics that some have argued are good approximations of NBA decision-making – NBA Efficiency and John Hollinger's PER.[18] Wins Produced is statistically significantly worse (at the 0.3% level or less) in more than half of the comparisons with these two metrics, i.e. its correlation with the given outcome (either team wins or player-specific adjusted plus/minus) is lower. The reason Wins Produced performs so poorly is that its theoretical model is flawed; its assumption of different production functions for own team and opponent possessions is at odds with the symmetry of possession production in basketball. Correcting this theoretical flaw leads to the Alternate Win Score metric, which performs better than Wins Produced in all five comparisons. The difference is statistically significant at the 4% level in one comparison and at the 0.1% level or less in the other four comparisons. The results leave little doubt that this simple adjustment correcting the theoretical flaw in Wins Produced better fits the data.

We also find that another advanced player evaluation model – adjusted plus/minus measured one season at a time – performs even worse than Wins Produced. One season does not provide enough variation in the combinations of players on the floor in order to precisely estimate the effects of individual players. This resulting "noisiness" is not problematic when using adjusted

[18] See Berri et al. (2006) and Berri (2006b).

plus/minus as a barometer, but as a predictor it leads to adjusted plus/minus performing much worse than any of the other metrics.

The results of this paper should not discourage the use of statistical analysis in basketball. John Hollinger's PER metric, for example, performs quite well. The important thing to take away is that statistical analysis must be done carefully and rigorously, with an appreciation of the complex and dynamic interactions that are at the heart of the game of basketball. The non-academic statistical analyst community in basketball, led by Dean Oliver, author of *Basketball on Paper,* challenges NBA decision-making, but with a degree of restraint and respect for conventional wisdom.

> "My null hypothesis is usually traditional coaching or management wisdom. So a hot hand exists, defense wins championships, and statistics are irrelevant until I prove otherwise (which I think I've done in many cases). Others may choose a different null hypothesis, but I think mine makes sense because I work with coaches and management and I'm not Billy Beane – it is my burden to prove things, not theirs." (Oliver, 2005).

This subtle shift of the burden of proof generates a healthy skepticism of findings from statistical models that challenge conventional wisdom. Basketball is a game where the five players on the court for a team jointly produce wins. Apportioning credit from this joint production function is difficult for both subjective and statistical analysis, so it is reasonable to presume that conventional wisdom can be a useful guide for how make these simplifications appropriately. Coaches, scouts, and general managers have been dealing with this joint production function for decades under intense scrutiny and competition, so it seems reasonable to assume that maybe they have learned something through that process.

That doesn't mean conventional wisdom is always right. But the healthy skepticism suggested by Oliver is helpful in uncovering inappropriate simplifications in our statistical models. That said, Wins Produced and adjusted plus/minus have played a significant role in advancing the field of basketball analytics. Despite their warts these models have generated deeper thinking about how wins are produced, what statistical models are measuring, and how statistical models can be evaluated. All of that is part of the process of helping statistical analysis become more rational than "irrational" NBA decision-makers.

Appendix 1: The Derivation of Wins Produced and Team Adjustments

Berri et al. (2006), Berri et al. (2007), and Berri (2008) create a model for player productivity that is grounded in the relationship between team wins and own team and opponent points per possession. Using data from the 1993-94 through 2004-05 seasons they estimate that an extra own team point scored produces 0.033 wins, holding opponent points, own possessions, and opponent possessions constant. Similarly, an extra opponent point reduces wins by 0.033, ceteris paribus.

Now using the above logic, it would be misleading to say that an extra own team possession, holding own team points, opponent points, and opponent possessions constant, generates wins, because holding own team and opponent points constant, an extra possession has no value.

However, there is no box score statistic that allocates opponent points to individual players, and there are two ways players generate more opponent points: by allowing opponents to score more points per possession and by allowing opponents more possessions. For the latter reason using possessions through shots or turnovers comes at the cost of allowing the other team another possession. Berri et al. (2006) estimates that teams score 1.02 points per possession, which implies that a possession is worth 0.034 wins (1.02 points/possession times 0.033 wins/point).

Hence, at the team level Wins Produced for team *T* (WP_T) is the following.

$$(A.1)\ WP_T = 0.033 \times \left(PTS_T - PTS_O\right) - 0.034 \times (POSS_T - POSS_O),$$

where PTS_T is points for team *T*, PTS_O is points for opponents *O*, $POSS_T$ is possessions for team *T*, and $POSS_O$ is possessions for opponents *O*. Since possessions for own team and opponents are approximately equal, Wins Produced at the team level is approximately equal to point differential multiplied by 0.033. To apply this formula for individual players, points and possessions need to be allocated to particular players. In other words, assumptions must be made about how individual players produce points and possessions.

Berri et al. (2006) assumes that own team points are produced by the player who scores the points and that opponent points are generated equally in proportion to the minutes played by a given player.[19] In order to discuss the possession production functions assumed by Berri et al. (2006), it is useful to introduce the general formula for possessions for own team *T* ($POSS_T$) from Kubatko et al. (2007).

$$\begin{aligned}(A.2)\ POSS_T = &\left(FGM_T + \lambda_M FTM_T\right) + \alpha\left[\left(FGA_T - FGM_T\right) + \lambda_{MS}\left(FTA_T - FTM_T\right) - ORB_T\right] \\ &+ \left(1-\alpha\right) DRB_O + TO_T,\end{aligned}$$

where FGA_T is field goal attempts for team *T*,
FGM_T is field goals made for team *T*,
FTA_T is free throw attempts for team *T*,
FTM_T is free throws made for team *T*,
ORB_T is offensive rebounds for team *T*,
DRB_O is defensive rebounds for opponent *O*,
TO_T is turnovers for team *T*,
λ_M is the fraction of made free throws that end possessions,[20]
λ_{MS} is the fraction of missed free throws that end possessions, and
α is a parameter between zero and one.

[19] This is only true for field goals. For free throws Berri et al. (2006) assume that opponent free throw points are proportional to personal fouls.

[20] Free throws that end possessions do not include first free throws of two, first and second free throws of three, or free throws due to technicals, flagrant fouls, and clear path fouls. Also, free throws after made field goals are not (double) counted, since the made field goal already has counted that possession. In the 2002-03 through 2005-06 seasons, 43.8% of free throws were possession ending free throws, although this breaks down to 45.2% for made free throws and 39.5% for missed free throws. The large difference comes from players being more likely to miss first free throws in a succession of free throws and more low-percentage free throw shooters shooting free throws after a made field goal.

A similar formula for possessions for opponent *O* ($POSS_O$) is the following.

$$(A.3)\ POSS_O = (FGM_O + \lambda_M FTM_O) + \alpha\left[(FGA_O - FGM_O) + \lambda_{MS}(FTA_O - FTM_O) - ORB_O\right] + (1-\alpha)DRB_T + TO_O,$$

For the own team production function Berri et al. (2006) assume that $\alpha = 1$ and that $\lambda_M = \lambda_{MS} = 0.47$, which results in the following formula for possessions for team *T* ($POSS_T$).

$$\begin{aligned}(A.4)\ POSS_T &= (FGM_T + 0.47\times FTM_T) + 1\times\left[(FGA_T - FGM_T) + 0.47\times(FTA_T - FTM_T) - ORB_T\right] \\ &+ 0\times DRB_O + TO_T \\ &= FGA_T + 0.47\times FTA_T - ORB_T + TO_T.\end{aligned}$$

For the opponents possession production function Berri et al. (2006) assume that $\alpha = 0$ and that $\lambda_M = \lambda_{MS} = 0.47$, which results in the following formula for possessions for team *O* ($POSS_O$).

$$\begin{aligned}(A.5)\ POSS_O &= (FGM_O + 0.47\times FTM_O) + 0\times\left[(FGA_O - FGM_O) + 0.47\times(FTA_O - FTM_O) - ORB_O\right] \\ &+ 1\times DRB_T + TO_O \\ &= FGM_O + 0.47\times FTM_O + DRB_T + TO_O.\end{aligned}$$

Berri et al. (2006) provide no justification for their assumptions about α or how it makes sense to assume different production functions for own team and opponent possessions. By assuming different production functions for own team and opponents possessions, Berri et al. (2006) assume that defensive rebounds are worth a full possession in the opponent production function, but nothing in the own team production function. Missed shots and offensive rebounds count a full possession in the own team production function, but nothing in the opponents. Hence, the theoretical model that serves as the foundation for Wins Produced assumes that basketball is something other than a symmetrical game where both teams produce possessions in the same manner. Moreover, teams are interchangeably "own team" and "opponent," so assuming different production functions implies that possession production for a given team in a given game is different depending on whether that team is the "own team" or "opponent."

Berri et al. (2006) make this mistake, because they adopt the misleading terminology of "possessions employed" and "possessions acquired" to refer to own team possessions and opponent possessions. This terminology is not problematic in the context of estimating possessions, as discussed in Kubatko et al. (2007). But in a theory of how possessions are produced, this terminology obscures the fact that on both sides of the court both teams jointly produce possessions in a symmetric manner. When the own team has possession of the ball, possessions are jointly produced through own team shots, own team offensive rebounds, own team turnovers, and opponent defensive rebounds. When the opponents have possession of the ball, the reverse occurs; possessions are jointly produced through opponent shots, opponent offensive rebounds, opponent turnovers, and own team defensive rebounds. It is hard to imagine a reasonable theory that can explain why these two production functions would be different, which implies that the assumption about α must be the same in the own team and opponents possession equations.

The assumption of two different production functions results in Berri et al. (2006) dramatically over-valuing rebounding versus scoring. (See Table 3 which shows the very high correlation of Wins Produced with rebounding and very low correlation with shots.) These assumptions lead to scorers who generate less than 1.02 points per shot reducing their Wins Produced with every shot they take.[21] Such an assumption may be defensible when analyzing teams, but at the individual-level this assumption implies that creating shots has no value and that shooting percentages are the same for the average shot as for the marginal shot. Theory implies that the efficiency of the marginal shot should be the same across all players on a team, but if the marginal shot has the same efficiency as the average shot, i.e. shot creation has no value, then teams should have their most efficient shooters to shoot *every* shot. Hence, the assumptions underlying Wins Produced imply that coaches are incredibly inefficient. Practically any team in the NBA could be the best team in NBA history by employing a simple change in strategy – have their most efficient shooter take every shot. The fact that no team has successfully employed such a simple strategy is a hint of the possible flaws underlying Wins Produced.

Ignoring the problematic assumptions by Berri et al. (2006), plugging (A.4) and (A.5) into (A.1) generates the following formula for Wins Produced for team *T*.

$$
\begin{aligned}
(A.6)\ WP_T &= 0.033\times\left(PTS_T - PTS_O\right) \\
&\quad - 0.034\times\left[\left(FGA_T + 0.47\times FTA_T - ORB_T + TO_T\right) - \left(FGM_O + 0.47\times FTM_O + DRB_T + TO_O\right)\right]
\end{aligned}
$$

Berri et al. (2006) apply this model to individual players by assuming that the production function for players is identical to that for teams (although they relax this assumption later). All of the own team statistics are replaced by the statistics for the individual player. Opponent free throws made (and the points generated by those free throws) are allocated to individual players according their proportion of the overall team fouls. Opponent turnovers that are steals are allocated according to the individual steal statistics. Block shots are estimated to reduce opponent two-point field goals made by 0.65 and so part of the opponent field goals made are allocated to players using individual block shot statistics. The remainder of opponent field goals made, the points generated by those field goals, and non-steal turnovers are allocated to individual players according to their proportion of the overall minutes played.[22]

It is useful to note that at this point that all of the terms in Wins Produced, except for (PTS_T – PTS_O), are approximations of ($POSS_T$ – $POSS_O$). Since possessions for both teams approximately add up to zero, all of these additional terms approximately add up to zero at the team level.

---

[21] The Alternate Win Score metric is derived from such an assumption that $\alpha = 0.7$, which is arbitrary but consistent with the fact that about 70 percent of missed shots are rebounded by the other team. It implies that shots that generate more than 0.72 points per shot are positive contributions to a team.

[22] This discussion ignores team rebounds and team turnovers, which are also allocated to individual players according to their proportion of the overall minutes played.

Berri et al. (2006) almost stopped at this point without incorporating assists, and it should be noted that even without assists, Wins Produced aggregated up to the team level explains approximately 95 percent of team wins. Incorporating assists has absolutely no effect on how well aggregated Wins Produced explains team wins. Any model like Wins Produced that aggregates up by team to point differential will explain 95 percent of team wins, so explaining team wins is completely useless for evaluating such models.

In fact, it is useful to consider an alternative theory that assumed that own team statistics should not be allocated to individual players. Such a theory would allocate all statistics to players according to their proportion of the overall minutes played. Such a metric would assume that all players on a given team were equally productive per minute played, yet would explain team wins just as well as Wins Produced.

Berri et al. (2006) estimates that each additional own team assist produces about 0.22 wins, but simply adding that to (A.6) would reduce how much of team wins could be explained by Wins Produced. So Berri et al. (2006) subtracts off a term from each player that allocates 0.22 times own team assists to each player according to their minutes played.

This same method can be used to produce the team adjustments for the non-Wins Produced metrics. For example, the team adjustment for Points per Game can be constructed by assuming that points per game is proportional to productivity per minute within a team.[23] Then, like with the treatment of assists in Wins Produced, the minutes-weighted sum of points per game can be subtracted off from individual players according to their proportion of the overall minutes played. Finally, the metric can then be put into a Wins Produced framework by allocating team efficiency to each player according to their minutes played. Adding a team adjustment to Points per Game in this way embodies a defensible assumption that team points and possessions cannot be allocated to individual players using box score statistics.

Given this adjustment for assists and by assuming that own team possessions are approximately equal to opponent possessions, any player evaluation metric can be given a team adjustment in the theoretical spirit of Wins Produced simply by allocating the residual from the following regression to each player according to the minutes played.

$$(A.7)\ TeamEFF_T = \beta_0 + \beta_1 PM_T + \varepsilon.$$

$TeamEFF_T$ is own team minus opponents points per possession (team efficiency) and PM_T is a given player metric aggregated up to the team-level. The residual from this regression is the part of team efficiency not explained the player evaluation metric. The team adjustment for Wins Produced that is described above is identical to this, except that (1) it is not pace-adjusted, i.e. it does not account for teams that play at a faster pace and (2) it allocates the difference in approximated own team possessions minus approximated opponent possessions to players. However, Berri et al. (2006) asserts that the team adjustment accounts for pace, which eliminates differences due to (1). Berri et al. (2006) also adds in an adjustment for rebounds not attributed to individual players that is simply the difference between approximated own team possessions

[23] Every 40 minutes a 20 points per game scorer plays generates 0.17 wins.

and approximated opponent possessions. Adding in that adjustment equates approximations for own team and opponent possessions and eliminates the differences due to (2).

References

Berri, David J. (2008). "A Simple Measure of Worker Productivity in the National Basketball Association." In *The Business of Sport*, Brad Humphreys and Dennis Howard, editors, 3 volumes. Westport, CT: Praeger, forthcoming in 2008.

----- (2006a). "Answering a Critic." *Wages of Wins Blog*, November 26, 2006, http://dberri.wordpress.com/2006/11/26/answering-a-critic/.

----- (2006b). "A Comment on the Player Efficiency Rating." Wages of Wins Blog, November 17, 2006, http://dberri.wordpress.com/2006/11/17/a-comment-on-the-player-efficiency-rating/.

Berri, David J., and Anthony Krautmann (2006). "Shirking on the Court: Testing for the Dis-Incentive Effects of Guaranteed Pay." *Economic Inquiry*, 44(3), July 2006, pp. 536-546.

Berri, David J., Martin B. Schmidt, and Stacey L. Brook (2007). "The Wins Produced Calculation." *Wages of Wins Blog*, 2007, http://www.wagesofwins.com/CalculatingWinsProduced.html.

----- (2006). *The Wages of Wins: Taking Measure of the Many Myths in Modern Sport*. Stanford, CA: Stanford University Press, 2006.

Coon, Larry (2007). "NBA Salary Cap FAQ." http://members.cox.net/lmcoon/salarycap.htm.

Gladwell, Malcolm (2006). "Game Theory: When it Comes to Athletic Prowess, Don't Believe Your Eyes." *The New Yorker*, May 29, 2006, pp. 86-87, http://www.newyorker.com/archive/2006/05/29/060529crbo_books1.

Hollinger, John (2005). *Pro Basketball Forecast: 2005-06 Edition*. Washington, DC: Brassey's, 2005.

----- (2004). *Pro Basketball Forecast: 2004-05 Edition*. Washington, DC: Brassey's, 2004.

----- (2003). *Pro Basketball Prospectus: All-New 2003-04 Edition*. Washington, DC: Brassey's, 2003.

Iliardi (2007). "Adjusted Plus-Minus: An Idea Whose Time Has Come." *82games.com*, October 28, 2007, http://www.82games.com/ilardi1.htm.

Kubatko, Justin (2007). *Basketball-Reference.Com: Basketball Statistics, Analysis and History*. http://www.basketball-reference.com/.

Kubatko, Justin, Dean Oliver, Kevin Pelton, and Dan T. Rosenbaum (2007). "A Starting Point for Analyzing Basketball Statistics," *Journal of Quantitative Analysis in Sports*, 3(3), Article 1, available at: http://www.bepress.com/jqas/vol3/iss3/1.

Leonhardt, David (2005). "To Play is the Thing." *New York Times*, August 28, 2005, http://www.nytimes.com/2005/08/28/weekinreview/28leon.html.

----- (2003). "Pro Basketball: Mavericks' New Math May Be an Added Edge." *New York Times*, April 27, 2003, http://query.nytimes.com/gst/fullpage.html?res=9904E1DD123DF934A15757C0A9659C8B63&sec=&spon=&pagewanted=all.

Lewin, David (2006). "2005-2006 Adjusted Plus-Minus Ratings." *82games.com*, 2006, http://82games.com/lewin2.htm.

Lewis, Michael (2003). *Moneyball: The Art of Winning and Unfair Game*. New York: W.W. Norton & Company Inc., 2003.

NBA.com (2007). “Efficiency Statistics.” http://www.nba.com/statistics/efficiency.html.

Oliver, Dean (2004). *Basketball on Paper*. Washington, DC: Brassey’s, 2004.

----- (2005). http://sonicscentral.com/apbrmetrics/viewtopic.php?t=299, 2005.

Price, Joseph and Justin Wolfers (2007). “Racial Discrimination Among NBA Referees.” NBER Working Paper #13206, June 2007, http://www.nber.org/papers/w13206.

Rosenbaum, Dan T. (2005a). “Using Statistics in Basketball: The Bar is Higher.” *Dan Rosenbaum’s NBA Blog*, August 29, 2005, http://danrosenbaum.blogspot.com/2005/08/using-statistics-in-basketball-bar-is.html.

----- (2005b). “Keeping Score: A Statistical Holy Grail: The Search for the Winner Within.” *New York Times*, April 10, 2005, http://query.nytimes.com/gst/fullpage.html?res=9B05E4DE153EF933A25757C0A9639C8B63.

----- (2004). “Measuring How NBA Players Help their Teams Win.” *82games.com*, 2004, http://www.82games.com/comm30.htm.

Siler, Ross (2007). “Larry Miller’s Batting Average.” *The Salt Lake Tribune*, October 16, 2007, http://blogs.sltrib.com/jazz/2007/10/larry-miller-batting-average_16.htm.

Stanley, Marcus (2005). Private e-mail correspondence, 2005.

Steele, Doug (2007). “Doug’s NBA and MLB Stats Home Page.” http://www.dougstats.com/.

Wooldridge, Jeffrey M. (2002). *Econometric Analysis of Cross Section and Panel Data*. Cambridge, MA: MIT Press, 2002.