Evaluating Individual Player Contributions in Basketball

Jarad Niemi*

Abstract
This paper introduces novel methodology for evaluating player effectiveness in basketball. We model separately each player's offensive and defensive contributions to their team and provide estimates that adjust for all players on the court. These offensive and defensive contributions are naturally combined to provide a plus-minus statistic for each player. We build a hierarchical model that estimates the underlying distributions for player contributions and shrinks player estimates toward average with more shrinkage given to players who play less. We illustrate the methodology using data from the 2009-2010 NBA regular season. Based on this season's data, we find the best player in the league to be Pau Gasol who provides his team with a 10.9 ± 4.3 point differential compared to an average player. The best offensive player is Steve Nash who provides his team with an additional 11.3 ± 3.0 points per game and the best defensive player is Raja Bell who reduces his opponents scoring by 12.7 ± 7.3 points per game.

Key Words: Bayesian analysis, Gibbs sampling, hierarchical model, plus-minus statistic, offensive rating, defensive rating

1. Introduction

In team sports around the world, statistics are maintained to measure individual player performance (Kubatko et al., 2007). These statistics are meant to give coaches, analysts, and spectators an understanding of a player's contribution to his team. For example, in basketball we have field goal percentage, offensive rebounds, and assists that indicate a player's offensive prowess and steals, blocks, and defense rebounds that indicate the strength of a player defensively. Analysts quickly notice short-comings of these statistics, e.g. a risky player might have a high number of assists per game, but an even higher number of turnovers per game. To remedy this, more statistics are typically added such as assist-to-turnover ratio. As more statistics are added, box scores quickly become hard to interpret. For this reason, attempts have been made to take combinations of box score statistics to provide meaningful summaries of player contribution. Examples are Dave Heeren's NBA efficiency rating (Heeren, 1988), Dave Berri's player value (Berri, 1999), and John Hollinger's player efficiency rating (Hollinger, 2005). These methods succeed in providing a one number summary of a player's contribution, but, as pointed out in Winston (2009) (and references therein), often have unfortunate implications, e.g. that the worst shooter in the NBA should shoot more to increase their rating.

Since typical box score statistics cannot capture the inherent complexities involved in sports, we consider the *plus-minus statistic*. The plus-minus statistic is simply how many points a player's team outscores their opponents by when that player is playing. Better players allow their team to score more points while forcing their opponents to score fewer and therefore their plus-minus statistic will generally be larger. The main concern with the plus-minus statistic is that it does not account for who is on the court at a given time, but *adjusted plus-minus* does. As implemented by Winston (2009), the adjusted plus-minus statistic numerically solves simultaneously for all players' plus-minus statistics that minimize the squared error in observed data minus predictions. The adjusted plus-minus

*University of California, Santa Barbara, Statistics & Applied Probability, Santa Barbara, CA 93106-3110

statistic has exactly the same interpretation of the original plus-minus statistic, but accounts for who is on the court with whom.

In a similar fashion, Rosenbaum (2004) casts the estimation of an adjusted plus-minus into a linear regression framework. The regression takes into account the difference in points scored while taking into account the players contemporaneously on-court and which side is playing at home. This analysis is restricted to players who play more than 250 minutes across two seasons where a *reference* player replaces those playing less than 250 minutes across the seasons. The analysis provides estimates of plus-minus for each player, but the offensive and defensive ratings have high uncertainty.

We extend this adjusted plus-minus idea to a model based framework where play-by-play data are analyzed for offensive play and defensive play separately and then combined to provide an overall player contribution estimate. In addition, we include all players in the analysis, but propose a model that shrinks player estimates toward the average with the amount of shrinkage being inversely proportional to the amount of time played. The paper continues as follows. Section 2 introduces the play-by-play data used for the analysis throughout. In section 3, we introduce the model and method of parameter estimation. Section 4 provides analyses based on the play-by-play data. Section 5 briefly compares the method to a couple of others, suggests possibilities for extending the model, and discusses applicability to other team sports.

2. NBA Season

The data for this analysis come from a single regular NBA season which comprises 82 games for each of the 30 teams. On average there are $\tilde{2}0$ substitution times which means there are around 20 sets of 10 (not necessarily unique) players that can score points or get scored on. Between each substitution, we record the five players for the away team, the five players for the home team, the amount of time those players are on the court and the number of points that are scored by the away and home teams during that time.

These play-by-play data contain all the known 10-player combinations. These combinations are clear whenever the on-court players and substitutions are clearly identified. Unfortunately, the online data published by ESPN.com and NBA.com do not state substitutions that occur during quarter breaks which results in player uncertainty. In some cases, players on court can be inferred from resulting on-court actions, e.g. shots taken and fouls committed. When the players cannot be inferred, the observations are removed and the resulting play-by-play data is published at `http://basketballgeek.com`.

These data are modified so that free throws are awarded to the set of players in the game before the foul, that resulted in the free throws, was committed. This adjustment is required due to a nuance in the way substitutions are handled during free throws which would assign points to players who have just entered the game (and had no role in the foul being committed) and also would result in scores with zero elapsed time.

Table 1 provides an example of the type of data used for this analysis. In total, we have $\approx$32,000 observations from the 2009-2010 NBA regular season and 443 players.

3. Individual Player Contributions

The method described here considers data on a *squad*, set of five players, basis. At all times two squads are concurrently on the court: a *home squad* and an *away squad* for the home and away teams, respectively. The goal here is to extract from these squads the contribution for each individual player after accounting for the players on-court teammates and opponents.

Table 1: Example of three observations of the 2009-2010 NBA season providing the names of the 5 players on the court for the home and away teams, the number of seconds the players were on the court, and the number of points scored by both the away and home squads.

Observation #	1	2	3
Away squad	Eddie House Kendrick Perkins Marquis Daniels Rasheed Wallace Ray Allen	Eddie House Kevin Garnett Marquis Daniels Paul Pierce Rasheed Wallace	Eddie House Kevin Garnett Paul Pierce Rajon Rondo Ray Allen
Home squad	Anthony Parker Daniel Gibson LeBron James Shaquille O'Neal Zydrunas Ilgauskas	J.J. Hickson Jamario Moon LeBron James Mo Williams Zydrunas Ilgauskas	Anderson Varejao Anthony Parker Daniel Gibson Jamario Moon LeBron James
Seconds	28	47	12
Away points	0	1	2
Home points	2	0	0

3.1 Data generating model

We jointly consider the points scored by the home and away squads. Let y_{iH} (y_{iA}) be the number of points scored by home (away) squad i, $i \in \{1, 2, \ldots, n\}$ where there were a total of n observations over the course of a season. We assume

$$y_i = (y_{iH}, y_{iA}) \overset{ind}{\sim} Po\left(y_{iH} \middle| Ht_i \prod_{p=1}^{P} \theta_p^{h_{ip}} \delta_p^{a_{ip}}\right) Po\left(y_{iA} \middle| At_i \prod_{p=1}^{P} \theta_p^{a_{ip}} \delta_p^{h_{ip}}\right) \quad (1)$$

where $Po(y|\lambda)$ represents the Poisson distribution with mean λ, H (A) is the average number of points per minute for the home (away) team, h_{ip} (a_{ip}) is an indicator of whether player p is on home (away) squad i, i.e. h_{ip} (a_{ip}) $= 1$ if yes, and 0 if no, t_i is the time in minutes between substitutions, and P is the number of players in the NBA. This approach combines the standard generalized linear model (McCullagh and Nelder, 1999, Ch. 6) with a Poisson process (Lawler, 1995, Ch. 3). The relationship with the generalized linear model is apparent if, for example, we take the logarithm of the home Poisson mean and set $\theta_j = \exp(\beta_j)$ where $\log(Ht_i)$ becomes the offset.

To understand the interpretation of this model, first consider the case where $\theta_{ip} = \delta_{ip} = 1\,\forall\, i, p$, i.e. all players are equal both offensively and defensively. In this case, home (away) squad i is expected to score Ht_i (At_i) points. The value for H (A) is the empirical average points per minute for all home (away) teams over the course of the season. Since t_i is the number of minutes of the game played by squad i, the expected points scored by the home (away) squad is proportional to the amount of time played.

In this model, the measure of player p's offensive (defensive) contribution is θ_p (δ_p). If all other parameters are constant, then the multiplicative increase in offensive (defensive) scoring while player p is in the game compared to the average player is θ_p (δ_p). Therefore if $\theta_p > 1$ ($\delta_p < 1$) that player's contribution to the team has a positive offensive (defensive) effect whereas if $\theta_p < 1$ ($\delta_p < 1$) that player's contribution has a negative effect.

3.1.1 Hierarchical structure

We assume independent Gamma priors with common shape and rate parameters for a player's offensive and defensive contributions, i.e. $p(\theta_p) = Ga(\alpha_\theta, \alpha_\theta)$ and $p(\delta_p) = Ga(\alpha_\delta, \alpha_\delta)$ where $p(x) = Ga(\alpha, \beta)$ is the gamma distribution proportional to $x^{\alpha-1}e^{-x\beta}$. Since $E[X] = \alpha/\beta$, the expectation of both priors is 1. Intuitively this means the average NBA player will have θ_p and δ_p equal to 1. In addition, this prior shrinks player contribution parameters toward the average with larger values of α_θ and α_δ resulting in more shrinkage.

We then build a hierarchical model by assuming the non-informative prior $p(\alpha_\theta, \alpha_\delta) = p(\alpha_\theta)p(\alpha_\delta) = \alpha_\theta^{-1}\alpha_\delta^{-1}$. This hierarchical structure allows the data to determine how much shrinkage should be applied by learning the underlying distribution of the θs and δs. For example, if the data indicate that all θs are in a small neighborhood around one, then α_θ will be large leading to small variability in offensive player contributions and therefore a large amount of shrinkage.

3.2 Parameter estimation

The goal of these methods is to obtain the posterior distribution of the model parameters conditional on the observed data, i.e. $p(\theta, \delta, \alpha_\theta, \alpha_\delta|y)$ where $\theta = (\theta_1, \ldots, \theta_p)'$, $\delta = (\delta_1, \ldots, \delta_p)'$, and $y = (y_1, y_2, \ldots, y_n)'$. Using Bayes' Rule and the conditional independence assumptions, we have

$$p(\theta, \delta, \alpha_\theta, \alpha_\delta|y) \propto p(\alpha_\theta, \alpha_\delta) \left[\prod_{p=1}^{P} p(\theta_p|\alpha_\theta)p(\delta_p|\alpha_\delta)\right] \left[\prod_{i=1}^{n} p(y_i|\theta, \delta)\right] \tag{2}$$

where $p(y_i|\theta, \delta)$ is given in equation (1). Since the posterior $p(\theta, \delta, \alpha_\theta, \alpha_\delta|y)$ is not available in an analytically tractable form, we approximate the posterior via a Markov chain Monte Carlo (MCMC) simulation approach, namely a Metropolis-within-Gibbs algorithm (Gelfand and Smith, 1990). Once converged, the MCMC procedure provides draws from the posterior distribution for all parameters.

The two general steps in the algorithm are 1) jointly sampling α_θ and α_δ and 2) cycling through all players jointly sampling θ_p and δ_p for that player. Due to conditional independencies, the first step is accomplished by using two univariate random walks to sample from $p(\alpha_\theta|\theta)$ and $p(\alpha_\delta|\delta)$ where the target distributions are

$$p(\alpha_\theta|\theta) \propto \left(\prod_{p=1}^{P} \theta_p\right)^{\alpha_\theta - 1} \exp\left(-\alpha_\theta \sum_{p=1}^{P} \theta_p\right) \alpha_\theta^{P\alpha_\theta - 1}\Gamma(\alpha_\theta)^{-P}$$

and

$$p(\alpha_\delta|\delta) \propto \left(\prod_{p=1}^{P} \delta_p\right)^{\alpha_\delta - 1} \exp\left(-\alpha_\delta \sum_{p=1}^{P} \delta_p\right) \alpha_\delta^{P\alpha_\delta - 1}\Gamma(\alpha_\delta)^{-P}.$$

The second step involves cycling through players and jointly sampling their offensive and defensive contribution parameters. The full conditional distributions for these parameters for player p, i.e. $p(\theta_p, \delta_p|y, \theta_{-p}, \delta_{-p})$ where x_{-p} indicates the vector x with component p removed, are independently Gamma distributed with updated hyper-parameters, i.e. $Ga(\theta_p|\alpha'_\theta, \beta'_\theta)\, Ga(\delta_p|\alpha'_\delta, \beta'_\delta)$. The hyper-parameters are provided in equation (3) and derived in Appendix A.

$$\begin{aligned}
\alpha'_\theta &= \alpha_\theta + \sum_{i=1}^{n} [h_{ip} y_{iH} + a_{ip} y_{iA}] \\
\alpha'_\delta &= \alpha_\delta + \sum_{i=1}^{n} [h_{ip} y_{iA} + a_{ip} y_{iB}] \\
\beta'_\theta &= \alpha_\theta + \sum_{i=1}^{n} \left[h_{ip} H t_i \prod_{k \neq p} \theta_k^{h_{ik}} \delta_k^{a_{ik}} + a_{ip} A t_i \prod_{k \neq p} \theta_k^{a_{ik}} \delta_k^{h_{ik}} \right] \\
\beta'_\delta &= \alpha_\delta + \sum_{i=1}^{n} \left[h_{ip} A t_i \prod_{k \neq p} \theta_k^{a_{ik}} \delta_k^{h_{ik}} + a_{ip} H t_i \prod_{k \neq p} \theta_k^{h_{ik}} \delta_k^{a_{ik}} \right]
\end{aligned} \tag{3}$$

To understand the interpretation of these updated hyper-parameters, consider $\alpha_\theta = \alpha_\delta = 0$. Then α'_θ (α'_δ) is the points scored by (on) player p's team while player p is on the court and β'_θ (β'_δ) is the expected points scored by (on) player p's team while player p is on the court taking into account all the players that are currently on the court and their current contribution parameters.

Code was written in C and called from the statistical environment R (R Development Core Team, 2009). Convergence was assessed using the method of (Brooks and Gelman, 1997) on multiple chains initialized from dispersed starting locations. This diagnostic is available using the function `gelman.diag` in the package `coda` in R.

To place these multiplicative parameters on a scale compatible with current plus-minus statistics, we multiply the average NBA team points per game by the parameter minus one. For the offensive contribution, this is equivalent to finding the expected number of extra points a team scores if an average player was replaced with this player for an entire game. Similarly, the defensive contribution is the expected number of extra points an opponent would score and the plus-minus contribution is the expected difference in points between the player's team and her opponent's team.

4. Analysis of the 2009-2010 NBA regular season

The 2009-2010 NBA regular season was analyzed using 10 chains with initial values overdispersed relative to the posterior. Each chain had 1,000 burn-in and 1,000 inferential iterations and no lack of convergence was detected. All results presented combine the inferential iterations of all 10 chains.

On average this season, the home team scored 2.11 points per minute (ppm), the away team scored 2.05 ppm, and 99.8 points per game (48 minutes) per team. Figure 1 provides posterior estimates for the hierarchical parameters and predictive distributions for player offensive, defensive, and plus-minus contributions. The predictive distributions indicate that a player contributing an extra 10 points per 48 minutes is an incredible offensive player while a defender reducing his opponent's scoring by 10 points in 48 minutes is incredible defensively. These estimates appear in line with other plus-minus statistics that also have most players in the range -10 to 10.

Table 2 provides estimates and credible intervals for the top 50 NBA players in terms of plus-minus. The table has been sorted according to the lower 2.5%-tile of the plus-minus credible interval. This analysis suggests that Pau Gasol was the most valuable NBA player providing a 10.90 (6.52, 15.18) point differential after adjusting for his teammates and opposing players. This differential occurs because Gasol provides 7.42 (4.16, 10.65)

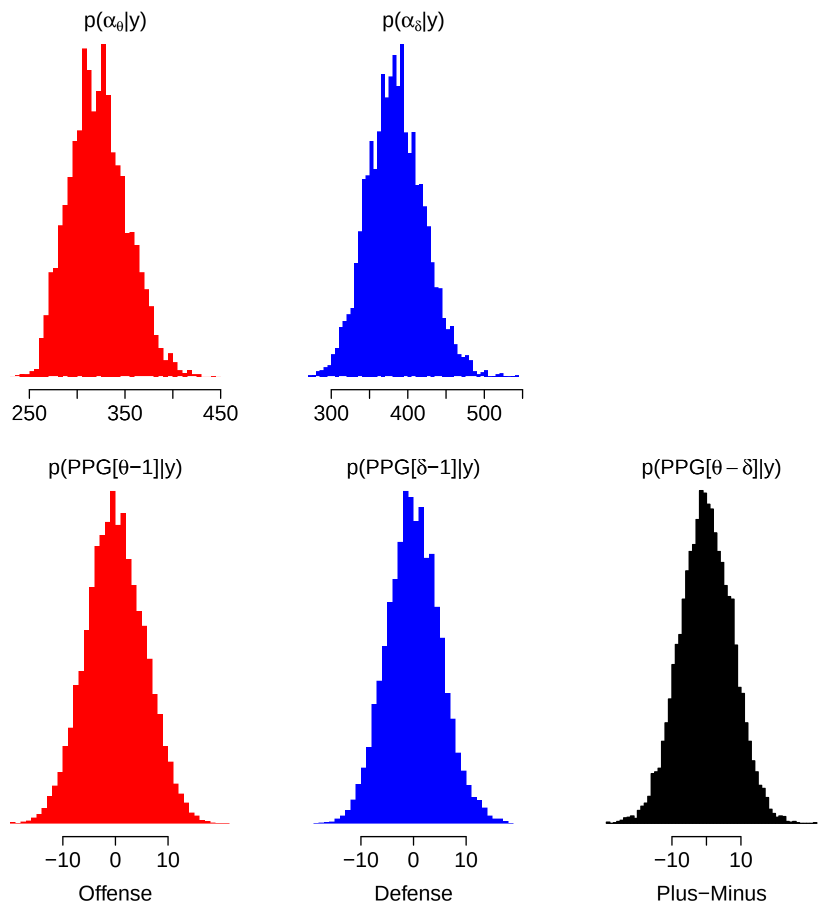

Figure 1: Posterior distribution (top row) for prior parameters and predictive distributions (bottom row) for a new player's offensive, defensive, and plus-minus contributions to a team.

Table 2: Estimates and credible intervals for plus-minus, offensive, and defensive contributions for the top 50 NBA players in the 2009-2010 regular season.

	Player	Plus-minus	Offensive	Defensive
1	Pau Gasol	10.90 (6.52,15.18)	7.42 (4.16,10.65)	-3.46 (-6.39,-0.55)
2	Dwight Howard	9.90 (6.32,13.59)	3.38 (0.65, 6.19)	-6.52 (-9.11,-3.91)
3	Vince Carter	10.07 (6.12,14.09)	1.09 (-1.84, 4.00)	-9.00 (-11.65,-6.24)
4	Andrei Kirilenko	9.66 (4.94,14.35)	2.85 (-0.61, 6.42)	-6.84 (-9.98,-3.49)
5	Kobe Bryant	8.78 (4.87,12.65)	4.00 (1.19, 6.96)	-4.75 (-7.46,-2.10)
6	Anderson Varejao	8.81 (4.64,13.10)	1.74 (-1.33, 4.93)	-7.10 (-9.97,-4.21)
7	Derek Fisher	9.16 (4.54,13.63)	6.54 (3.28, 9.83)	-2.59 (-5.59, 0.64)
8	Deron Williams	8.38 (4.53,12.23)	6.57 (3.75, 9.36)	-1.80 (-4.50, 0.85)
9	Matt Barnes	8.71 (4.53,12.91)	3.77 (0.74, 6.88)	-4.94 (-7.78,-1.97)
10	Paul Millsap	8.45 (4.39,12.58)	5.57 (2.58, 8.64)	-2.87 (-5.72,-0.06)
11	LeBron James	8.03 (4.21,11.78)	4.56 (1.89, 7.33)	-3.45 (-6.04,-0.83)
12	Matt Bonner	9.47 (3.92,14.83)	3.58 (-0.28, 7.56)	-5.82 (-9.48,-2.13)
13	Ron Artest	7.95 (3.91,11.91)	2.68 (-0.23, 5.68)	-5.28 (-7.96,-2.45)
14	Greg Oden	11.66 (3.67,19.75)	4.68 (-1.22,10.80)	-6.96 (-12.37,-1.41)
15	Ray Allen	7.36 (3.52,11.13)	2.25 (-0.45, 5.04)	-5.10 (-7.70,-2.50)
16	Mike Bibby	7.79 (3.48,12.20)	2.44 (-0.68, 5.62)	-5.38 (-8.29,-2.35)
17	Manu Ginobili	7.46 (3.36,11.58)	3.10 (0.04, 6.14)	-4.38 (-7.23,-1.41)
18	Chris Andersen	8.06 (3.28,12.90)	9.85 (6.38,13.33)	1.81 (-1.53, 5.25)
19	Kyle Korver	9.45 (3.21,15.56)	6.95 (2.46,11.62)	-2.51 (-6.71, 1.85)
20	Jameer Nelson	7.71 (3.20,12.24)	1.58 (-1.74, 4.86)	-6.16 (-9.29,-2.94)
21	Kevin Durant	6.70 (3.12,10.24)	3.87 (1.26, 6.49)	-2.85 (-5.29,-0.36)
22	Tim Duncan	7.03 (3.05,10.93)	-0.30 (-3.05, 2.49)	-7.30 (-10.01,-4.55)
23	Kevin Garnett	7.41 (3.01,11.86)	1.75 (-1.48, 5.00)	-5.70 (-8.66,-2.60)
24	Jason Williams	7.55 (2.96,12.26)	4.46 (1.07, 7.99)	-3.11 (-6.36, 0.10)
25	Channing Frye	7.21 (2.91,11.56)	10.23 (7.13,13.51)	3.06 (0.06, 6.08)
26	Al Horford	6.69 (2.86,10.51)	3.40 (0.67, 6.22)	-3.30 (-5.88,-0.58)
27	Brandon Roy	6.97 (2.86,11.19)	4.39 (1.33, 7.47)	-2.60 (-5.42, 0.33)
28	Josh Smith	6.76 (2.83,10.67)	3.95 (1.16, 6.85)	-2.83 (-5.43,-0.04)
29	Ryan Anderson	8.87 (2.77,14.89)	2.21 (-2.11, 6.58)	-6.65 (-10.65,-2.41)
30	Jared Dudley	7.04 (2.52,11.53)	8.47 (5.24,11.81)	1.44 (-1.63, 4.57)
31	Rashard Lewis	6.54 (2.45,10.69)	2.44 (-0.57, 5.47)	-4.15 (-6.93,-1.20)
32	Nick Collison	7.01 (2.35,11.68)	-1.06 (-4.40, 2.32)	-8.07 (-11.24,-4.77)
33	Richard Jefferson	5.95 (2.19, 9.86)	-1.08 (-3.79, 1.77)	-7.05 (-9.64,-4.37)
34	Lamar Odom	6.07 (2.13, 9.99)	1.12 (-1.66, 4.08)	-4.92 (-7.63,-2.15)
35	Kendrick Perkins	6.41 (2.08,10.80)	1.00 (-2.06, 4.20)	-5.41 (-8.40,-2.33)
36	Rajon Rondo	5.40 (1.69, 9.05)	1.40 (-1.28, 4.05)	-4.00 (-6.54,-1.30)
37	Wesley Matthews	6.64 (1.68,11.78)	6.44 (2.85,10.15)	-0.21 (-3.73, 3.33)
38	Jason Kidd	5.32 (1.65, 9.09)	1.97 (-0.76, 4.73)	-3.38 (-6.04,-0.75)
39	Carlos Boozer	5.48 (1.64, 9.40)	3.63 (0.79, 6.48)	-1.81 (-4.49, 0.88)
40	Jamal Crawford	5.60 (1.64, 9.60)	4.17 (1.32, 7.06)	-1.42 (-4.23, 1.38)
41	Dirk Nowitzki	5.32 (1.62, 9.10)	3.68 (0.99, 6.43)	-1.62 (-4.17, 1.08)
42	Steve Nash	5.57 (1.52, 9.66)	11.27 (8.33,14.30)	5.68 (2.84, 8.58)
43	Andrew Bynum	5.73 (1.28,10.37)	1.45 (-1.87, 4.83)	-4.29 (-7.51,-1.02)
44	Shawn Marion	5.24 (1.26, 9.29)	1.60 (-1.26, 4.57)	-3.63 (-6.47,-0.79)
45	Andrew Bogut	5.62 (1.22, 9.98)	0.31 (-2.79, 3.55)	-5.33 (-8.32,-2.25)
46	Joe Johnson	4.96 (1.19, 8.83)	3.00 (0.25, 5.77)	-1.94 (-4.61, 0.75)
47	Rodrigue Beaubois	7.87 (1.03,14.81)	5.89 (0.86,11.05)	-1.99 (-6.65, 2.83)
48	Leandro Barbosa	7.85 (1.01,14.66)	12.04 (7.13,17.11)	4.17 (-0.40, 8.98)
49	Nene Hilario	4.89 (0.99, 8.69)	5.37 (2.57, 8.18)	0.50 (-2.15, 3.26)
50	Dwyane Wade	4.82 (0.98, 8.75)	3.77 (0.98, 6.57)	-1.05 (-3.77, 1.69)

additional points for the Lakers and restricts his opponents to 3.46 (0.55, 6.39) fewer points per game (ppg). Since these intervals do not contain zero, Gasol's offensive, defensive, and plus-minus are all significantly better than average.

The 2009-2010 league MVP was LeBron James who provided an additional 4.56 (1.89, 7.33) ppg for the Cavaliers while holding his opponents to 3.45 (0.83, 6.04) fewer ppg resulting in an 8.03 (4.21, 11.78) ppg differential. The overlapping plus-minus intervals for James and Gasol suggest that the two are not statistically significantly different from each other and an estimate of the probability that Gasol's plus-minus is greater than that for James is 84%. So although this analysis can distinguish players from average, it cannot distinguish among the top 10-20 players in the league. In fact, Greg Oden may be the best plus-minus player in the league owing mainly to his defense. Unfortunately injuries limited him to 21 games in the 2009-2010 season and therefore his contributions have large uncertainties.

The best offensive player in the league (when sorted by lower 2.5%-tile) is Steve Nash who provides 11.27 (8.33,14.30) points per game for the Suns, but also allows his opponents to score an additional 5.68 (2.84, 8.58) points per game. The best defensive player is Raja Bell who limits his opponents to 14.26 (7.12, 21.18) fewer points per game, but reduces his team's scoring by 12.70 (5.25, 19.78) points per game. In the waning moments of a close-fought battle, these are two players a coach would want to consider for strategic offense-defense substitutions.

Complete results for the 2007-08 through 2009-10 regular seasons are available at `http://www.pstat.ucsb.edu/faculty/niemi/research/basketball/player-contributions/`.

5. Discussion

The approach presented in this article is a hierarchical model where player contributions for both offense and defense affect the points per minute scored by the away and home teams. The contribution parameters as well as their distributions are estimated using a Markov chain Monte Carlo approach. Compared to other plus-minus statistics, the advantages of this hierarchical statistical modeling approach are uncertainty quantification, parameter shrinkage to the average player, estimation of the underlying distribution of player contributions, and separate assessment of a player's offensive and defensive contribution to the team.

In this article, the uncertainty of player contribution parameters as well as their prior distribution are quantified. The parameters for the prior distribution have little uncertainty, i.e. the difference between $Ga(300, 300)$ and $Ga(400, 400)$ is minor on the scale of the multiplicative effects we are estimating. In contrast, there is considerable uncertainty in player contribution estimates especially if interest centers on distinguishing individual players. Using one season of data, the best players can be clearly distinguished from average players, but difficulty arises in separating the top players. A next step to decrease this uncertainty would be to create a dynamic model for the player contribution parameters to allow for information to be shared across and, perhaps, within seasons.

An appealing aspect of utilizing a prior distribution for the player contribution parameters is shrinking estimates toward average. In the full conditional distributions given in equation (3), the prior has the effect of adding α_θ and α_δ points for a player's team and his opponent's respectively. This has little effect on players like James who were on the court for $\approx$12,000 total points, but has a large effect on players like Bell who were only on the court for $\approx$100 total points. On one hand, this makes Bell's defensive ability that much more impressive since his defensive contribution was shrunk to the average by a large amount. On the other hand, even Bell might have gotten lucky while he was on the court and therefore a better estimate of his ability is to shrink it toward the average.

Finally, many plus-minus statistics only provide an estimate of the point differential a player creates when they are in the game relative to an average player. In contrast, analyses using player statistics often have misleading interpretations, e.g. the worst shooter in the league should shoot more to increase their contribution. The approach described here takes a middle ground where players themselves are modeled rather then their statistics, but information about their offensive and defensive contributions can be inferred.

References

Berri, D. J. (1999), "Who is 'most valuable'? Measuring the player's production of wins in the National Basketball Association," *Managerial and Decision Economics*, 20, 411–427.

Brooks, S. P. and Gelman, A. (1997), "General methods for monitoring convergence of iterative simulations," *Journal of Computational and Graphical Statistics*, 7, 434–455.

Gelfand, A. E. and Smith, A. F. M. (1990), "Sampling-based approaches to calculating marginal densities," *Journal of the American Statistical Association*, 85, 398–409.

Heeren, D. (1988), *The Basketball Abstract*, Prentice Hall, Englewood Cliffs, NJ.

Hollinger, J. (2005), *Pro Basketball Forecast*, Potomac Books Inc., Dulles, VA.

Kubatko, J., Oliver, D., Pelton, K., and Rosenbaum, D. (2007), "A starting point for analyzing basketball statistics," *Journal of Quantitative Analysis in Sports*, 3, 1–22.

Lawler, G. F. (1995), *Introduction to Stochastic Processes*, Chapman & Hall/CRC.

McCullagh, P. and Nelder, J. A. (1999), *Generalized Linear Models*, Chapman & Hall/CRC, 2nd edn.

R Development Core Team (2009), *R: A Language and Environment for Statistical Computing*, R Foundation for Statistical Computing, Vienna, Austria, ISBN 3-900051-07-0.

Rosenbaum, D. T. (2004), "Measuring how NBA players help their teams win." `http://www.82games.com/comm30.htm`.

Winston, W. (2009), *Mathletics: How Gamblers, Managers, and Sports Enthusiasts Use Mathematics in Baseball, Basketball, and Football*, Princeton University Press.

A. Full conditional distributions

Assume $\theta_p, \delta_p \stackrel{iid}{\sim} Ga(\theta_p; \alpha_\theta, \beta_\theta) Ga(\delta_p, \alpha_\delta, \beta_\delta)$ and data observed according to equation (1).

$$
\begin{aligned}
&p(\theta_p, \delta_p | y, \theta_{-p}, \delta_{-p}) \\
&\quad \propto \left[\prod_{i=1}^{n} Po\left(y_{iH}; Ht_i \prod_{p=1}^{P} \theta_p^{h_{ik}} \delta_p^{a_{ik}} \right) Po\left(y_{iA}; At_i \prod_{p=1}^{P} \theta_p^{a_{ik}} \delta_p^{h_{ik}} \right) \right] p(\theta, \delta) \\
&\quad \propto \prod_{i=1}^{n} \left(Ht_i \prod_{k=1}^{P} \theta_k^{h_{ik}} \delta_k^{a_{ik}} \right)^{y_{iH}} \exp\left(-Ht_i \prod_{k=1}^{P} \theta_k^{h_{ik}} \delta_k^{a_{ik}} \right) \\
&\quad \cdot \prod_{i=1}^{n} \left(At_i \prod_{k=1}^{P} \theta_k^{a_{ik}} \delta_k^{h_{ik}} \right)^{y_{iA}} \exp\left(-At_i \prod_{k=1}^{P} \theta_k^{a_{ik}} \delta_k^{h_{ik}} \right) \\
&\quad \cdot \theta_p^{\alpha_\theta - 1} \exp(-\theta_p \beta_\theta)\, \delta_p^{\alpha_\delta - 1} \exp(-\delta_p \beta_\delta) \\
&\quad \propto \prod_{i=1}^{n} \left[\left(\theta_p^{h_{ip}} \delta_p^{a_{ip}} \right)^{y_{iH}} \right] \exp\left(-\theta_p \sum_{i=1}^{n} h_{ip} Ht_i \prod_{k \neq p} \theta_k^{h_{ik}} \delta_k^{a_{ik}} - \delta_p \sum_{i=1}^{n} a_{ip} At_i \prod_{k \neq p} \theta_k^{h_{ik}} \delta_k^{a_{ik}} \right) \\
&\quad \cdot \prod_{i=1}^{n} \left[\left(\theta_p^{a_{ip}} \delta_p^{h_{ip}} \right)^{y_{iA}} \right] \exp\left(-\theta_p \sum_{i=1}^{n} a_{ip} At_i \prod_{k \neq p} \theta_k^{a_{ik}} \delta_k^{h_{ik}} - \delta_p \sum_{i=1}^{n} h_{ip} Ht_i \prod_{k \neq p} \theta_k^{a_{ik}} \delta_k^{h_{ik}} \right) \\
&\quad \cdot \theta_p^{\alpha_\theta - 1} \exp(-\theta_p \beta_\theta)\, \delta_p^{\alpha_\delta - 1} \exp(-\delta_p \beta_\delta) \\
&\quad \propto \theta_p^{\sum_{i=1}^{n} h_{ip} y_{iH} + a_{ip} y_{iA}} \exp\left(-\theta_p \sum_{i=1}^{n} h_{ip} Ht_i \prod_{k \neq p} \theta_k^{h_{ik}} \delta_k^{a_{ik}} + a_{ip} At_i \prod_{k \neq p} \theta_k^{a_{ik}} \delta_k^{h_{ik}} \right) \\
&\quad \cdot \delta_p^{\sum_{i=1}^{n} h_{ip} y_{iA} + a_{ip} y_{iH}} \exp\left(-\delta_p \sum_{i=1}^{n} h_{ip} At_i \prod_{k \neq p} \theta_k^{a_{ik}} \delta_k^{h_{ik}} + a_{ip} Ht_i \prod_{k \neq p} \theta_k^{h_{ik}} \delta_k^{a_{ik}} \right) \\
&\quad \cdot \theta_p^{\alpha_\theta - 1} \exp(-\theta_p \beta_\theta)\, \delta_p^{\alpha_\delta - 1} \exp(-\delta_p \beta_\delta) \\
&\quad \propto \theta_p^{\alpha_\theta + \sum_{i=1}^{n} [h_{ip} y_{iH} + a_{ip} y_{iA}] - 1} \exp\left(-\theta_p \left[\beta_\theta + \sum_{i=1}^{n} h_{ip} Ht_i \prod_{k \neq p} \theta_k^{h_{ik}} \delta_k^{a_{ik}} + a_{ip} At_i \prod_{k \neq p} \theta_k^{a_{ik}} \delta_k^{h_{ik}} \right] \right) \\
&\quad \cdot \delta_p^{\alpha_\delta + \sum_{i=1}^{n} [h_{ip} y_{iA} + a_{ip} y_{iH}] - 1} \exp\left(-\delta_p \left[\beta_\delta + \sum_{i=1}^{n} h_{ip} At_i \prod_{k \neq p} \theta_k^{a_{ik}} \delta_k^{h_{ik}} + a_{ip} Ht_i \prod_{k \neq p} \theta_k^{h_{ik}} \delta_k^{a_{ik}} \right] \right)
\end{aligned}
$$