Spring 2024

Defensive Impact Wins: Developing a New Method to Rate Individual Defense in NBA Games

Dylan J. Stiles
University of New Hampshire, Durham

Defensive Impact Wins: Developing a New Method to Rate Individual Defense in NBA Games

Dylan Stiles

Abstract

With the analytics revolution in sports in the past 20 years, it seems that everything that can be quantified is. In basketball though, trying to break the game down into a set of numbers comes with a unique problem. While we've come up with a good set of advanced numbers to measure offensive efficiency, defense is fundamentally harder to quantify. The game is played five on five, but it has often been popular or convenient to model defense as a set of five one on one games. As defenses became more complex into the 2010s, this methodology became more insignificant. Other metrics seemed to rely too much on the idea that basketball is a team sport, and have a hard time identifying good defensive players on bad defensive teams, and vice versa. Previous examples of defensive analytics include simply counting points per possession by each defender, defensive rating (points allowed per 100 possessions while a certain defender is in the game), and defensive win shares. I first collected team-level defensive data from the 2022-23 season split by play type from the NBA's website. Using binomial regression, I regressed wins and losses against stop rates, foul rates, and forced turnover rates, creating 27 possible models. I then took all of the slopes that were both significant and allowed (did not say something like "forcing more turnovers is associated with fewer wins") and multiplied them by the players' rates above league average. After adding those up, I centered it by subtracting the mean so that a league average defender is worth zero wins. There were issues that hurt the success of this

model, as just 12 of the 27 slopes I fit actually made it in. Six of these slopes were missing because the playtype numbers are not stored on NBA.com at the player level. Some others were disallowed or non significant, things that could have possibly been prevented by using more than one season of data. Unlike many metrics, this doesn't seem to be applicable to the team level at all - elite defensive teams had their entire roster add up to about +5.5 wins, and elite defensive teams usually win about 50 games.

Introduction and Literature Review

There are problems with any attempt to quantify the defensive contributions of an individual player in basketball. The modern NBA keeps a lot of data. We can see shots made and missed, rebounds, assists, and turnovers for every player in the league. These are "counting stats." They can be collected simply by watching the games – and they are, by official scorekeepers in every NBA arena. Using them, we can quantify offense and even begin to isolate the effects of single players. No such method currently exists for defense, because these counting stats do not exist. In fact, it is common to think they can't exist.

Trying to fairly give credit and blame to individual players on defense is like trying to count up things that didn't happen. Every driving lane that closes, every pass that isn't made, every contested look turned down, and anything else that happens in one possession will lead to either one make, one miss, a foul, or a turnover. The following will be an effort to count these plays prevented, to see which defenders bring their teams the most wins.

The Wages of Wins by Berri, Schmidt, and Brook

This sports economics book contains analysis of the data from many sports, but most important here is the Win Score model. This model isn't the most useful for defensive studies as it contains only counting stats and the data is mostly offensive, but it does show the framework for constructing these metrics. Berri found the correlation for teams between wins and many stats including points, rebounds, field goal attempts, turnovers, and fouls. To simplify the model, Berri compared everything to points, saying that a rebound was nearly the same as a point, but an assist was only worth half as much. The major benefit of this model was imposing a penalty for shot attempts, not just shot misses. This was justified by the fact that shooting the ball spends a possession. To have made that expense worth it, you need to hit enough of your shots. By adding and subtracting appropriately (adding points, adding half assists, subtracting field goal attempts) Berri arrives at Win Score. This is a simplified version of his more well known stat, Wins Produced and Wins Produced per 48 Minutes. The latter of these is the most analogous to what I am attempting to do – a metric that rates individuals by their box-score contributions to winning, normalized for playing time. Of course, because these are box-score contributions, they are context-driven. An example Berri cites is that Kobe Bryant's reduction in productivity in 2004-05 might be associated with the Lakers lack of a second ball handler (or "true point guard"). This put the ball in Bryant's hands more, resulting in more turnovers and lower wins metrics. While not a direct comparison for the eventual defensive impact statistic, this finally provided me with some type of analytical plan. Regress team statistics I was interested in against wins, for up to as many seasons as NBA data provided for. Use these slopes to provide weights for different stats. Add up those stats, weighted, per player, and divide by minutes played.

Characterizing the Spatial Structure of Defensive Skill in Professional Basketball by Franks, Miller, Bornn, and Goldsberry

This is a 2015 study published in *The Annals of Applied Statistics* that illustrates exactly why it is so hard to measure defense in basketball (and other goal-based sports like hockey and soccer) The goal of this method is to solve what the authors called the *identifiability problem* of defense in the NBA. This is essentially the problem I proposed in the introduction: who gets credit for the fact that the Warriors' possession ended in a miss? The model proposed here is intended to quantify shot frequency and shot efficiency. This is a high-level model that first models the position of the defender relative to the position of the offensive player and the hoop. All the interactions of a possession are summarized in a hidden Markov model. This is a transitive model meant to convey switch interactions, so any defender can switch to guard any offensive player at any time. The EM algorithm is used to fit the hidden Markov model for 30 possessions. This model tells the authors that defenders tend to be positioned 2/3 of the distance from the hoop to their man, and slightly towards the ball. Defenders get closer to their man when he has the ball, which also matches conventional wisdom. This process was repeated 100 times, using 30 possessions in 100 games, and the weights were found to be stable. This paper was quote advanced as it was based on the 2013-14 season, when switches were not as common as they are today. It also is not a model I plan to try to replicate, as I don't have access to the same level of 25fps tracking data for all 10 players and the ball It did a good job of bringing out expected defensive shot charts for individual players. However, the authors point out that this does not remove entirely the confounding effect that emerges from a team sport. This is the challenge of any meaningful model of defense.

Counterpoints: Advanced Defensive Metrics for NBA Basketball by Franks, Miller, Bornn, and Goldsberry

This paper, presented at the 2015 MIT Sloan Sports Analytics Conference, seems to be an extension of the previous work by these authors. Using the previous spatial relations model for defense that produced defensive shot charts. The authors used this data along with matchup stats for the entire 2013-14 season. “We know estimated defensive matchups for every offensive player who converted a field goal during the 2013-14 season,” they write. This allowed them to produce five new advanced metrics. Volume Score represents the total number of field goal attempts faced by a defender. Disruption Score is interpreted as the degree to which a defender is able to make his man miss. Defensive shot charts were covered in the previous paper, but represent relative scoring efficiency by area against individual players. Shots Against is a weighted average of shots taken against defenders, per 100 possessions. Finally, Counterpoints is a weighted average of points against defenders per 100 possessions. This metric is important because it assigns defensive blame proportional to time spent defending the scorer and is in a real unit of points. Although this is a well-specified metric, some of the player rankings made me question its precision in this form. (There is no season in which Jrue Holiday was the worst guard defender in the league.) The authors acknowledge themselves this this measure is limited by the fact that we are assuming we have full information about defensive schemes on a play-to-play basis.

A comprehensive review of plus-minus ratings for evaluating individual players in team sports by Hvattum

This paper covers the origins and efficacy of plus-minus metrics across different sports. Plus-minus metrics are a shockingly recent addition to NBA data collection, first becoming popular around 2003 and with a version being sold to the Dallas Mavericks around this time. Even back then, the flaws were apparent. Everyone knew simple plus-minus was context-dependent, biased for both team and opponent quality. Additionally, the metric said nothing about what an individual player contributed to his team during his time on the floor, only whether his team outscored their opponent in these minutes. This is not directly comparable to any model I would like to build – I started by examining some of the adjusted plus-minus stats that exist and found that they are all too dependent on context, even when taking in counting stats. This paper covers some of these models, such as WINVAL (likely the model Winston and Sagarin sold to Dallas), adjusted plus-minus (Rosenbaum) and how they are calculated. This paper also mentions ESPN's Real Plus-Minus (RPM) but this is a proprietary formula. The paper concludes by saying that many of these plus-minus ratings are made with no consideration for the previous iterations, but they do solve some of the inadequacy of pure plus-minus ratings.

Win Shares and Rookie Contracts in the NBA by Lucas Kobat

This is a reviewed honors thesis from the University of South Dakota that examines the relationship between salary and win shares, and concludes that the best way to acquire producers of win shares without spending much money is through the draft. This paper is not closely related to what I am studying, except in one way. This thesis contains, on page 22, what looks like an exact step-by step guide to calculating defensive win shares per player (DWS). The method starts with calculating player defensive ratings by the Dean Oliver method (which appears to differ from box-score defensive ratings). From there, one simply calculates Marginal Defense and Marginal Points Per Win for each player. The quotient found by dividing those is a player's defensive win shares. Here, however, we see why this stat too is imperfect. Dean Oliver's player defensive ratings contain the team's defensive rating, which invariably will penalize players for being on bad defensive teams, and reward players for being on good ones. In some ways this is a good thing. The six highest defensive win share seasons of all time all belong to Bill Russell, widely considered one of the greatest defenders ever. But the association with team performance is too high. Scrolling Basketball Reference's top 250 DWS seasons of all time, we see nearly exclusively players from playoff teams. It also appears some players are rewarded for sharing the floor with Hall of Fame teammates. Towards the back of that list, for example, we find Rasho Nesterovic's 2003-04 season. Nesterovic played center for the Spurs that season, where he lined up next to Tim Duncan. He was certainly not a bad defender, but in only one other season did he pass 55% of his 2003-2004 DWS. And he spent most of his career playing with either Duncan or Kevin Garnett, two of the ten best defensive players ever. The Spurs were a great defensive team in '04 with him, but when he left town prior to the 2006-07 season, their defense was just as good. I did not start writing looking to pick a fight with Rasho Nesterovic, I simply scrolled down and clicked on a name I did not recognize. Finding such a clear example of team quality being rewarded so quickly just illustrates some possible shortcomings in DWS.

Sabermetric Analysis: Wins-Above-Replacement by Karl Hendela

The last source covered extensively in this review has nothing to do with basketball. However, the model I would most like to emulate that I have examined is Berri's Wins Produced, which has a unit of wins. It's great to examine defense in terms of preventing points, but ultimately we care so much about defense because it leads to wins. Additionally, a point in 2024 isn't worth the same as a point in 2004, as scoring has increased dramatically in the past 20 years. Modeling to produce wins is the best way to interpret someone's defensive contributions – how many wins was your defense worth? For another reference, we examine the most famous wins added sabermetric of all: baseball's Wins Above Replacement (WAR). For as much handwringing about WAR as you can see by interacting with baseball media, it's a great stat. Teams with the most WAR on their roster have the most talent on their roster, and they generally win more games. I recently entered the WAR totals for all 30 teams in the past 5 full seasons using the two major formulas for WAR into a spreadsheet for an unrelated project. The highest total I saw was from the 2019 Astros, who came one game shy of winning a championship. The lowest I saw, depending on formula, was either the 2023 Athletics or 2023 Rockies. Both teams were uncompetitive. But why are there two formulas? Two major baseball data sites, Baseball-reference.com and Fangraphs.com, use their own versions of WAR. Both are somewhat explained somewhere on the site, but the actual formulas are proprietary. This paper, published in Seton Hall's Journal of Undergraduate Research in 2020, tries to fix that. This method, called SHU-WAR, is calculated mostly the same as bWAR and fWAR. Batting, baserunning, pitching and fielding runs are calculated with adjustments for park factors and position, as well as league averages. These are added up, added to the replacement player adjustment, and divided by twice the league average runs per game. While the formulas used for bWAR and fWAR are private, SHU-WAR had over a 90% correlation with both. This is comparable to the correlation between bWAR and fWAR.

The defensive wins metric I will be fitting can't work exactly like WAR, because basketball is a different game. Still, a clearly explained process that produces a wins-added based measure is my goal, so a study of WAR felt like a necessary inclusion.

Data Used

The metric I am fitting is based most closely on Berri's Wins Produced. To fit this, we will need both team stats and player stats. The data in this study was provided on NBA.com, from the 2022-23 NBA season. The team stats available are defensive metrics against nine different playtypes.

- **Isolation –** The simplest plays in basketball, when one offensive player tries to score on one defensive player with his teammates not directly involved in the play.
- **Handoff** – A common action where one player, usually a big man, hands the ball to a ball handler behind him to either set or fake a screen.
- **Pick and Roll** – A very common action where one player sets a screen for the ball handler, usually to force the defense to switch assignments, and then "rolls" to the basket. This action is split on the NBA.com stats page by whether the ball handler keeps the ball or passes it to the roll man.
- **Offscreen –** Somewhat catch-all term for plays involving screens. Not able to be part of this metric in current form.
- **Post-Up –** Plays in which the ball is dumped into a player, usually a big man, close to the basket.
- **Spot-**Up – Plays in which a player catches the ball to quickly shoot it.
- **Putback -** Plays in which an offensive rebound leads directly to a second shot attempt
- **Transition** – Fullcourt, fast break plays by the team that just gained possession. Usually occurs off of turnovers, sometimes off defensive rebounds by fast players.

The putback and transition plays were not kept at the player level, so although there were regressions fitted for them, they cannot be part of the metric and will not be reported. For all other plays, we are interested in only three variables, which are either provided or easily transformed from the NBA website.

- **Stop Rate –** The percentage of plays on which the offense does not score. The basic measure of success for all plays.
- **Shooting Foul Rate –** The percentage of plays on which the offense draws a shooting foul and shoots free throws. These are worse then simply not getting a stop, since free throws lead to far more expected points than other shot attempts.
- **Turnover Rate –** The percentage of plays on which the defense forces a turnover. These are better than other stops, since not only is the offense not scoring, but the team forcing a turnover is more likely to score on its next play.

Now, we fit separate binomial GLMs for each term, for each playtype. The response is wins and losses, as both are needed in R to fit the model. We record all the fitted slopes, but the metric will only be able to use those which are of the correct sign and significant. Stop rate and turnover slopes must be positive, and foul slopes must be negative. It does not make sense to conclude that forcing a turnover is expected to produce a loss.

Results - Slopes

- **Handoff**
  - Stoprate – 0.04642, significant
  - Turnover – -0.0005736, negative nonsignificant
  - Foul - -0.1153, significant
- **Isolation**
  - Stoprate - 0.04730, significant
  - Foul – -0.04322, nonsignificant
  - Turnover – 0.05237, significant

- **Offscreen**
    - Stoprate - -0.01016, negative nonsignificant.
    - Foul – 0.09325, positive significant
    - Turnover - -0.06448, negative significant
        - All offscreen variables will be excluded
- **Pick and Roll Ball**
    - Stoprate – 0.14631, significant
    - Foul - -0.17094, significant
    - Turnover – 0.002575, nonsignificant
- **Pick and Roll Roll**
    - Stoprate – 0.06203, significant
    - Foul - -0.10368, significant
    - Turnover - -0.007045, negative nonsignificant
- **Post-Up**
    - Stoprate – 0.03822, significant
    - Foul - -0.006812, nonsignificant
    - Turnover - -0.01515, negative nonsignificant
- **Spot-Up**
    - Stoprate – 0.14577, significant
    - Foul - -0.12208, significant
    - Turnover – 0.09622, significant

Fitting The Metric

To fit this metric and its multiple versions we follow several steps. We will describe the steps for the first version of the metric in detail.

1. Import player playtype spreadsheets, as opposed to the team spreadsheets used to fit the models.
2. Convert player stop rates, foul rates, and turnover rates to deviations by subtracting the average of the corresponding team statistics.
3. Multiply the significant and allowed GLM slopes by the deviations. This leaves us with 12 numbers for each player, which we add up.
4. This is the "season defensive wins" number, the amount of wins a player would be expected to produce if they played defense for the entire season.
5. This is not a meaningfully interpretable number, so we multiply by the proportion of minutes played in the season for each player, done by multiplying by the number of minutes they played and then dividing by 82*48, which is approximately the number of minutes in a season excluding overtime.
6. Now, center the metric by subtracting its average. This means that a league average defender is worth zero wins.
7. This can be interpreted as, "the number of wins added on defense in 2022-23."

There are problems with missing data in this metric before we even examine the results. Through missing data from the NBA.com page, nonsignificant slopes, and illegal slopes, we are left with 12 of 27 desired parameters in this version of the metric. This is the most mathematically honest version of the metric.

Results - Version 1

Player	Team	Dwins1
Draymond Green	GSW	3.481358
Shai Gilgeous-Alexander	OKC	3.145087
Tyler Herro	MIA	3.098615
CJ McCollum	NOP	3.071841
Kristaps Porzingis	WAS	2.919318
Jevon Carter	MIL	2.845542
Zach LaVine	CHI	2.747877
D'Angelo Russell	MIN	2.73098
Jayson Tatum	BOS	2.687579
Bam Adebayo	MIA	2.62589
Bobby Portis	MIL	2.52449
Anthony Edwards	MIN	2.50348
Tyus Jones	MEM	2.496806
Stephen Curry	GSW	2.405139
Trae Young	ATL	2.324286
Austin Reaves	LAL	2.306509
Giannis Antetokounmpo	MIL	2.296053
Tobias Harris	PHI	2.228009
Jarrett Allen	CLE	2.155035

These are the top 20 players in 2022-23 for this version of the metric, and we can quickly see problems. Draymond Green, at the top of the list, is going to be a Hall of Fame player mostly due to his defense. Past him though, Tyler Herro at 3 and CJ McCollum at 4 are regarded as some of the worst defenders in the league. Other famously bad defenders are Zach Lavine and Trae Young. We do see some great ones like Kristaps Porzingis, Jayson Tatum, Bam Adebayo, and Giannis Antetokounmpo. An interesting trend about the bad defenders on this list is that they are all guards. This leads us to consider whether there is a positional bias in this model, which will be investigated further.

Results – Version 2

For the second version of this model, we include the four slopes which are allowed but not significant.

Player	Team	Dwins2
Draymond Green	GSW	3.628902
CJ McCollum	NOP	3.299977
Tyler Herro	MIA	3.158675
Shai Gilgeous-Alexander	OKC	3.080251
D'Angelo Russell	MIN	2.994847
Kristaps Porzingis	WAS	2.936256
Jayson Tatum	BOS	2.880734
Zach LaVine	CHI	2.812577
Bam Adebayo	MIA	2.673954
Tyus Jones	MEM	2.667613
Jevon Carter	MIL	2.648531
Bobby Portis	MIL	2.637545
Anthony Edwards	MIN	2.521689
Trae Young	ATL	2.443525
Stephen Curry	GSW	2.431781
Giannis Antetokounmpo	MIL	2.348075
Tobias Harris	PHI	2.297712
Austin Reaves	LAL	2.270053
Jarrett Allen	CLE	2.238331

This gives us different numbers, but most of the same players. We have some high-volume guards who are known to play a lot of minutes, have the ball a lot, and play very little defense. This spawned two final versions of this model.

The Fouls Problem

In trying to fit this model, we have to consider two statements which are absolutely true, but contradict each other in ways that make modeling harder. First, committing a foul is bad. It should have a negative effect on your chances of winning the game. It gives your opponent some of the easiest points they will get. Second, this does not mean that the players that foul the most are bad defenders, and it does not mean that the players that foul the least are good defenders. To foul someone you need to usually be trying to guard them. Additionally, centers tend to foul more than guards regardless of how good or bad they are at contesting shots, since they face more shots at the rim, where most fouls are drawn. This identifies a possible systemic flaw in this model – guards who play very little defense will commit very few fouls, have a foul rate that is far better than league average, and have inflated scores here. We now fit two more versions of this model where we remove fouls altogether.

Results – Versions 3 and 4

In these versions, we start with versions 1 and 2 respectively, and remove all foul variables.

Player	Team	Dwins3
Draymond Green	GSW	2.810798
D'Angelo Russell	MIN	2.40101
CJ McCollum	NOP	2.391897
Jevon Carter	MIL	2.088245
Shai Gilgeous-Alexander	OKC	1.945033
Zach LaVine	CHI	1.851948
Bam Adebayo	MIA	1.813462
Stephen Curry	GSW	1.504283
Ayo Dosunmu	CHI	1.447113
Tobias Harris	PHI	1.443734
Bruce Brown	DEN	1.427785
Tyler Herro	MIA	1.427022
Kristaps Porzingis	WAS	1.394207
Giannis Antetokounmpo	MIL	1.333964
Isaiah Joe	OKC	1.307855
Jalen Green	HOU	1.263786
Dyson Daniels	NOP	1.26281
Caleb Martin	MIA	1.237585
Jalen McDaniels	CHA	1.212716

Player	Team	Dwins4
Draymond Green	GSW	2.472821
D'Angelo Russell	MIN	2.067952
CJ McCollum	NOP	2.065308
Jevon Carter	MIL	1.750056
Shai Gilgeous-Alexander	OKC	1.62142
Zach LaVine	CHI	1.521478
Bam Adebayo	MIA	1.491265
Stephen Curry	GSW	1.173463
Ayo Dosunmu	CHI	1.115966
Tobias Harris	PHI	1.112939
Bruce Brown	DEN	1.099911
Tyler Herro	MIA	1.097595
Kristaps Porzingis	WAS	1.058767
Giannis Antetokounmpo	MIL	0.991129
Isaiah Joe	OKC	0.96746
Dyson Daniels	NOP	0.933474
Jalen Green	HOU	0.927215
Caleb Martin	MIA	0.918293
Jalen McDaniels	CHA	0.891716

We see some minor corrections to this problem now. Trae Young has left the top 20 and Tyler Herro fell down multiple spots. Overall though, we still see some problems.

Conclusions

This is a good first step towards fitting a model to better understand NBA defense. I have never seen a method like this used for a defensive metric, and finalizing a successful model here could fill that gap. There are a few ideas on how this model could be updated. First, including data from multiple seasons would likely stabilize the slopes. This could lead to fewer illegal or nonsignificant coefficients. Second, we could attempt to remove this model's guard-bias by doing some positional adjustments. We could do this a few ways. First, we could use a different parameter in place of stoprate. Options include effective field goal percentage and points per posession, the latter being likely preferred. We could also adjust the way fouls are scored here. Instead of subtracting the team league average for all players, we could subtract the point guard player average foul rate for point guards, the center average foul rate for centers, etc. Position is not a column in the NBA.com outputs, so we would need to manually input it. This would likely give us the best version of this model we could create with standard NBA.com data and is a step towards separating player defense from their team's success.

Works Cited

Berri, D. J., Schmidt, M. B., & Brook, S. L. (2006). *The Wages of Wins: Taking Measure of the Many Myths in Modern Sport* (1st ed.). Stanford University Press. Chapters 6-7

Alexander Franks, Andrew Miller, Luke Bornn, Kirk Goldsberry. "Characterizing the spatial structure of defensive skill in professional basketball." The Annals of Applied Statistics, 9(1) 94-121 March 2015. https://doi.org/10.1214/14-AOAS799

Hvattum, Lars Magnus. “A Comprehensive Review of Plus-Minus Ratings for Evaluating Individual Players in Team Sports.” *International Journal of Computer Science in Sport*, vol. 18, no. 1, July 2019, pp. 1–23. *DOI.org (Crossref)*, https://doi.org/10.2478/ijcss-2019-0001.

Kobat, Lucas R., "Win Shares & Rookie Contracts in the NBA" (2018). *Honors Thesis*. 14. https://red.library.usd.edu/honors-thesis/14

Hendela, Karl (2020) "Sabermetric Analysis: Wins-Above-Replacement," Locus: The Seton Hall Journal of Undergraduate Research: Vol. 3 , Article 7

The Official Site of the NBA for the Latest NBA Scores, Stats & News. | NBA.Com. https://www.nba.com/. Accessed 9 May 2024.

“Basketball Statistics & History of Every Team & NBA and WNBA Players.” Basketball-Reference.Com, https://www.basketball-reference.com. Accessed 9 May 2024.