AN APPROXIMATE DYNAMIC PROGRAMMING APPROACH TO DETERMINE THE OPTIMAL SUBSTITUTION STRATEGY FOR BASKETBALL

by

David W. Hughes
A Dissertation
Submitted to the
Graduate Faculty
of
George Mason University
in Partial Fulfillment of
The Requirements for the Degree
of
Doctor of Philosophy
Systems Engineering and Operations Research

Committee:

___________________________________ Dr. John Shortle, Dissertation Director

___________________________________ Dr. Rajesh Ganesan, Committee Member

___________________________________ Dr. Andrew Loerch, Committee Member

___________________________________ Dr. Guoqing Diao, Committee Member

___________________________________ Dr. Ariela Sofer, Department Chair

___________________________________ Dr. Kenneth S. Ball, Dean, Volgenau School of Engineering

Date:______________________________ Spring Semester 2017
George Mason University
Fairfax, VA

An Approximate Dynamic Programming Approach to Determine the Optimal Substitution Strategy for Basketball

A Dissertation submitted in partial fulfillment of the requirements for the degree of Doctor of Philosophy at George Mason University

by

David W. Hughes
Master of Science
Stanford University, 2009
Bachelor of Science
United States Military Academy, 2001

Director: John Shortle, Professor
Department of Systems Engineering and Operations Research

Spring Semester 2017
George Mason University
Fairfax, VA

DEDICATION

This dissertation is dedicated to my beautiful, loving wife Meghan and my three amazing children – Alyssa, JD and Kelly.

ACKNOWLEDGEMENTS

I would like to thank my family and friends who helped make this possible. A special thank you goes out to LTC Sam Huddleston for encouraging me to pursue a PhD in something I'm passionate about: sports. To my good friend, Omar Mance – I am so grateful that you linked me up with a Division I college basketball program so that I could get access to coaches, players, and all their data. To my dissertation committee members, Dr. Ganesan, Dr. Loerch, and Dr. Diao – thank you for meeting with me whenever I needed guidance and for your insightful and challenging comments and questions. To my dissertation director, Dr. Shortle – I really appreciate your dedication to my research. You consistently met with me, week in and week out, for the past two years. I would not have finished if it were not for you! Finally, to my wife and children – Thank you for putting up with the long work weeks and the late nights. Your love and support helped get me through these past three years.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

LIST OF EQUATIONS

LIST OF ABBREVIATIONS AND/OR ACRONYMNS

Aerobic Zone 1 .... AZ1
Aerobic Zone 2 .... AZ2
Anaerobic Threshold Zone .... ATZ
Approximate Dynamic Programming .... ADP
Dynamic Programming .... DP
Efficiency Rating .... EFF
Expected Possession Value .... EPV
Heart Rate .... HR
High Intensity Zone .... HIZ
Integration Definition for Function Modeling IDEF0
Intermittent Recovery .... IR
Markov Decision Process .... MDP
National Basketball Association .... NBA
National Collegiate Athletic Association .... NCAA
Operations Research .... OR
Transition Probability Matrix .... TPM
Volume of Oxygen .... VO_2
Yo-Yo Intermittent Recovery Level 1 .... Yo-Yo IR1
Yo-Yo Intermittent Recovery Level 1 .... Yo-Yo IR2

ABSTRACT

AN APPROXIMATE DYNAMIC PROGRAMMING APPROACH TO DETERMINE THE OPTIMAL SUBSTITUTION STRATEGY FOR BASKETBALL

David W. Hughes, Ph.D.

George Mason University, 2017

Dissertation Director: Dr. John Shortle

The purpose of this research is to determine a coach's optimal basketball lineup throughout the course of a game while accounting for players' changing endurance levels. Coaches currently use a variety of substitution strategies in basketball. These include player rotations, heuristics, playing the "next best" player, and others. However, such strategies may not be optimal in terms of distributing rest over the course of a whole game. In part, this is because live, in-game monitoring of physiological factors, such as heart rates, is not currently authorized by the NCAA. Even with such data, it may not be obvious when to rest a skilled player in order to ensure that player is fresh for later portions of the game. This research develops an approximate dynamic programming (ADP) model of a coach's basketball lineup decisions accounting for changing players' endurance levels and the uncertainty of the defensive intensity played against them. The ADP model is quantified via offline endurance data as well as a value hierarchy of player skill sets created from discussions with the head coach of a Division I program. Policy implications resulting from the model are analyzed and compared to other solution methods. The analysis shows that lineups generated from the ADP model consistently provide more overall value when compared with other substitution strategies, such as the "next best" (or greedy) strategy. This difference tends to be greater for games in which low-intensity defense is played. Knowing each players' current endurance level, their value at each endurance level, and their likelihood of transitioning between endurance levels could help a coach determine a better substitution strategy.

CHAPTER ONE – INTRODUCTION

College basketball coaches make countless, sequential decisions under uncertainty throughout a basketball game in order to best position their team for a win. Some of these decisions include determining the starting lineup, choosing a defensive/offensive formation for each possession, deciding when to call a timeout, and making a player substitution. The coach's lineup decision each minute of the game has the potential to greatly affect the outcome of the game. Thus, it is safe to assume that the coach's objective is to have the best lineup in the game for each possession, given the circumstances of the game. For example, if the team is down by 10 with 5 minutes left to play, a coach may decide to put his/her best offensive lineup in. Alternatively, if the team is up by 10 with 5 minutes to go, the coach may decide to put in his/her best defensive lineup. If the game is close, the coach may go with his/her best balanced lineup.

To determine these "best" lineups for that particular coach, the coach needs to know what he/she wants in an offensive and defensive player. These "wants" may vary drastically from coach to coach. From an offensive standpoint, one coach may want players that score points whereas another coach may desire players that take uncontested shots (and not necessarily make them).[1] Defensively, one coach may value players who

[1] While putting value on simply taking an uncontested shot may seem unusual, college coaches actually do just this. Philosophically, the argument is that players who take uncontested shots are more likely, in the long-run, to make them when compared to taking and making contested shots.

create turnovers, while another coach likes players who are in the “right” defensive position. A coach who knows what he/she wants from his/her players can evaluate each player accordingly. Adding in player constraints (i.e. coach requires one center, two forwards, and two guards in the game at all times), the coach can then determine the makeup of their preferred lineup(s). However, there is a huge shortcoming with this approach. It does not specify how long theses lineups can be kept in the game. The endurance/fatigue of players is often overlooked.

Basketball coaches certainly recognize that their players fatigue throughout the course of the game. However, the coach does not know exactly how fatigued a player is at any given moment because live, in-game monitoring of physiological factors (such a heart rates (HRs)) is not authorized by the NCAA. To counteract this, a coach uses substitution strategies, such as heuristics or player rotations, to guide their substitution decision-making. For example, a coach may determine that the starting center has to be substituted out of the game at the 4-minute mark to keep him/her fresh for later portions of the game. This “rule of thumb” may produce sub-optimal results, especially given the uncertainty of the opponent’s defensive intensity. If the opponent is playing soft defense, the starting center may still have “gas in the tank” at the 4-minute mark and the backup center (who is not as skilled) would come in prematurely. Conversely, if the opponent is using a high intensity defense (i.e. full-court pressure) for the first four minutes, the center might be completely fatigued by the 4-minute mark and the backup center would have been more valuable on the court earlier in the game.

College coaches have a lot of experience in making substitution decisions. They know their players' skills sets, tendencies, and fitness levels. A lot of their knowledge of each player comes from observing daily practice. Some coaches even require their players to wear heart rate monitors during practice so that they can determine who is fatigued or if any player is not giving 100% effort on a certain drill. Unfortunately, the players are not allowed to wear heart rate monitors in games because the NCAA has not approved it. But since coaches have an idea of who is in shape and who is not, they do their best to keep the players on the court as fresh as possible. During a game, they can see and hear when a player is fatigued and thus when they are in need of a substitution. They can even go against their heuristics or set rotation if they see that a player is really fatigued. While making a substitution at the moment when a player is completely fatigued may be the best decision for that particular instant in time, it may not be the optimal decision when considering the game in its entirety. Substituting earlier (or even later to allow other players to recover) may produce better results. Using an approximate dynamic programming (ADP) approach that accounts for each player's unique fatigue and recovery rate, I am going to show that a global strategy for substitutions provides more overall lineup value to the coach when compared with other substitution strategies. The hope is that this new global strategy ultimately leads to more wins.

It is important to note that this research does not address substituting players to improve matchups that dynamically change in the course of a game. For example, if the opponent puts in a very tall center, then the coach might want to put his/her tallest center in the game. It also does not account for late game substitution strategies, like substituting a player on the defensive end to protect the better offensive player from picking up his/her fifth foul or inserting good free-throw shooters in the final minutes of the game when the team is winning and opponent is trying to foul in order to lengthen the game. My research is more geared toward recommending a better substitution strategy for the steady-state game, where players do not get injured or are not in foul trouble.

1.1 Research Questions

I would like to improve the way coaches make substitution decisions in basketball. The current methods of using heuristics, player rotations, or myopic decision making appear sub-optimal. Is it possible for a coach to be proactive and know precisely when to substitute a player given their current endurance/fatigue level (even though they cannot use live, in-game monitoring of physiological factors) and the uncertainty of the opponent's defensive intensity level? In other words:

What are the optimal lineups over the course of a game when accounting for each players' fatigue/recovery rates and the uncertainty of the opponent's defensive intensity?

This fundamental research question leads to other questions that this research attempts to answer such as: (1) how does a coach quantify the value of each player or lineup, (2) how does one quantify the risk of leaving a player/lineup in too long, (3) should a coach begin/end the game with their best lineup, (4) how much better is the

ADP strategy than the greedy strategy or simply using player rotations? A depiction of these research questions, as well as a few more, are shown in Figure 1.

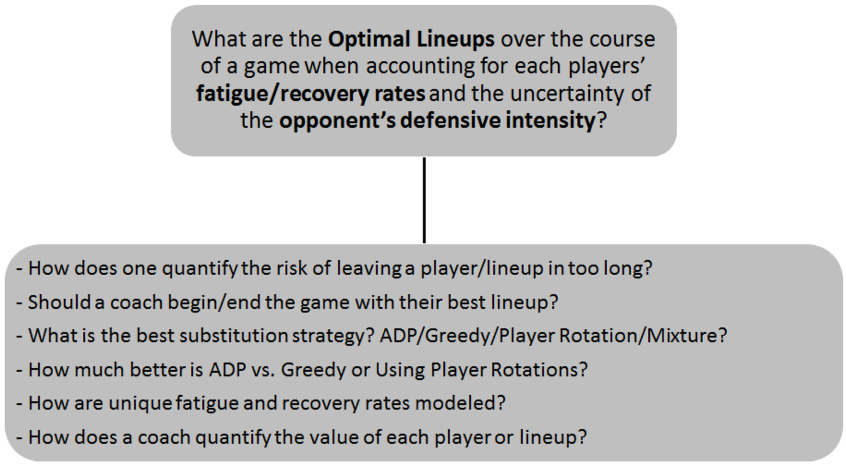

Figure 1: Research Questions

1.2 Dissertation Organization

The remainder of this dissertation is organized as follows. Chapter Two provides the literature review, comprising an overview of basketball metrics currently used, an introduction to current substitution strategies, and a discussion of how endurance/fatigue is measured in sports. Chapter Two also briefly describes three modeling approaches: value modeling, dynamic programming (DP), and approximate dynamic programming (ADP). Chapter Three features the methodology, including purpose and a detailed description of the model. Chapter Three also contains limitations, constraints, and assumptions as well as a section on simulating the model. Model quantification and results are found in Chapters Four and Five, respectively. Chapter Six describes the research contributions, and finally, future work is shared in Chapter Seven.

CHAPTER TWO – LITERATURE REVIEW

This research spans a number of different topics and disciplines. This chapter gives an overview of the literature related to this work. The first section describes the metrics used in basketball that aid in evaluating the effectiveness of a player or lineup. Without metrics, it would be hard to quantify why a certain player (or lineup) should be in the game over another player (or lineup). Next, there is a discussion of current substitution strategies. That is followed by a summary of how endurance/fatigue is measured in sports. Finally, three relevant modeling approaches (Value Modeling, Dynamic Programming (DP), and Approximate Dynamic Programming (ADP)) are discussed.

2.1 Basketball Metrics Overview

Basketball statistics have been recorded and published in daily box scores for decades. However, new statistics are being developed all the time in an attempt to quantify the total effectiveness of a player. For example, beginning with the 1973-1974 NBA season, steals, blocks, and turnovers were first kept as official statistics.[2] Then, in 2002, an Efficiency Rating (EFF) was developed for the NBA. EFF is a measure of how well a player performs on a per game basis.[3] It attempts to quantify the value of a player beyond the three big basketball statistics of points, rebound, and assists (which are

---

[2] http://www.databasebasketball.com/about/aboutstats.htm

[3] EFF = (PTS + REB + AST + STL + BLK − Missed FG − Missed FT − TO) / GP

offensively biased). EFF includes defensive elements (steals and blocks) and also contains metrics that subtract offensive value from a player (number of missed field goals and free throws and number of turnovers). Statistics (or metrics), like EFF, help fans, players, coaches, and statisticians quantify how well a player played or how good a player is over the course of a season or career. But as Kirk Goldsberry pointed out in his “Databall” article: “[t]here is no singular metric that explains basketball any more than there is a singular metric that explains life.” But that does not keep the data analysts and statisticians from being fixated on creating such a singular metric.

Most of the current metrics are tailored for the individual player and are offensively biased. For example, the “triple-double” statistic is widely heralded as the mark of a great player. When a player earns a “triple-double,” that means that the player achieved double digit values in three statistics, which is usually points, assists, and rebounds. Points and assists are purely offensive statistics and rebounds are comprised of both offensive and defensive rebounds. So, it is clear to see how a “standard triple-double” is offensively biased. However, in recent years, there have been unbiased metrics developed that measure the effectiveness of the entire lineup and/or give credit for contributing complementary skills. One example of this is the standard +/- Rating. This metric is based on the premise that a good player helps his/her team. It focuses on how the player's team does when they are on the court versus how the team does when they are off the court. For example, if a player’s +/- Rating is +10, then that means that when that player was on the court, his/her team outscored the opponent by 10 points. It does not matter if that player scored, assisted, or rebounded when they were in the game. It only cares if the team scored more or less points than the opponent when that player was on the court. The strength of this metric is that it is unbiased, and it captures complementary skills that a player brings to the team that are not captured in the box score. However, there is one weakness with this metric. It does not account for the ability of the players they are on the court with and the players they are playing against. If a player is playing against a very easy team, everyone on the team would have a great +/- Rating. Conversely, if a player is playing with the four worst players on the team, his/her +/- Rating would be much worse. To address this weakness, Adjusted +/- Ratings[4] were developed. The Adjusted +/- Rating is just like the +/- Rating, except that it adjusts each player's rating based on the ability of the players they are on the court with and the players they are playing against. Another example, found in Kuehn (2016), "develops a player evaluation framework that stresses the importance of accounting for complementarities between teammates when evaluating players. This is done by developing a probabilistic model of a basketball possession as a progression of events, where the probability of each event's occurrence is determined by the offensive players' skills, the defensive players' skills, and the complementarities between the skills of teammates. Evaluating players using this framework allows one to assess the substitutability between different game actions, the lineup-specific value a player brings to a team, and the players that are the best and worst teammates. It also allows [one] to separately identify the individual effect from the effect teammates have on a player's statistical production." A third example is discussed in Goldsberry (2014). With the

[4] Excel Solver is needed to solve for this metric.

addition of player tracking cameras in the ceilings of every NBA arena, an ocean of new statistical information is available to a limited population. PhD students from the Harvard statistics department have been granted access to this data and developed a metric called expected possession value (EPV), which estimates the predicted point value of any moment during any possession of an NBA game. “They propose that if [one] could build a model that accounts for a few key factors – like the locations of the players, their individual scoring abilities, who possesses the ball, their on-ball tendencies, and their position on the court – one could start to quantify performance value in the NBA in a new way.” All three of these “outside-the-box” techniques of measuring players/lineups has challenged, and even changed, the way coaches think about substitutions.

A summary of current basketball metrics is listed in Table 1 below. The first column of the table lists the metrics. The second column states whether the metric can be used to evaluate individual performance, and the third column specifies if the metric can be used to measure the value of five players that make up a particular lineup. The fourth column indicates if the metric gives credit for complementary skills. The fifth and final column is a description or definition of the metric. To view the formulas for most of these metrics (the ones that are available to the public), please refer to Table 23 in the Appendix A.

Table 1: Player Performance/Effectiveness Metrics

Player Performance/Effectiveness Metric	Individual Metric?	Line-up Metric?	Credit for Complementary Skills?	Description / Definition
Effective Field Goal Percentage (EFG)[1]	Yes	Yes	No	Measures the quality of an individual's OR the team's shooting ability
Turnovers Committed per Possession (TPP)[1]	Yes	Yes	No	Measures how frequently an individual or team gives up control of the basketball before attempting a shot
Offensive Rebounding Percentage (ORP)[1]	No	Yes	No	Measures how often a team rebounds their team's missed shots
Free Throw Rate (FTR)[1]	No	Yes	No	Measures how often a team gets to the foul line and converts at the foul line
Opponent's Effective Field Goal Percentage (OEFG)[1]	Yes	Yes	No	Measures the quality of an individual opponent OR opponent team's shooting ability
Defensive Turnovers Committed per Possession (DTTP)[1]	Yes	Yes	No	Measures how frequently the opponent gives up control of the basketball before attempting a shot
Defensive Rebounding Percentage (DRP)[1]	No	Yes	No	Measures how often a team rebounds their opponent's missed shots
Opponent's Free Throw Rate (OFTR)[1]	No	Yes	No	Measures how often an opponent gets to the foul line and converts at the foul line
NBA Efficiency Rating / Game[1,2]	Yes	Yes	No	Good stats are worth +1 and bad stats are worth -1
Player Efficiency Rating (PER)[1,2,4]	Yes	Yes	No	John Hollinger's all-in-one basketball rating, which attempts to boil down all of a player's contributions into one number. Using a detailed formula, Hollinger developed a system that rates every player's statistical performance
Game Score Rating[1,2]	Yes	Yes	No	Used to rank player performances during a game
Win Scores[1,2]	Yes	Yes	No	Win scores are ultimately converted to a Wins Produced metric
+/- Ratings[2]	Yes	Yes	Yes	Based on the premise that a good player helps his team. It focuses on how the player's team does when he is on the court vs. how he does when he is off the court.
Adjusted +/- Ratings[1,2]	Yes	Yes	Yes	Based on the premise that a good player helps his team. It focuses on how the player's team does when he is on the court vs. how he does when he is off the court. It also adjusts each player's rating based on the ability of the players he is on the court with and the players he plays against.
WINVAL Points Rating[1, 2]	Yes	Yes	Yes	Similar to Adjusted +/- Ratings. Each player's offensive ability and defensive ability are included in the rating system
Expected Possession Value - Added[3]	Yes	Yes	Yes	As a player handles the ball, the EPV is constantly updating. If the player drives to the basket for an uncontested layup, the EPV is basically 2 points. If the player is shooting a contested shot, the EPV is much lower. The simplest method to divvy up credit for the changes in EPV during a play is to attribute them to the person in possession of the basketball at the time of those changes. Using this approach, a player would earn "points added" (or "points subtracted") for a particular sequence, corresponding to the amount of EPV gain (or loss) from the time the player began his play until the time his play terminates (with a pass, shot, or turnover).

1. Winston, Wayne L. Mathletics: How Gamblers, Managers, and Sports Enthusiasts Use Mathematics in Baseball, Basketball, and Football. 2009.
2. NBA Stuffer.com. "Player Evaluation Metrics." Accessed online, 2016. http://www.nbastuffer.com/component/option,com_glossary/func,display/Itemid,90/catid,42/
3. Goldsberry, Kirk. "Databall." Accessed online, 2014. http://grantland.com/features/expected-value-possession-nba-analytics/
4. Wikipedia.org. "Player Efficiency Rating." Accessed online, 2016. https://en.wikipedia.org/wiki/Player_efficiency_rating#Calculation

Any logical combination of these metrics (as long as they are mutually exclusive) could be used with my research methodology to quantify the value of a lineup. Even just one metric (Adjusted +/- Rating, WINVAL Points Rating, or EPV-Added) could be used, if that is what the coach wanted to use to evaluate his/her lineups.

2.2 Current Substitution Strategies

While there is very little formal research currently being done in the area of basketball substitution strategies (Gómez et al. 2016, Clay and Clay 2014, Adkins et al.

2007), there is a plethora of coaches willing to share their substitution philosophies via the internet. Sartini (2009) shares his philosophy that substituting players is a very important part of the game time coaching strategy. He argues that it is important because (1) not having the proper personnel on the floor in crunch time can results in lost games, (2) substituting properly can boost players morale and confidence, (3) substituting keeps key players fresh and out of foul trouble when the coach really needs them, and (4) substituting can change the momentum and pace of the game.

Some coaches use a substitution strategy that attempts to keep the best available lineup in the game at all times (Adkins et al. 2007). That means when a player is fatigued, the coach substitutes the next best player from the bench. So, if a starting guard comes out of the game because he/she is tired, the coach looks to substitute the next best guard off the bench. This approach may be optimal for that next minute of the game. However, this short-sighted approach may not lead to the best solution when considering the entire game.

Other coaches use heuristics to guide their player substitutions. Wilkes (2016) describes an example of a heuristic by stating that the coach needs his/her best player on the floor as much as possible. Wilkes (2016) encourages coaches to play the best player for nine minutes, rest him/her for two minutes, and then put the best player back in for the final nine minutes of the first half. Then, repeat in the second half. Another example of a heuristic is described by Kloppenburg et al. (2016). They advocate players playing until the opponent scores a set amount of points, such as six points. Once the opponent scores six points, then substitute. As mentioned in the introduction, the decision to play the best player for nine straight minutes or until the other team scores six points may be too long, or it may be too short. It all depends on the player's endurance capacity and how quickly that endurance is decayed as a result of the opponent's defensive intensity.

Still other coaches primarily use a player rotation as their substitution strategy. According to Sartini (2009), Clay and Clay (2014), Bourgase (2016), Gels (2016), and Riches (2016), a player rotation is a consistent substitution strategy that allows the players to get comfortable playing at certain times of the game because they know when they are coming in and with whom they are playing. For example, a coach may rotate three relatively equal lineups every four minutes.

One high school coach uses a very unique substitution strategy – self-substitution. Brown (2015) attempts to overwhelm his opponents with "fresh legs" over the course of the entire game. The team's style of play is a full-court press, transition offense, and multiple half-court defensive looks. As with this style of play, they are constantly emphasizing to the players that to be successful, each player has to understand that the maximum time they can likely give on the court (if they are each going all-out), is 5 to 7 minutes of activity at a time. Thus, when a player decides he is too fatigued to continue, he raises his fist. This signal alerts a pre-determined substitute to head to the scorer's table. The self-substitution strategy relies on the player making an honest assessment of his/her current fatigue level.

A few coaches believe in the Grinnell System, which was developed in the early 1990s by coach Dave Arseneault at Grinnell College in Grinnell, Iowa.[5] The Grinnell

[5] https://en.wikipedia.org/wiki/Grinnell_System

System relies on shooting three-point field goals, applying constant pressure with a full-court press, and substituting players frequently. The 15-man roster is divided into three groups of five, and new shifts are substituted every 45 to 90 seconds (similar to an ice hockey shift). Each shift plays at full speed and then rests while the next group does the same. Players rarely play more than 20 minutes a game.

Some coaches do not appear to have any set strategy when it comes to substitutions. These coaches like to change things up constantly and make changes based on the opposition, pace of the game, and gut feel. According to Fagan (2010), one day after his team crumbled in a 104-97 loss to the very bad Washington Wizards, Eddie Jordan's players acknowledged the lack of an actual pattern in the team's substitution pattern. In the loss, a game in which the Philadelphia Sixers held an 18-point, second-half lead, Jordan played his entire active roster. At game's end, only two players, center Marreese Speights (plus 10) and forward Rodney Carney (plus 5), posted positive +/- Ratings. The pair combined for 19 minutes, 49 seconds of playing time – none in the game's fourth quarter. Reserve power forward Elton Brand (even +/- Rating) posted the next-best number. He left the game with 8:49 remaining and did not return. When the head coach was questioned about the rotation, his answer was somewhat defensive and very enlightening at the same time. "For the most part, I don't like tinkering with the lineup….We normally play eight players, and most teams go with an 'eight-player' rotation. Then you pick the guy at nine, 10, 11 that you think can match up for you, that had good practices, that you think can do things against a zone, or against a bigger matchup, so I don't think our rotation is that skewed like you guys think. We go eight players and then we go nine, 10, 11 – just like most other teams. It depends on a gut feeling, how guys practice, and matchups." It is fairly clear that NBA coach Eddie Jordan had no substitution strategy whatsoever.

Attali (2013) argued that players with "hot hands" (meaning the player has scored a basket or two in a row) affected the behavior of players and coaches. According to his research, which is based on an entire NBA season's worth of data, "even a single successful shot sufficed to increase a player's likelihood of taking the next team shot, increased the average distance from which this next shot was taken, decreased the probability that this next shot was successful, and decreased the probability that the coach replaced the player." This showed that coaches either make decisions based on their gut or what they see with their own eyes, or they assume that most recent performance predicts future performance. They do not rely on data or models to help guide their substitution decision-making.

Finally, a few coaches are beginning to look to analytics to supplement their coaching decisions. Gómez et al. (2016), with the use of stepwise multiple linear regression models, identifies the "temporal effects of substitutions related to coach-controlled, on-court, and situational variables in elite basketball." Their findings "allow optimizing coaches' plans and team management of on-court and bench players throughout the game." Next, Clay and Clay (2014) conducts a statistical analysis to "examine the impact of size of rotation on team performance and success among 7,154 NCAA Division I Men's Basketball games collected over multiple seasons." While their research does not allow for the comparison of early game versus late game results, they think it is “reasonable to hypothesize that much of the difference in fatigue levels plays out in the fourth quarter of the game.” Thus, having top offensive players with fresh legs, quick hands, and sustained energy levels on the court at the end of a tight game is essential because they are more likely come up with the offensive rebounds, big steals, and other defensive plays when it really counts. As sports science continues to grow and become more widely accepted, other coaches will consider using analytic tools to supplement their substitution decision-making.

Even though analytical substitution strategies have not been extensively used in basketball, Myers (2012) advocated for the use of operations research (OR) techniques for the timing of substitutions in soccer, particularly for fatigued players. While both soccer and basketball entail intermittent exercise with bouts of short, intense activity, the main difference between the two sports is that soccer only allows three substitutions per game (and a player cannot reenter once they exit) whereas basketball allows for infinite substitutions during the course of a game. Myers proposed a “data mining technique of decision trees to develop a decision rule to guide managers on when to make each of their three substitutions in a soccer match. While past research investigated the overall effect of various strategy changes during the course of a match, no studies focused solely on the crucial element of timing.” Myers’ proposed decision rule demonstrated “between 38% and 47% effectiveness when followed versus 17% and 24% effectiveness when not followed based on over 1200 observations collected from some of the world’s top professional leagues and competitions.”

Please see Table 2 for a summary of current substitution strategies.

Table 2: Summary of Current Substitution Strategies

	Title	Coach	Formal Research	Basketball Related	Heuristic	Player Rotation	Gut	Self-Sub	Greedy	Other	Quantitative Technique
1	Wilkes, Glenn (2016) “Mastering the Art of Substitutions”	x		x	x						
2	Riches, Coach (2016) “Four Ways to think about Substitutions Strategically”	x		x	x	x					
3	Kloppenburg et al. (2016) “Making Sound Substitutions”			x	x	x		x			
4	Sartini, Ken (2009) “Do You Sub to Keep your Players Fresh and Out of Foul Trouble?”	x		x		x					
5	Bourgase, Brock (2016) “Substitution Strategies”	x		x		x					
6	Gels, James (2015) “Basketball Coaching…Game Strategy”			x		x	x				
7	Fagan, Kate (2010) “Jordan’s Substitution Pattern Comes into Question”			x			x				
8	Brown, Bruce (2015) “A Self-Substitution System for Basketball”	x		x				x			
9	Arseneault, Dave (1991) "The Grinnell System"	x		x						x	
10	Adkins et al. (2007) “Basketball Drills, Plays, and Strategies”		x	x					x		Greedy Optimization
11	Attali, Yigal (2013) “Perceived Hotness Affects Behavior of Basketball Players and Coaches”		x	x			x				Regression
12	Gómez et al. (2016) “Exploring the effects of substituting basketball players in high-level teams”		x	x						x	Regression
13	Clay & Clay (2014) “Player Rotation, On-Court Performance and Game Outcomes in NCAA Men's Basketball.”		x	x		x					Regression
14	Myers, Brett R (2012) “A Proposed Decision Rule for the Timing of Soccer Substitutions”		x							x	Decision Trees

With so little formal research in this area, particularly in basketball, it is time to find out if there is a substitution strategy that is better than heuristics, player rotations, self-substitutions, gut feelings, or the myopic approach. Certainly, a coach would love to have a little more clarity with regards to when to make substitutions. But that decision of when to make the substitution hinges largely on the endurance capacity of the player and how quickly the player fatigues/recovers.

2.3 Measuring Endurance/Fatigue in Sports

Fatigue is well researched in the sport of soccer. That is partially due to the fact that player tracking data is approved for use during games. This tracking data not only records player movements, but it also collects physiological data on the players. Another reason that endurance/fatigue of a soccer player is of interest to researchers is due to the fact that substitutions are limited and final. Only three substitutions can be made by each team during a game and once a player is taken out of the game, that player cannot re-enter. Therefore, it is essential that a soccer player has a lot of endurance.

Reilly (1997) pointed out that "soccer entails intermittent exercise with bouts of short, intense activity punctuating longer periods of low-level, moderate-intensity exercise." (I would argue that the same could be said of basketball.) He found that anaerobic efforts are evident when soccer players are actively handling the ball or shadowing fast-moving opponents. However, the largest strain was placed on the aerobic metabolism. His research found that, "on average, competitive soccer corresponded to an energy expenditure of about 75% maximal aerobic power," and that the energy expenditure varied with playing position (being highest among midfield players) and different modes of motion (such as running backwards and sideways). He discussed different ways to calculate energy expenditure, but pointed out that the most widely used strategy was to measure HR during match-play and plot the observations on HR-VO_2 regression lines determined during incremental running on a treadmill. He concluded that "HR in itself provided a useful index of overall physiological strain." He wrapped up by stating that strategies in the areas of training, nutrition, and soccer tactics may be used to reduce the effects of fatigue.

Bangsbo, Marcello-Iaia, and Krustrup (2008) discussed the two Yo-Yo intermittent recovery (IR) tests that were used to evaluate an individual's ability to perform intense exercise over and over again. "The Yo-Yo IR level 1 (Yo-Yo IR1) test focused on the capacity to carry out intermittent exercise leading to a maximal activation of the aerobic system, whereas Yo-Yo IR level 2 (Yo-Yo IR2) determined an individual's ability to recover from repeated exercise with a high contribution from the anaerobic system. Evaluations of elite athletes in various sports involving intermittent exercise showed that the higher the level of competition the better an athlete performed in the Yo-Yo IR tests….The Yo-Yo IR tests showed to be a more sensitive measure of changes in performance than maximum oxygen uptake. The Yo-Yo IR tests provided a simple and valid way to obtain important information of an individual's capacity to perform repeated intense exercise and to examine changes in performance."

Mohr, Krustrup, and Bangsbo (2003) explained that the goal of their study was to "assess physical fitness, game performance, and development of fatigue" during professional soccer matches. "Computerized time-motion analyses were performed 2-7 times during the competitive season on 18 top-class and 24 moderate professional soccer players. In addition, the players performed the Yo-Yo intermittent recovery test." Results showed that top-class players sustained high intensity running and sprinting longer than moderate players, and also performed better on the Yo-Yo intermittent recovery test. The 13 substitute players "covered 25% more ground during the final 15 min of high-intensity running than the other players." Additionally, "defenders covered a shorter distance in high-intensity running than players in other playing positions." Finally, fatigue occurred after periods of high-intensity running as well as right after halftime and at the end of games (Mohr, Krustrup, and Bangsbo 2005).

Achten and Jeukendrup (2003) discussed applications of HR monitors and their limitations. They concluded that “the most important application of HR monitoring is to evaluate the intensity of the exercise performed.” They found that “HR showed an almost linear relationship with VO_2 at submaximal intensities and can therefore be used to accurately estimate the exercise intensity.” They also noted that the relationship between HR and VO_2 varies by individual. Therefore, to get a precise estimate of exercise intensity, the relationship should be determined for each individual. Finally, they pointed out that when HR monitoring is compared with other indications of exercise intensity, monitoring HR was easy, relatively cheap, and can be used in most situations.

A few college basketball programs are beginning to collect and analyze physiological data. For example, in 2015, a university located in Virginia started tracking live HRs of its players during practice in order to monitor the impact of workouts on the individual and assess their overall fitness trends. It is well documented in the literature that HRs can be used to assess an individual’s fitness level as well as determine when a person fatigues. Unfortunately, HR monitors have not been approved by the NCAA for use in an actual game. Despite this challenge, I plan to use the HR data from practices and link it to potential endurance level changes of an individual player in a game.

2.4 Modeling Approaches

There are three relevant modeling approaches that are described in this section. First, there is a value modeling overview, which is developed from Kirkwood (1997) and

Parnell (2010). That is followed by a discussion of dynamic programming (DP) and approximate dynamic programming (ADP). The DP and ADP discussion is mostly adapted from Powell (2011), Balakrishna (2009), and Gosavi (2003). Other useful references include Denardo (1982) and Bellman (2003).

2.4.1 Value Modeling Overview

Value modeling provides a way for a stakeholder or decision-maker to think about what is important to them (what they value), so that their most important functions and objectives are represented in the model. First, a qualitative value model is developed, and then that model is used to build the quantitative value model. The quantitative value model provides a measurable way to evaluate how well different solution options meet the overall objectives of the stakeholder/decision-maker (Parnell 2010).

2.4.1.1 Qualitative Value Modeling

According to Parnell (2010), a qualitative value model consists of a fundamental objective, functions (if used), objectives, and value measures. Often times, this model is represented pictorially as a value hierarchy or value tree. Kirkwood (1997) describes the criteria for a good value hierarchy as "completeness, nonredundancy, decomposability, operability, and small size." The criteria for completeness and nonredundancy mean that the model is collectively exhaustive and mutually exclusive (there is independence between all value measures and all value measures are sufficient in scope to evaluate the fundamental objective). Operability means that it is easily understood by its users, and small size refers to the fact that it should contain as few measures as possible while still being collectively exhaustive and mutually exclusive. An example structure of a value hierarchy is found below in Figure 2.

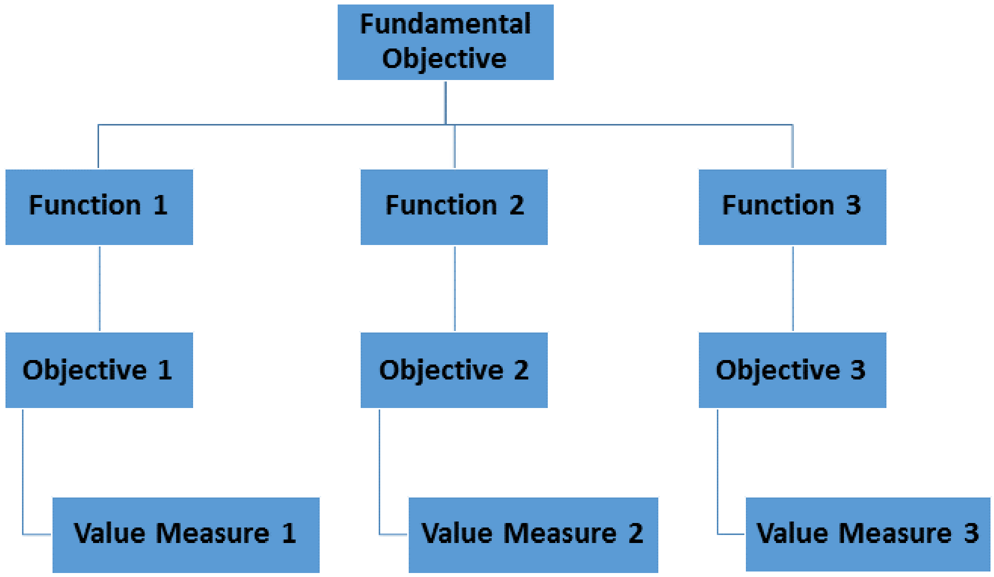

Figure 2: Example Structure of a Value Hierarchy

2.4.1.2 Quantitative Value Modeling

After the value hierarchy is approved by the key stakeholders/decision makers, a quantitative value model is developed. This model allows one to determine how well potential solutions line up with stakeholder values. The mathematical model used to assess the value of potential solutions is called the additive value model and is given by

Equation 1: Additive Value Model Equation

$$v(x) = \sum_{i=1}^{n} w_i v_i(x_i)$$

where $v(x)$ is the total value of a potential solution, i = 1 to n for the number of value measures, x_i is the score of the potential solution on the ith value measure, $v_i(x_i)$ is the single-dimensional value of the potential solution on the ith value measure, w_i is the measure weight of the ith value measure and $\sum_{i=1}^{n} w_i = 1$ (Parnell 2010).

The emphasis for the quantitative value model is at the bottom of the value hierarchy. However, it is not an easy task to develop value measures that are quantifiable and measurable. And once they are determined, value functions have to be developed as a means for evaluating how well potential solutions meet these values. It is also challenging to accurately portray a decision maker's thoughts and beliefs about individual weights of each value measure. These challenges can be overcome by an experienced systems engineer or analyst, and the resulting quantitative value model provides a nice framework for evaluating potential solutions (Parnell 2010).

2.4.2 Dynamic Programming Overview

Dynamic programming (DP) focuses on finding optimal actions to take, given a discrete system state, such that the potential resulting system states are favorable. Dynamic programs attempt to solve for an optimal action (or decision) over a finite or an infinite time horizon. In other words, it tries to solve for the optimal action/decision regardless of the current state. DP models have a Markovian structure with regard to state transition. Recursive equations are used to arrive at an optimal decision for each system state (Birge and Louveaux 2011). The size of the state space must be appropriately constrained for DP, and it also requires state transition probabilities, which can be a challenge to obtain.

DP is used to solve short-term, sequential decisions under uncertainty. DP solves these problems using a recursive formulation (Balakrishna 2009), which is depicted by Bellman's equation (Powell 2011) —

Equation 2: Bellman's Equation

$$V_t(S_t) = \max_{x_t \in X_t}(C_t(S_t, x_t) + \gamma E\{V_{t+1}(S_{t+1})|S_t\})$$

where $V_t(S_t)$ is the total value of a particular state at iteration t, $C_t(S_t, x_t)$ is the immediate contribution value of a particular state S_t and action/decision x_t pair, γ is a discount factor and is set at a constant value of 0.9, and $E\{V_{t+1}(S_{t+1})|S_t\}$ is the accumulated future total value of state S_{t+1} given the current state S_t.

Dynamic programs are comprised of state and decision variables, exogenous information processes, a state transition function, and an objective function. Actions/decisions are made based on state variables. The state transition function, which may be deterministic or stochastic, describes how the state variables change given an action/decision and any pertinent exogenous information processes. Finally, given an objection function and information about a particular state, a good decision can be made.

2.4.2.1 Markov Decision Processes

DP models with discrete states and decisions belong to a class of models known as Markov Decision Processes (MDP). DP is one technique used to solve Markov Decision Process problems. There are three properties that describe Markov processes and MDP. Those three properties are: (1) the jump property – transitions between states commonly occur; (2) the memoryless property – the transition from any state hinges on only the current state and does not depend on the states visited prior to the current state; and (3) the unit time property – all state transitions occur after unit time (Gosavi 2003).

2.4.2.2 The Curses of Dimensionality and Modeling

Bellman's equation provides a compact, recursive solution technique that is used to solve short-term, sequential decision-making problems under uncertainty. But in order to use Bellman's equation, transition probability and contribution/reward matrices must be defined and developed. For the transition probability matrices (TPMs), the probability of transitioning from each and every state to each and every other state must be specified for all possible decisions. For the transition contribution/reward matrix, Gosavi (2003) explains a contribution/reward must be specified for each state and decision pair.

For either matrix, specifying the data needed to solve a dynamic program using Bellman's equation is often very difficult, if not impossible, if the state space is extremely large. Additionally, the computational burden of solving a dynamic program is not only related to the size of the state space, but also to the size of the outcome space and decision space (Powell 2011). Even super computers with large amounts of memory would have a tough time storing and recursing through the matrices necessary to solve a DP problem with hundreds of thousands of states, outcomes, and decisions. This challenge is known as the curse of dimensionality (Powell (2011) and Gosavi (2003)).

It is extremely difficult to develop the relevant dynamics of some systems. In complex DP problems, specifying the state transitions, transition contribution/reward matrices, and transition times, given various actions (or decisions) can be an arduous task. Thus, Gosavi (2003) argues that this curse of modeling can also be challenging.

2.4.3 Approximate Dynamic Programming Overview

Approximate dynamic programming (ADP) techniques can deal with large state spaces and is most useful where the time between decisions is on the order of seconds or minutes. As was discussed in the section on dynamic programming, in order to find an exact solution using Bellman's equation, a complete model (including transition probability and contribution/reward matrices) is required. Even with a complete model, the computational burden of *synchronously* solving for an exact solution may be impossible. Therefore, according to Powell (2011), this curse of dimensionality must be addressed using *asynchronous* approximate dynamic programming, which is a method of updating one state at a time. After updating that one state, the entire value function is updated. Then, in the next iteration, a new state is randomly chosen. The new state is updated, and then the entire value function is again updated. This asynchronous method is continued until all states are explored and have converged on a solution. The computational burden is greatly reduced using this method, thus allowing the problem to be solved much quicker. See Figure 3 below for the steps of asynchronous ADP.

Step 0. Initialize an approximation for the value function $\overline{V}^0(S)$ for all states S. Let $n = 1$.
Step 1. Randomly choose a state s^n.
Step 2. Solve

$$\hat{v}^n = \max_{a \in \mathcal{A}} \left(C(s^n, a) + \gamma \sum_{s' \in \mathcal{S}} \mathbb{P}(s' | s^n, a) \overline{V}^{n-1}(s') \right).$$

Step 3. Use $\hat{v}^n$ to update the approximation $\overline{V}^{n-1}(s)$ for all s.
Step 4. Let $n = n+1$. If $n < N$, go to step 1.

Figure 3: Asynchronous Approximate Dynamic Programming Steps

Due to the fact that a basketball substitution model has the potential to have an extremely large state space and that substitution decisions need to be made frequently (at least once per minute), approximate dynamic programming using asynchronous update may be an efficient approach for this problem. Of course, this would hinge on whether transition probability and contribution/reward matrices can be determined and modeled.

CHAPTER THREE – METHODOLOGY

This chapter describes the proposed research in detail. It begins with defining the purpose of this dissertation, and then describes the model in-depth. Next, it discusses the limitations, constraints, and assumptions of the model. Finally, it concludes with a simple example that demonstrates how all the pieces of the model fit together so that the model can be simulated.

3.1 Purpose

The purpose of this research is to determine a coach's optimal basketball lineup throughout an *r*-minute game while accounting for player's changing endurance levels that are impacted by the uncertainty of the opponent's defensive intensity level. As mentioned in the introduction, current substitution strategies (i.e. heuristics, player rotations, or myopic decision making) appear sub-optimal. Is it possible for a coach to be proactive and know precisely when to substitute a player given their current endurance/fatigue level (even though they cannot use live, in-game monitoring of physiological factors) and the uncertainty of the opponent's defensive intensity level? An opponent's defensive intensity impacts each player's endurance differently. As endurance levels are decayed (due to the opponent's defensive intensity), it is important for the coach to know how a player's value (or contribution/reward) changes. For example, is the coach's top player still the top player when that player is at a low endurance level? Or is a lower skilled bench player at a higher endurance level better? It is equally important for the coach to know how quickly a player's endurance level recovers when they are pulled out of the game and put on the bench. Knowing each players' current endurance level and their value at each endurance level could help a coach determine a substitution strategy. However, the uncertainty about the level of defensive intensity played against the team leaves the coach wondering if his/her substitution strategy is optimal.

3.2 Model

To answer the research questions, a uniquely tailored value model as well as an approximate dynamic programming (ADP) model is developed. The value model is created from discussions with the head coach and their coaching staff. A value hierarchy with quantifiable metrics helps the coach measure the value of each player and/or lineup. The ADP model then uses the value model results to populate the contribution/reward matrix that is needed to solve Bellman's equation. The ADP model also accounts for changing players' endurance levels based on whether they are playing on the court or recovering on the bench.

3.2.1 General Model Information

A game is divided into r short intervals of time. For example, college games are 40 minutes in length, so $r = 40$ corresponds to one-minute intervals. At the start of each time interval, at most one player may be substituted for another, or no substitutions need to be made. This means the coach is limited to no more than one player swap at a time.

Basketball teams are comprised of m players, of which c players are on the court and b players on are the bench. Therefore, $c + b = m$. The college basketball rule book requires five players on the court at all times ($c = 5$). Therefore, $m \geq c$. According to scholarshipstats.com, a men's college team has, on average, 16 players on the team. (Only 13 are allowed to be on scholarship.[6]) However, most teams only play 7-9 players per game.

A player's current endurance level is represented by l distinct endurance levels. In this research, only three ($l = 3$) endurance levels {High (H), Medium (M), and Low(L)} are utilized.

3.2.2 Value Model

A meeting with the key stakeholders allows for the value hierarchy to be developed. See Figure 4. This qualitative value model consists of the fundamental objective, which is to provide the "best" lineup for each possession. The word "best" depends on the strategy used. The functions are listed below the fundamental objective. From my discussion with the coaching staff, they approve of three main functions: (1) conduct efficient offense, (2) defend the basket, and (3) provide experience. Some of these functions are further refined into sub-functions. For example, the first main function (conduct efficient offense) comprises three sub-functions: (1) take uncontested shots, (2) limit turnovers, and (3) get the offensive rebound. The coaching staff could have developed more sub-functions under this main function (like accrue assists, score points, etc.), but they determined that these three sub-functions were the most important for them to focus on. Under each of the functions (or sub-functions, if used), an objective is listed which either maximizes, minimizes, or optimizes the function (or sub-function, if used).

[6] http://www.scholarshipstats.com/basketball.htm

Finally, value measures are defined under each objective. For example, the value measure for the objective ‘maximize defensive matrix score’ is average defensive matrix score[7]. The value measure allows for each player to be scored in a measurable way, and this facilitates comparisons amongst players using the additive value model equation (see Equation 1) that was described in the Quantitative Value Modeling section. Players with higher average defensive matrix scores rank higher than players with lower scores.

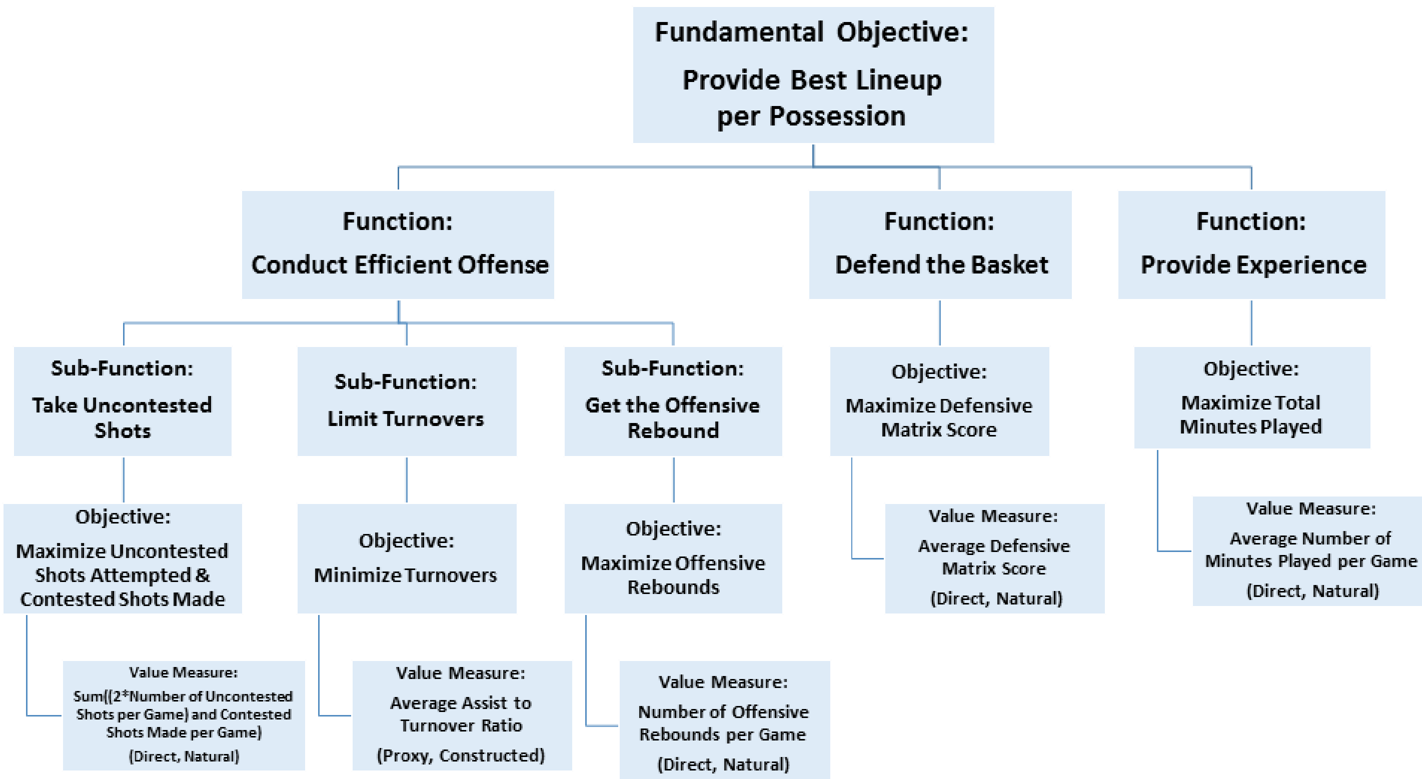

Figure 4: Value Hierarchy

[7] The university that I am working with does not want to reveal the details of their defensive matrix. It is proprietary and will remain a “black box.”

3.2.3 ADP Model

There are five elements to an approximate dynamic program, consisting of state variables, decision variables, exogenous information processes, a transition function, and an objective function. See Figure 5.

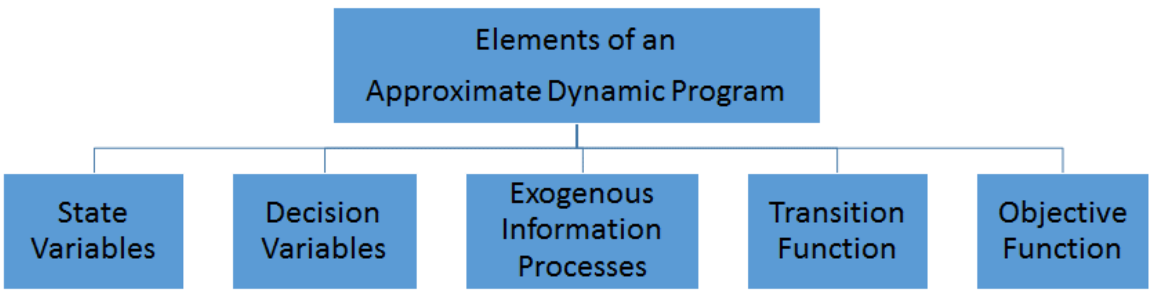

Figure 5: Elements of an Approximate Dynamic Program

According to Powell (2011), state variables describe what one needs to know at a point in time. Decision variables are controlled by the decision maker and represent the primary challenge of approximate dynamic programming. Exogenous information processes are the variables that describe information that arrive exogenously and represent the sources of randomness. The transition function describes how the state changes from one point in time to another. And finally, the objective function allows one to determine which decision (or action) produces the maximum value for a given state. A discussion of each of these elements of the ADP approach and formulation is found in the rest of this section.

3.2.3.1 State Variables

A state is defined as a combination of a player's in/out status (1 if in, 0 if out) and his or her current endurance level. The size of the state space depends on *m*, *c*, and *l*. Assuming eight players on the team (*m* = 8), of which five are on the court (*c* = 5), there are 56 possible lineups (8 choose 5). With three endurance levels (*l* = 3), the state space explodes to 367,416 states ($56 * 3^8$). As an example, two possible states are [1H,1H,1H,1H,1H,0H,0H,0H] and [1H,0L,1M,1H,1L,0H,1H,0M], where H, M, and L denote the current endurance level of a given player. The number of {1's} in each state equal *c*.

3.2.3.2 Decision Variables

To simplify the problem, the coach is limited to no more than one player swap at a time. Given this constraint and parameters *c* and *b*, the number of feasible *d* decisions is calculated using the equation: $d = (c * b) + 1$. If *m* = 8 and *c* = 5, there are 16 possible decisions. (The '+1' decision is the "*do nothing"* decision.) For example, the 16 possible decisions from state [1H,1H,1H,1H,1H,0H,0H,0H] are to do nothing, swap player 1 with 6, 1 with 7, 1 with 8, 2 with 6, and so forth.

3.2.3.3 Exogenous Information Processes

This research focuses on a problem involving stochasticity. Thus, state transitions are affected by quantifiable randomness. Model stochasticity presents as exogenous changes to the state space based on information arriving over the time interval from *t* to *t+1* and is represented by the random variable W_{t+1}. See Figure 6. A sample realization is represented by ω_i where $W_{t+1}(\omega_i)$ is an observation of randomness inherent in the system arriving between time *t* to *t+1*. The sample can be collected through real-world observation of physical processes, sampling known distributions, or applying computer simulation of a complex process.

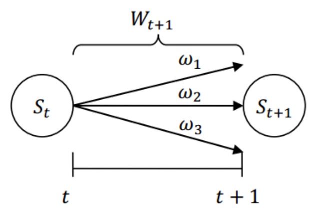

Figure 6: State Transition Model Under Uncertainty

Modeled uncertainty for this research is collected through real-world observation and is limited to the likelihood of decay for the five players on the court that are playing against an uncertain opponent's defense as well as the likelihood of recovery for the players on the bench. The opponent's defensive scheme probabilities are determined from historical data. These defensive scheme probabilities directly impact how each player decays, and therefore also ultimately affect the state transitions.

3.2.3.4 Transition Function

A transition model *M* depicted in Equation 3 incorporates the effects of the new information arriving between time steps and describes the evolution of the state variable. The model for evolving from state S_{t+1} is dependent on the current state S_t, the decision *x* at time *t*, and the random variable W_{t+1}

Equation 3: Transition Model *M*

$$S_{t+1} = S^M(S_t, x_t, W_{t+1})$$

For example, if two possible decisions exist (x_1, x_2) from state S_t then the transition in Figure 6 is modified as follows:

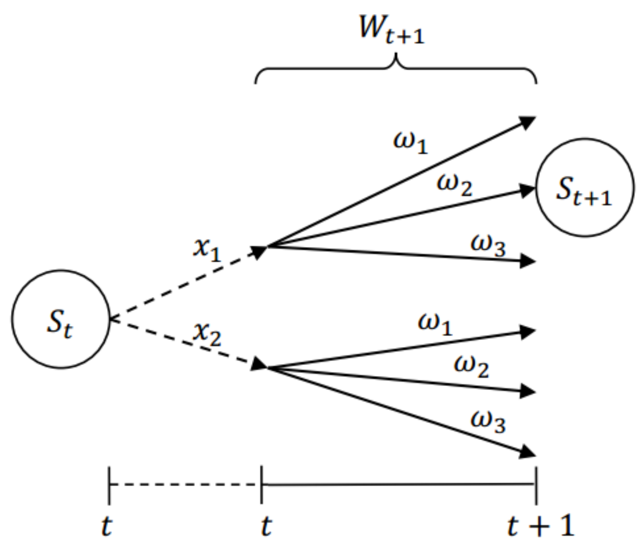

Figure 7: Modified State Transition Model

The model doubles in size from three branches to six branches with the introduction of two different decisions. Note that there is no passing of time associated with the decision, as the decision is made instantaneously at time *t* and prior to observing the random variable W_{t+1}.

The endurance level of a player evolves according to a Markov chain, depending on whether the player is playing on the court or sitting on the bench. (To see how this Markov chain is constructed, refer to Section 4.3.) D_{ij}^k is the probability of decaying from endurance level *i* to endurance level *j* for player k. R_{ij}^k is the probability of recovering from endurance level *i* to endurance level *j* for player k.

Thus, the probability of going from one state to another state is defined as the independent probabilities from the appropriate decay or recovery matrix multiplied by each other. See Figure 8.

Given:

D_{ij}^{k} is the probability of decaying from endurance level *i* to endurance level *j* for player *k*

R_{ij}^{k} is the probability of recovering from endurance level *i* to endurance level *j* for player *k*

$$A_{ij}^{k}(\boldsymbol{S_t}, \boldsymbol{x_t}) = \begin{cases} D_{ij}^{k} & \text{if player } k \text{ is on the court} \\ R_{ij}^{k} & \text{if player } k \text{ is on the bench} \end{cases}$$

Therefore:

$$\Pr\{(\boldsymbol{S_t} \rightarrow \boldsymbol{S_{t+1}})|\boldsymbol{x_t}\} = \prod_{\boldsymbol{k=1}}^{\boldsymbol{m}} \boldsymbol{A}_{\boldsymbol{S_t(k)},\boldsymbol{S_{t+1}(k)}}^{\boldsymbol{k}}(\boldsymbol{S_t}, \boldsymbol{x_t})$$

where:

$\boldsymbol{S_t(k)}$ = endurance state of player *k* in system state $\boldsymbol{S_t}$

$\boldsymbol{x_t}$ = substitution decision at time *t*

Figure 8: Calculating the Probability of Transitioning between States Given a Decision

For example, for a three player team (*m* = 3) with two on the court (*c* = 2) and one on the bench (*b* = 1), the probability of transitioning from the state [1H,1H,0H] to state [1H,1H,0H] (when the decision is to replace no one) is 0.0825 (0.33*0.25*1 = 0.0825). Refer to decay and recovery matrices in Figure 9. The probability of transitioning from the state [1H,1H,0H] to [1H,1H,0L] (when the decision is to replace no one) is zero (0.33*0.25*0 = 0). It is clear to see that infeasible states (moving from a high endurance level to a low endurance level while on the bench) are captured by calculating state transitions in this way.

Player 1 DECAY	H	M	L
H	0.33	0.44	0.22
M	0	0.33	0.67
L	0	0	1

Player 1 RECOVERY	H	M	L
H	1	0	0
M	0.33	0.67	0
L	0.22	0.44	0.33

Player 2 DECAY	H	M	L
H	0.25	0.50	0.25
M	0	0.25	0.75
L	0	0	1

Player 2 RECOVERY	H	M	L
H	1	0	0
M	0.25	0.75	0
L	0.05	0.20	0.75

Player 3 DECAY	H	M	L
H	0.15	0.55	0.30
M	0	0.15	0.85
L	0	0	1

Player 3 RECOVERY	H	M	L
H	1	0	0
M	0.20	0.80	0
L	0	0.15	0.85

Figure 9: Decay/Recovery TPMs for Players 1, 2, & 3

3.2.3.5 Objective Function

According to Powell (2011), I can solve this stochastic optimization problem of the form:

Equation 4: Objective Function

$$\max_{\pi} \mathbb{E}^{\pi}\left\{\sum_{t=0}^{T} \gamma^{t} C_t^{\pi}\big(S_t, X_t^{\pi}(S_t)\big)\right\}$$

by recursively computing Bellman's optimality equations (see Equation 5).

Equation 5: Bellman's Optimality Equations

$$V_t(S_t) = \max_{x_t}(C_t(S_t, x_t) + \gamma\mathbb{E}\{V_{t+1}(S_{t+1})|S_t\})$$

For most problems, solving Equation 4 directly is computationally intractable. However, it does provide the basis for identifying the properties of optimal solutions and finding and comparing "good" solutions. Recall that Bellman's optimality equation has two components, an immediate contribution due to the decision made at time '*t*', and the accumulated future total value of making decision '*x*' in state '*S*' at time '*t*'. Equation 5 provides a mechanism for solving large stochastic optimization problems in a relatively simple way.

In my formulation, the immediate contribution value is dependent on which players are in the game and what their current endurance level is at in the post-decision state. The value of each player comes from the value hierarchy, and the endurance level scales a player's value based on a coach's subjective assessment. For example, if the top player on the team has a value of 100 (out of a possible score of 100) at high endurance, a coach might scale that player by .666 when he/she is at medium endurance and .333 when the player is at low endurance. Thus, the three values for that top player would be 100, 66.6, and 33.3. When m =8, the three values for all eight players help determine the contribution/reward matrix that is necessary to solve the first component of Bellman's equation.

The accumulated future total value is also dependent on which players are in the game and what their current endurance level is at in the post-decision state. The expected future value of landing in all possible post-decision states is different depending on the coach's decision.

Summing both the immediate contribution and the accumulated future total value for each decision allows me to determine which decision produces the maximum value for a given state. Using a gamma (γ) value of 0.9, the maximum values of each state converge after a certain number of iterations. Convergence means that the problem is solved to optimality, and the best decision to make from that particular state is found by seeing which decision gave that converged number. A policy vector π is then produced which prescribes the best decision to make from each and every state. As such, given that m = 8, c = 5, and l = 3, the optimal policy vector π^* specifies one decision/action (from the 16 possible decisions/actions) to make for each of the 367,416 states. It is important to note that in order to get the "greedy" solution, simply change the gamma value from 0.9 to zero (0) and re-solve Bellman's equation.

Figure 10 below shows the Integration Definition for Function Modeling (IDEF0) model for determining the optimal substitution strategy. Using a single box to represent the top level system function, it illustrates the things that cross the system boundary as inputs (on the left) and outputs (on the right), the physical mechanisms that enable the top level system function (on the bottom), and the controls that determine how the function operates (on the top) (Parnell 2010). This figure clearly shows the different inputs that must be defined and determined before solving for the optimal policy vector π^*. All of these inputs are described in the previous sub-sections of the ADP Model section, and are further elaborated on in the Scope section and in Chapter Four.

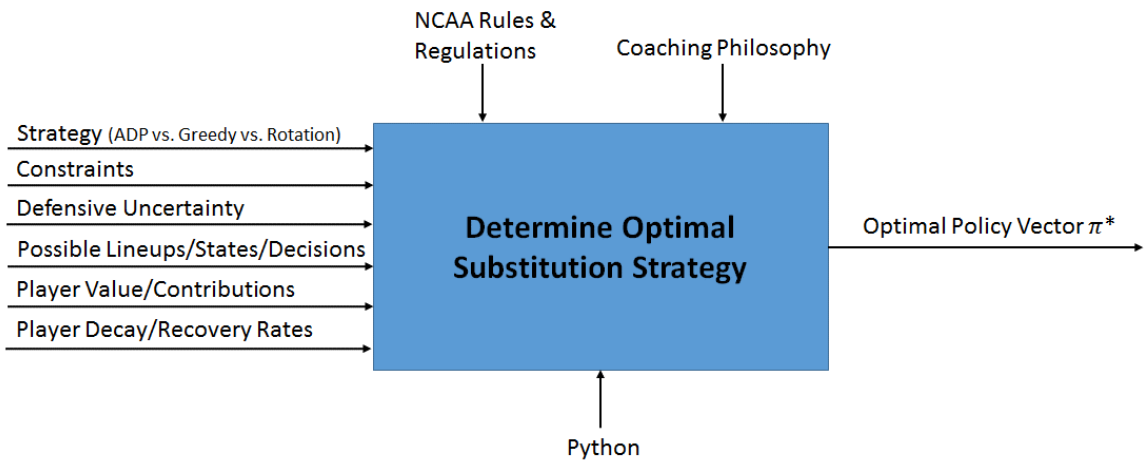

Figure 10: IDEF0 Model for Determining the Optimal Substitution Strategy

3.3 Scope

This section describes the bounds of the research in terms of limitations, constraints, and assumptions. Using Morey (2005) as a guide, limitations are imposed by a lack of capability or resources to consider study objectives more fully. Limitations are (at least in part) beyond the control of the study designer. Constraints are defined as deliberate design choices to define or reduce the design space, and are controlled by the author. Assumptions are unverified statements taken as fact for study purposes in order to proceed efficiently. They often address limitations or constraints.

3.3.1 Limitations

Because heart rate (HR) monitors are not approved for in-game use at the college level (or any level for that matter), it is not possible to know precisely how fatigued (or at what endurance level) a player is at any point in a game. To address this limitation, we have obtained off-line, physiological endurance data for players on a basketball team. This endurance data was collected during practice and included results from four endurance tests that were conducted on a treadmill using a heart rate monitor. The description and discussion of the conversion of this endurance data to individual decay/recovery rates is described in Section 4.3.

It is difficult to get player decay/recovery data on more than one college basketball team because only a few Division I basketball programs currently use HR monitors during practice. My current relationship with the college basketball team in Virginia was fortuitous because I was a friend of one of the assistant coaches on the team, and this particular team actually used HR monitors during practice. Even if I knew of other college basketball programs that used HR monitors, getting access to teams is extremely challenging. Hopefully, as the results of this research become known and as more college basketball teams begin to use HR monitors in practice, the impact of this research will not be limited to such a small population of college teams.

Modeled uncertainty for this research is limited to the opponent's defensive intensity. Although other sources of variability exist that impact endurance levels, such as player injuries, foul trouble, or timeout decisions of the coach, addressing the uncertainty of the defensive intensity is the primary concern.

Many aspects of a player's effectiveness are tied to the human condition. How well a player slept last night (Mah et al. 2011), the quality of the food they eat (Reilly 1997), or whether or not they suffer from jet lag (Song et al. 2016) are just some examples that can influence a player's effectiveness. This research will not consider these aspects.

3.3.2 Constraints

As alluded to in the introduction, this research focuses on the college basketball game. The college game, with few paid analysts available to them (because the NCAA regulates how many paid coaches are on a team's staff), has much to gain from this research. Furthermore, my connection with a Division I basketball program makes it a logical choice for me to constrain the analysis to the college game.

I chose to model a team of eight players (m = 8, c = 5, b = 3) as opposed to 13 or 16. The reason is because the 5-8 players that are not modeled are often red-shirt freshman or players that are simply not good enough to play (yet) in the coach's opinion. Another reason is that the number of possible lineups for a 16-man team is over 4,300 ([m = 16] choose [c = 5] equals 4,368). An eight-person team only has 56 potential lineup combinations (8 choose 5 equals 56), so this helps to significantly constrain the problem, while not losing significant/realistic lineup combinations.

A player's current endurance level can be represented by l distinct endurance levels. In this research, only three (l = 3) endurance levels (High, Medium, and Low) are utilized. Given eight players (m = 8, c = 5, b=3) each with three possible endurance levels (l = 3) means that each one of the 56 possible lineups would have over 6,500 states (l^m = 3^8 = 6,561 states). 6,561 states per lineup means that there are 367,416 total states (56*6,561 = 367,416). When one considers that a 12-man team (with three endurance levels) has over 420 million total states, it seems reasonable to constrain the problem to 8 players (m = 8) and 3 endurance levels (l = 3).

While 56 lineups may appear to constrain the problem appropriately, it actually does not constrain it enough. My analysis of the one Division I basketball program from a university in Virginia (who wishes to remain anonymous) reveals that the coach used 77 lineups (m = 9) in the course of a 14 game conference season. However, 47 of those lineups accounted for less than 15% of the total playing time. Therefore, I further constrained the problem to the top 30 lineups, which accounted for over 85% of the total playing time. See Figure 11 below. This reduced the state space by over 46% (from 367,416 to 196,830 states).

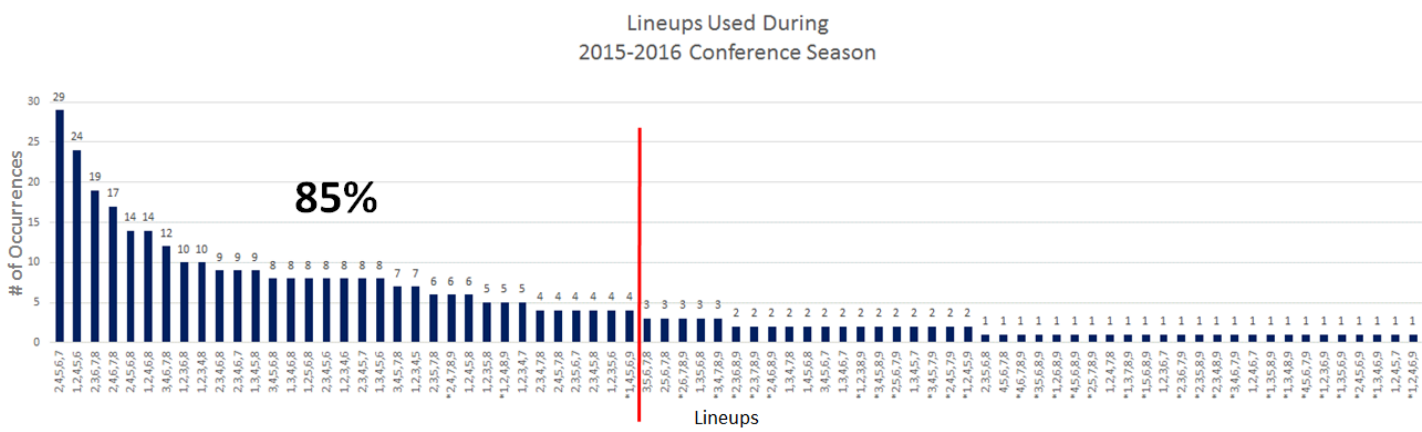

Figure 11: Frequency of Lineups Used

The model does not determine whether a team wins or loses, or even scores for that matter. It simply determines the optimal decision to make from any given state. The thought is that if the value of the players on the court during an r minute game is optimized, then the likelihood of scoring (if on offense), preventing a score (if on defense), and winning the game is higher.

3.3.3 Assumptions

First Beat is a heart rate (HR) monitoring package that tracks and stores HR activity. It automatically estimates a maximum HR based on user inputs (such as age, height, weight, etc.). However, this max HR is updated if a new maximum is reached. There are five training zones that *First Beat* tracks (High Intensity Zone, Anaerobic Threshold Zone, Aerobic Zone 2, Aerobic Zone 1, and Recovery Training Zone). The High Intensity Zone (HIZ) tracks when the HR is between 90-100% of the max HR. The Anaerobic Threshold Zone (ATZ) tracks HRs between 80-90% of the max HR. Aerobic Zone 2 (AZ2) and Aerobic Zone 1 (AZ1) track HRs between 70-80% and 60-70% of the max HR, respectively. There are four common defenses played in college basketball (full-court pressure, half-court pressure, man-to-man, and zone). Each defense has a different level of intensity. I assume that if a team plays against a full-court pressure defense, that is equivalent to being in the HIZ. If they play against the half-court pressure defense, then that maps to the ATZ. And finally, if the team plays against the man-to-man or zone defense, that is the same as being in the AZ2 or AZ1, respectively. Please see Table 3 below.

Table 3: Defenses and their Training Zone Equivalent

Defensive Type	Defensive Name	Training Zone Equivalent
1	Full-Court Pressure	S_1
2	Half-Court Pressure	S_2
3	Man-to-Man	S_3
4	Zone	S_4

where $S_d \in$ {HIZ, ATZ, AZ2, AZ1}

No players get into foul trouble or get injured during the game, requiring them to be substituted.

All player's endurance levels start at 'High' to begin the game. A player's endurance stays at 'High' indefinitely as long as his/her HR is between their resting HR and 60% of their max HR. This means that a player does not begin to fatigue (decay) until their HR is greater than 60% of their max HR. Another way to state this is that a player never fatigues if they are running on a treadmill at 59.99% of the max HR.

A player's decay rate (resulting from playing against an opponent's defense) is linear and constant. Decay is modeled via Markov chains, and the quantification of these Markov chains using player endurance data is described and discussed in Section 4.3.

Once a player is in the 'Low' endurance state, he/she stays in that state as long as they are in the game. Thus, a player that just entered the 'Low' endurance state is valued the same as if that player was in that state for 10 straight minutes.

A player's decay rate is not dependent on the preceding defensive intensity that the player just experienced. This means that no matter what defensive intensity the player just experienced on the court, they decay according decay rate of the current defensive intensity. In other words, it is immaterial if, in t - 1, a player just experienced a high intensity defense (full-court pressure) or a low intensity defense (zone). Their possible decay transition is derived from the current (time t) defensive intensity decay matrix.

A player's recovery rate is also linear and constant. Recovery is modeled via Markov chains, and the quantification of these recovery Markov chains using player endurance data is described and discussed in Section 4.3.

Once a player is in the ‘High’ endurance state, a player stays in that state as long as they remain on the bench. Additionally, a player that just entered the ‘High’ endurance state is valued the same as if that player was in that state for 10 straight minutes.

A player’s recovery rate is not dependent on the preceding defensive intensity that the player just experienced. This means that no matter what defensive intensity preceded the coach’s decision to put them on the bench, they recover at their prescribed recovery rate. In other words, if a player goes to the bench at a low endurance level after a low intensity defense (zone) as opposed to a high intensity defense (full-court pressure), their recovery rate is the same in both instances.

Decay and recovery rates that are derived from the four endurance tests translate to in-game decay and recovery rates.

A player fatigues/recovers at the same rate in minute one as in minute 40, and a player only decays when playing offense, which occurs once every minute.

Normal college halftime lasts 15 minutes during the regular season, so players recover at their designated recovery rate for 15 minutes after the 20-minute mark of the game.

There are no TV timeouts (which occur every 4 minutes in the college game and are 2 minutes in length) and no normal timeouts (4 per game and last 30 seconds each). This means that potential player recovery is not being accounted for.

The coach will specify the lineup constraints before the game. For example, the coach may only want to use 15 specific lineups (out of the 56 possible for a team where $m = 8$) for a particular game. Also, assume that there are no player position constraints.

For example, the coach does not require that two guards, two forwards, and one center have to be on the court at all times. (The specified lineups would mitigate the need to specify positional constraints.)

The probability of going from one state to another state is defined as the independent probabilities from the appropriate decay or recovery matrix multiplied by each other. See Figure 8.

Players' value scores are derived from a coach's value hierarchy. This value score assumes the player is at high endurance. To calculate the value of players when they are at medium and low endurance, the coach develops a scalar for each player. For example, a player that has a total value score of 100 out of 100 (at high endurance) might be scaled by two-thirds (100*(2/3) = 66.66) when they are at medium endurance and by one-third (100*(1/3) = 33.33) when they are at low endurance. Another player that has a total value score of 90 out of 100 (at high endurance) may only be scaled by three-quarters (90*(3/4) = 67.50) when they are at medium endurance and by half (90*(1/2) = 45.00) when they are at low endurance.

A player rotation scheme, when used in the model, is assumed to consist of three lineups (see Table 4) and is described as follows. The first lineup in the rotation is the best five players (as determined from the values scores that are derived from the coach's value hierarchy) from the team of eight, and that lineup plays the first four minutes of the game. The second lineup in the rotation keeps the best two players in the game, but substitutes the remaining three players on the court with the three worst players on the bench. This second lineup plays for two minutes (minutes 5 and 6 of the first half). Then, the third lineup enters the game for minutes 7 and 8. The third lineup is comprised of the third and fourth best players along with the three worst players. Repeat this rotation until halftime (20 minute mark). Then, after halftime, start with the first lineup again. The first lineup ends up getting 12 minutes of playing time per half, while the second and third lineup get four minutes each per half. However, the top four players receive 16 minutes of playing time per half, the fifth best player gets 12 minutes per half, and the three worst players play eight minutes each per half. See Table 5 for a rollup of total playing time for each player. It is important to note that a player rotation does not care about the defensive uncertainty. The idea is to rotate lineups (that play well together) in a predictable way to keep the players fresh and engaged.

Table 4: Player Rotation

Lineup 1 (4 Min)	Lineup 2 (2 Min)	Lineup 3 (2 Min)
Player 7 (Rank #1)	P7 (R1)	P6 (R3)
P2 (R2)	P2 (R2)	P8 (R4)
P6 (R3)	P1 (R6)	P1 (R6)
P8 (R4)	P4 (R7)	P4 (R7)
P3 (R5)	P5 (R8)	P5 (R8)

Table 5: Player Rotation Playing Times per Game

Player #	Minutes Played/Game
7	32
2	32
6	32
8	32
3	24
1	16
4	16
5	16

3.4 Simulating the Model

Using an Integration Definition for Function Modeling (IDEF0) framework, I model how state transitions are simulated. Refer to Figure 12 to see the IDEF0 model for simulating state transitions.

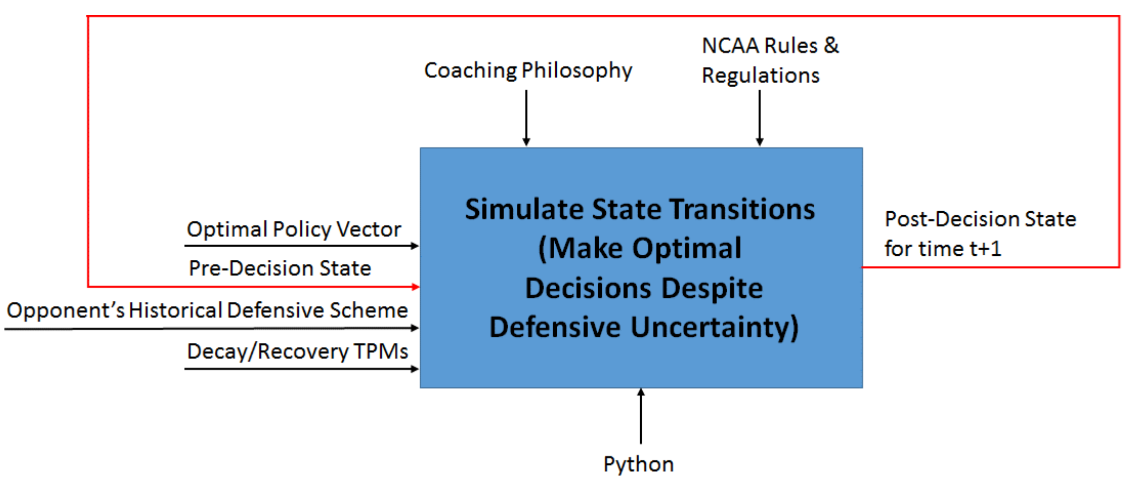

Figure 12: IDEF0 Model for Simulating State Transitions

With the optimal policy vector π^*, it is now possible to simulate state transitions. When uncertainty plays out, that impacts the endurance of the players currently in the game. As mentioned in Table 3, which is found in the Assumptions section, each defense is paired with a different level of intensity. I assume that if the player plays against a full-court pressure defense, that is equivalent to being in the HIZ. If the player plays against the half-court pressure defense, then that maps to the ATZ. And finally, if the player plays against the man-to-man or zone defense, that is the same as being in the AZ2 or AZ1, respectively. Therefore, the opponent's likelihood of playing certain defensive schemes directly affect the rates at which the players decay in the simulation. For example, if the opponent uses 50% zone defense and 50% man-to-man defense, then each minute of the game decay the players on the court according to their individual AZ1 or AZ2 decay TPM (which would be derived from endurance test #3 and #4).

Here is a simple example to demonstrate how state transitions are simulated. Assume that $m = 3$, $c = 2$, $b = 1$, and $l = 3$. Also, assume that the defense only plays man-to-man defense. Therefore, each player only has one decay matrix that came from the Aerobic Zone 2 test. (The decay matrices are in the top right corner of Figure 13). Assume we solve for the ADP and greedy optimal policy vectors. Then, if we start in the state [1H,1M,0M], that means that players 1 and 2 are in the game at a high and medium endurance, respectively, and that player 3 is on the bench at a medium endurance. The total value of state [1H,1M,0M] comes from the value matrix and is 16 (10 for player 1 and 6 for player 2). Player 3 does not contribute to the value because he is on the bench. When the uncertainty plays out, player 1 remains at a high endurance, but player 2 decays from a medium endurance to a low endurance. Player 3 remains at a medium endurance on the bench. In Figure 13, it is clear to see that the new state is now [1H,1L,0M] and its value is 13. In approximate dynamic programming, this is often referred to as the pre-decision state. At this point, the optimal policy vector π^* should be used to determine the best decision (or action) moving forward. According to ADP, the best decision is to replace no one and therefore the post-decision state value remains at 13. However, according to the greedy solution (shown in red text below the line), the best decision is to replace player 2 (swap player 2 with player 3). This decision increases the post-decision value to 14. When the uncertainty plays out again, the pre-decision state changes. Then, using the appropriate policy vector to determine the next best decision, the new post-decision state is determined and that value is calculated. From this simple example, it is easy to see that while ADP was minus 1 (13 vs. 14) in the first post-decision state, it was plus 1 (17 vs. 16) in the second post-decision state. This method of comparison between the ADP and greedy approach can be used to see which one produces more value over one (or many) *r*-minute basketball game(s). This same approach is used for an *m*-person team (with *c* players on the court, *b* players on the bench, and *l* endurance levels).

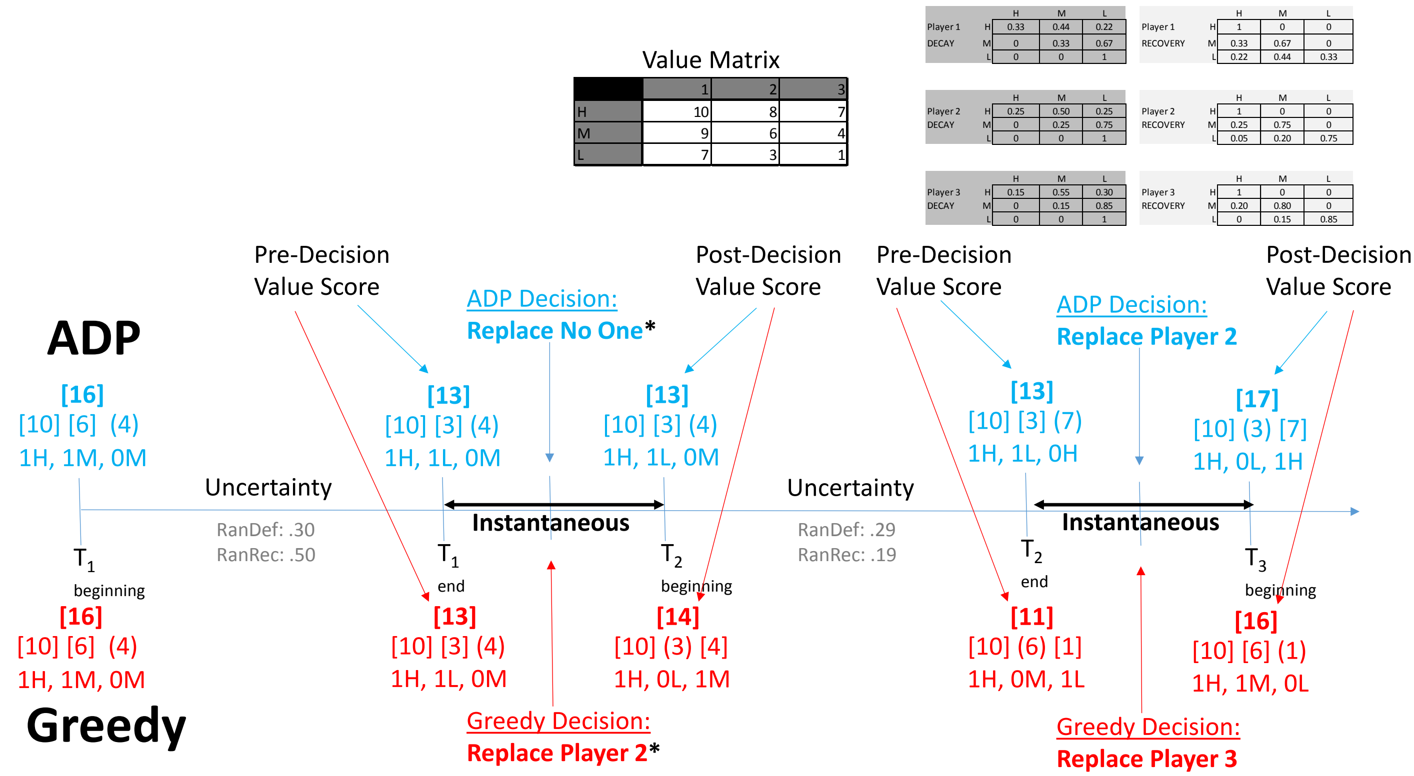

* Different decisions from the same state (1H, 1L, 0M)

Figure 13: Simulating State Transitions (ADP vs. Greedy)

CHAPTER FOUR – MODEL QUANTIFICATION

This chapter includes the quantitative value model results that follow from the value hierarchy, a summary of the four endurance tests, and a description/discussion of the conversion of the endurance data to decay/recovery transition probability matrices (TPMs). This chapter concludes with the endurance test results.

4.1 Quantitative Value Model Results

Given the coach-approved value hierarchy (Figure 4), I gather the raw data I need from all 14 conference games so that I can populate the value measures for each player. This allows me to calculate individual player's value scores relative to one another, using the additive value model that is discussed in the Quantitative Value Modeling section. (The coach requested that I use a balanced weighting scheme [0.4 – Offense, 0.4 – Defense, 0.2 – Experience].) A '1' means that player was the best in that category (or function). A '0' means that player was the worst in the category. Refer to Table 6. The total value for each player fills in the high endurance portion of the value matrix. The coach made a decision to scale each player by two-thirds for the medium endurance and one-third for the low endurance. See Table 7 for the value matrix (which enables the contribution/reward matrix to be calculated).

Table 6: Value Scores and Rank for Each Player

	Weights				
	0.40	0.40	0.20		
Player	Offense	Defense	Experience	Total Value	Rank
1	0.17	0.80	0.36	0.460	6
2	0.37	1	1	0.748	2
3	0.29	0.72	0.84	0.572	5
4	0.34	0	0	0.136	7
5	0	0.02	0.23	0.054	8
6	0.68	0.67	0.98	0.736	3
7	1	0.81	0.80	0.884	1
8	0.52	0.64	0.76	0.616	4

Table 7: Value Matrix for Each Player

		Player							
		1	2	3	4	5	6	7	8
Endurance Level	H	0.460	0.748	0.572	0.136	0.054	0.736	0.884	0.616
	M	0.307	0.499	0.381	0.091	0.036	0.491	0.589	0.411
	L	0.153	0.249	0.191	0.045	0.018	0.245	0.295	0.205

4.2 Endurance Tests Summary

As was briefly mentioned in the Limitations section, we obtained off-line, physiological endurance data for players on a basketball team. This endurance data was collected during practice and included results from four endurance tests that were conducted on a treadmill using a heart rate monitor. The first endurance test (Test #1: High Intensity Zone Test) was designed to see how long a player can run on a treadmill at 14.0 mph (a 4:17 minute/mile pace). The decay timing started once the player's HR was at 90% of their max HR. After the player stopped the test by slowing the treadmill down or jumping off of it, the decay timing stopped and was recorded (see Figure 14 for the IDEF0 model for calculating decay rate for each test) and the recovery timing began. (For the IDEF0 model for calculating recovery rate, see Figure 15.) The recovery time stopped and was recorded once the player's HR was at or below 60% of their max HR. Refer to Figure 16 for a pictorial description of Test #1's timeline. Tests #2, #3, and #4 were very similar, but had slight differences. Test #2 (Anaerobic Threshold Zone Test) used a treadmill speed of 12.6 mph and the timing started when the player's HR was at 80% of their max HR. Test #3 (Aerobic Zone 2 Test) used a treadmill speed of 11.2 mph and the timing started when the player's HR was at 70% of their max HR. Test #4's (Aerobic Zone 1 Test's) treadmill speed was 9.8 mph and the timing started when the player's HR was at 60% of their max HR. See Table 8.

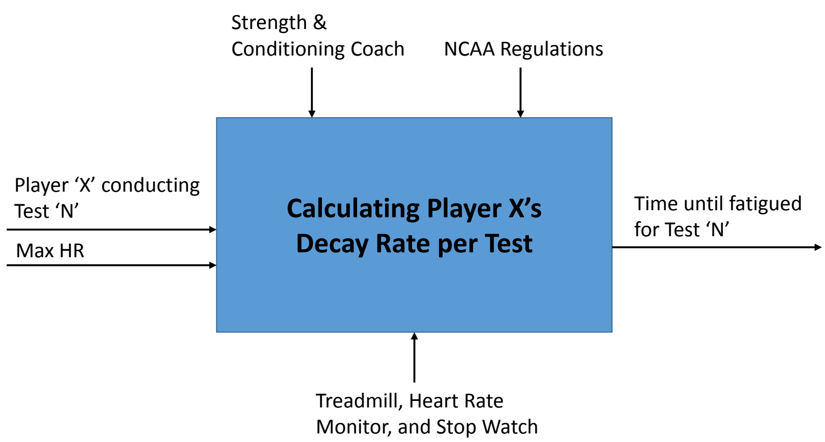

Figure 14: IDEF0 Model for Calculating Player X's Decay Rate per Test

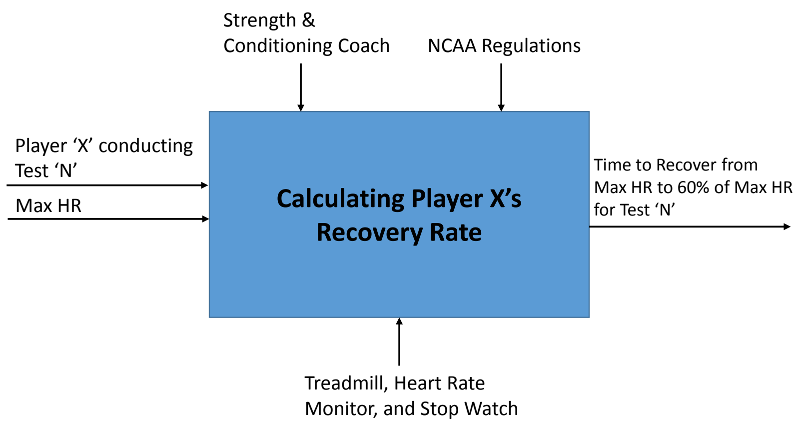

Figure 15: IDEF0 Model for Calculating Player X's Recovery Rate

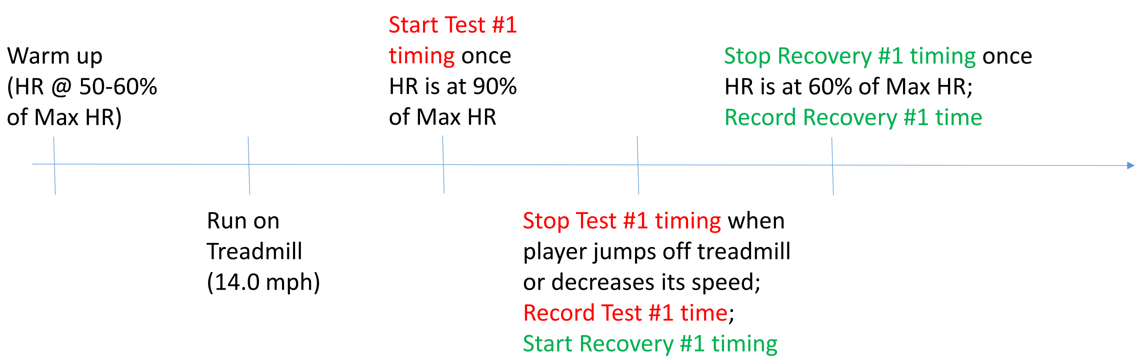

Figure 16: Endurance Test #1 (High Intensity Zone)

Table 8: Endurance Test Parameters

Test #	Zone	Speed (mph)	Timing Begins (% of Max HR)
1	High Intensity Zone	14.0	90%
2	Anaerobic Threshold Zone	12.6	80%
3	Aerobic Zone 2	11.2	70%
4	Aerobic Zone 1	9.8	60%

Figure 17 shows that I tried out the tests to get a sense of what the players had to go through. I tried to accomplish all four endurance tests in one day, but found that it was very difficult to do so. In order to not completely ruin the strength & conditioning coach's planned workout for that day, one endurance test a day is recommended.[8]

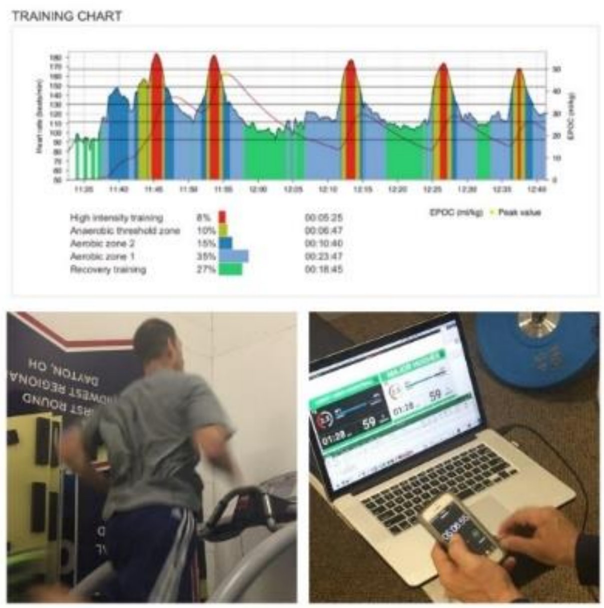

Figure 17: MAJ David Hughes Conducting One of the Endurance Tests

Once all four tests were conducted by Player 'X', four decay and recovery times were recorded. (Notional results are shown in Table 9). In the next section, I will describe a method for converting the decay and recovery times into decay/recovery transition probability matrices (TPMs).

[8] It should be noted that HR data of each basketball player from the university located in Virginia (that wishes to remain anonymous) was already being collected on a daily basis well before I arrived. These four endurance tests were currently being administered by the strength and conditioning coach. I was simply given access to all the HR data of each player on the team, and they let me know the timeframe of the endurance tests so I could extract the appropriate times of each player for all four tests.

Table 9: Example Decay/Recovery Results from the Four Endurance Tests for Player 'X'

	Player 'X' (Minutes)
Test #1 (14.0 mph)	2
Test #2 (12.6 mph)	4
Test #3 (11.2 mph)	6
Test #4 (9.8 mph)	9

	Player 'X' (Minutes)
Test #1 Recovery	5
Test #2 Recovery	5
Test #3 Recovery	4
Test #4 Recovery	4
Average Recovery	4.5

The purpose of these four endurance tests is to measure an individual player's decay and recovery rate. The four endurance tests help to determine the likelihood of transition between levels of endurance. In essence, the decay/recovery tests help serve as a proxy measure for decays and recoveries that occur in a real game. So, as uncertainty plays out on the court, the endurance levels transitions can be simulated.

4.3 Converting Endurance Data to a Markov Chain

This section describes how the endurance data is transformed into decay and recovery transition probability matrices for each player, thus allowing for the calculation of the probability of going from one state to another state as described in Figure 8.

4.3.1 Decay Transition Probability Matrix

Decay (with $l = 3$) is modeled by a Markov chain as shown in the example in Figure 18 (numbers in the matrix are notional).

		H	M	L
Player 'X'	H	0.67	0.33	0.00
DECAY	M	0.00	0.67	0.33
	L	0.00	0.00	1.00

Figure 18: Example Decay TPM for Player 'X'

The values in Figure 18 follow from the assumption that a player's decay rate (resulting from an opponent's defensive intensity) is linear and constant. As previously mentioned, players' decay rates are derived from four endurance tests. I assume that the decay time measured by a particular endurance test is the average time it takes for a player to transition from the highest endurance state to the lowest endurance state while operating in a given HR zone. For example, if a player takes six minutes to step off the treadmill (set at 9.8 mph) in the Aerobic Zone 1 endurance test, then I assume that it takes, on average, six minutes to transition from the 'high' state to the 'low' state. I further assume the intermediate state is "halfway" in between, so it takes three minutes to move from 'high' to 'medium' and three minutes to move from 'medium' to 'low' (refer to top portion of Figure 19).

Now, in the model, fatigue evolves in a probabilistic manner, not a deterministic manner. I choose the parameters of the Markov chain such that the *average* time it takes to go from one state to the next is the same as the *measured* time obtained from the data. This yields the transition probabilities shown in the bottom portion of Figure 19. For example, if a player is in the 'high' state, then the player remains in the 'high' state with probability 2/3 and decays to the 'medium' state with probability 1/3. The number of time steps until the player transitions from ‘high’ to ‘medium’ is a geometric random variable with a mean of 1 / (1/3) = 3 minutes. This assumes 1-minute per time step.

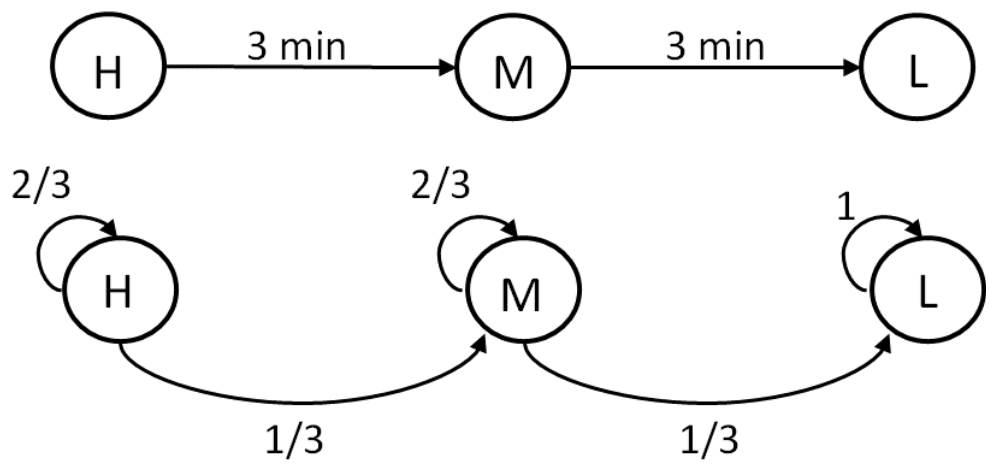

Figure 19: Decay Transition Times between Endurance Levels and Their Associated Transition Probabilities

It is also assumed that once a player is in the ‘Low’ endurance state, he/she stays in that state as long as they are in the game. Thus, a player that just entered the ‘Low’ endurance state is valued the same as if that player was in that state for 10 straight minutes. It is important to note that it is possible to directly transition from ‘High’ to ‘Low’ in one time step if a player lasts less than two minutes during one of the endurance tests.

For each player, an overall decay matrix is created from the four decay TPMs that correspond to the four training zones/defenses (refer to Table 3). From these four decay matrices, a single decay matrix is obtained for each player as long as the opponent’s defensive scheme probabilities (see Table 10) are known. (These defensive scheme probabilities come from historical data.)

Table 10: Opponent's Historical Defensive Scheme Probabilities

Opponent's Defense	% of Time Defense Used
Full-Court Pressure	25%
Half-Court Pressure	25%
Man-to-Man	25%
Zone	25%

Refer to Equation 6 to calculate the probabilities found in the overall decay TPM.

Equation 6: Overall Decay TPM Probability

$$Pr_{ij}^{overall} = \sum_{d=1}^{4} w_{S_d}\left(P_{ij}^{S_d}\right)$$

Simply weight each of the four decay TPMs by the corresponding defensive scheme probability and then sum up the four transition pairs. Please see the example in Figure 20. In this notional example, an opponent never uses full-court and half court pressure, and therefore, the HIZ and ATZ matrix are zeroed out. The .75 for the High to High transition in the Overall Decay TPM comes from weighting the AZ2 and AZ1 matrix by .50 each and then summing the High to High transitions from each of matrix (($P_{HH}^{overall}$ = 0*0) + (0*.50)+(.50*.67) + (.50*.83) = .75).

Player 'X' DECAY H.I.Z.		H	M	L
	H	0.00	1.00	0.00
	M	0.00	0.00	1.00
	L	0.00	0.00	1.00

Player 'X' DECAY A.T.Z.		H	M	L
	H	0.50	0.50	0.00
	M	0.00	0.50	0.50
	L	0.00	0.00	1.00

Player 'X' DECAY A.Z.2.		H	M	L
	H	0.67	0.33	0.00
	M	0.00	0.67	0.33
	L	0.00	0.00	1.00

Player 'X' DECAY A.Z.1.		H	M	L
	H	0.83	0.17	0.00
	M	0.00	0.83	0.17
	L	0.00	0.00	1.00

Training Zones	Defensive Equivalent	% of Time Defense Used
High Intensity Zone	Full-Court Pressure	0%
Anaerobic Threshold Zone	Half-Court Pressure	0%
Aerobic Zone 2	Man-to-Man	50%
Aerobic Zone 1	Zone	50%

Player 'X' DECAY Overall		H	M	L
	H	0.75	0.25	0.00
	M	0.00	0.75	0.25
	L	0.00	0.00	1.00

Figure 20: An Example of How to Calculate Player X's Overall Decay TPM

4.3.2 Recovery Transition Probability Matrix

Similar to decay, recovery (with $l = 3$) is modeled by a Markov chain as shown in the example in Figure 21.

Player 'X' RECOVERY		H	M	L
	H	1.00	0.00	0.00
	M	0.33	0.67	0.00
	L	0.00	0.33	0.67

Figure 21: Example Recovery TPM for Player 'X'

The values in Figure 21 follow from the assumption that a player's recovery rate is also linear and constant. This means that if it takes six minutes (on average) for him/her to recover to 60% of their max HR after stepping off the treadmill, then assume it takes three minutes to move from 'Low' to 'Medium' and another three minutes to move from ‘Medium’ to ‘High’. See top portion of Figure 22. This assumption would lead to the probabilities listed in the bottom portion of Figure 22.

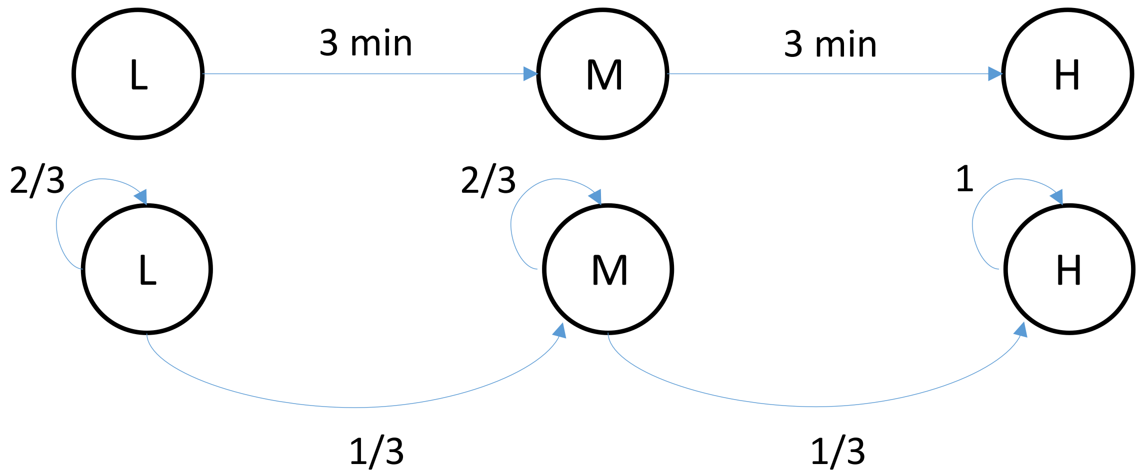

Figure 22: Recovery Transition Times between Endurance Levels and Their Associated Transition Probabilities

Once a player is in the ‘High’ endurance state, a player stays in that state as long as they remain on the bench. Additionally, a player that just entered the ‘High’ endurance state is valued the same as if that player was in that state for 10 straight minutes. It is important to note that it is possible to directly transition from ‘Low’ to ‘High’ in one time step if a player’s HR recovers to 60% of their max HR less than two minutes (on average) during the recovery tests.

4.4 Endurance Test Results

The strength and conditioning coach at the Division I basketball program that I worked with provided me the heart rate data from the four physiological tests that each player conducted. From this data, I extract all four endurance test times (in minutes) for each player. See Table 11.

Table 11: Endurance Test Results (in Minutes) for Each Player

	Player							
	1	2	3	4	5	6	7	8
Test #1 (14.0 mph) - High Intensity Zone	1.22	0.92	2.12	0.97	0.52	1.52	1.77	0.48
Test #2 (12.6 mph) - Anaerobic Threshold Zone	2.65	1.97	4.45	2.18	1.10	3.20	4.23	1.02
Test #3 (11.2 mph) - Aerobic Zone 2	6.62	5.33	5.73	6.45	2.90	5.92	6.27	2.22
Test #4 (9.8 mph) - Aerobic Zone 1	13.73	10.30	11.35	12.35	6.37	9.55	11.75	5.82

	Player							
	1	2	3	4	5	6	7	8
Test #1 Recovery	2.17	3.23	1.98	3.08	3.53	2.03	2.53	4.77
Test #2 Recovery	2.27	3.38	1.72	3.15	3.68	2.32	2.63	4.67
Test #3 Recovery	2.47	3.53	1.62	3.32	3.62	2.23	2.72	4.85
Test #4 Recovery	2.20	3.25	1.65	3.02	3.25	2.15	2.15	4.53
Average Recovery	2.28	3.35	1.74	3.14	3.52	2.18	2.51	4.71

From this endurance test information, I determine each player's decay and recovery TPM in the manner that is described in Section 4.3. This enables the calculation of the probability of going from one state to another state (see Figure 8). I now have all the data I need to solve for the optimal policy vector π^*(see Figure 10).

CHAPTER FIVE – RESULTS

This chapter presents a discussion about the model complexity and then the results from two case studies. The first case study is a toy problem (3 choose 2 (3c2)) that illustrates concepts. The second is a real-world problem (8 choose 5 (8c5)). This chapter concludes with discussion and insights from the results.

5.1 Model Complexity: The Curses of Modeling and Dimensionality

It is extremely difficult to develop the relevant dynamics of some systems. In complex dynamic programming problems (like what is described in this dissertation), specifying the state transitions, transition contribution/reward matrices, and transition times, given various actions/decisions can be an arduous task. That is why Gosavi (2003) argues that this curse of modeling can be a challenging task. In order to use Bellman's equation, transition probability and contribution/reward matrices must be defined. If the state space is extremely large, specifying the data needed to solve a dynamic program using Bellman's equation is difficult, if not impossible. While my problem has a very large state space, figuring out ways to appropriately constrain it (i.e. limiting the number of players on the team to eight and reducing the number of useable lineups to 30) help keep the problem solvable. For the transition probability matrices, the probability of transitioning from each and every state to each and every other state must be specified for all possible decisions. This is one of the most challenging aspects of my research.

Determining each player's probability of transitioning between endurance levels (for both decay and recovery) is a necessary step in order to define a transition probability matrix. I determine this by collecting eight times (four decay times and four recovery times) from four endurance tests. Making the assumption that I can model transitions between endurance levels using a Markov chain is a necessary first step that must be solved before calculating the transition probability matrix. The next crucial step is making the assumption that the probability of going from one state to another state is defined as the independent probabilities from the appropriate decay or recovery matrices multiplied by each other. If not for these two critical assumptions, I cannot use the methodology of asynchronous ADP that is described in Approximate Dynamic Programming Overview section of Chapter Two. For the contribution/reward matrix, a contribution/reward must be specified for each state/decision pair. Using a value hierarchy and additive value model to calculate the contribution/reward is a novel way to populate the contribution/reward matrix. For both matrices, I develop defensible techniques to specify the data needed.

The computational burden of solving a dynamic program is not only related to the size of the state space, but also to the size of the outcome space and decision space (Powell 2011). Super computers with large amounts of memory have a tough time storing and recursing through the matrices necessary to solve an approximate dynamic programming problem with hundreds of thousands of states, outcomes, and decisions. This challenge is known as the curse of dimensionality (Powell 2011 and Gosavi 2003). For example, solving for a policy vector with 196,830 states (30 possible lineups: 8

choose 5 with no lineup restrictions and three endurance levels) is impossible with my personal laptop that has 8 GB of RAM and an eight core processor. Fortunately, I have access to a CPU at George Mason University that has 192 GB of RAM and a 56 core processor. However, I discover that precomputing the 16 action-specific TPMs (which are 196,830 states by 196,830 states each) requires over 38.7 billion calculations per action-specific TPM and at least one terabyte of RAM. Since I only have about 20% of the required RAM on the Mason CPU, I have to compute the action-specific TPMs on the fly. This causes the solving time to exponentially explode. Essentially, states in the TPM are calculated each time they are randomly visited, which may be 50-100 times. So, instead of calculating a row (state) in the action-specific TPM just once upfront and storing that row (state) in memory, I calculate on an as needed basis. Since I do have 192 GB of RAM available to me, I store as many rows (states) as I can in a cache. Thus, if I have 33% of all states in cache, then one out of every three iterations are not re-calculated. Establishing a cache does save computing time, but nowhere near the time savings if all rows of all action-specific TPMs are precomputed and stored in RAM. Since I cannot obtain a CPU with a terabyte of RAM, I decide that this is the only feasible way to solve this size of a problem. Table 12 and Figure 23 show the computing (solving) times of different scenarios.

Table 12: Computing Times (in seconds) for Different Scenarios

m	c	b	l	# Lineups (m choose c)	# Actions/Decisions ((c*b)+1)	# States ((l^m)*(#Lineups))	Cache Required?	Computing Time (in seconds)
3	2	1	3	3	3	81	No	2
4	2	2	3	6	5	486	No	133
5	3	2	3	10	7	2,430	No	1,257
6	3	3	3	20	10	14,580	No	11,772
7	5	2	3	21	11	45,927	No	51,694
8	5	3	3	15*	16	98,415	No	108,780
8	5	3	3	30*	16	196,830	Yes	1,123,249
8	5	3	3	56	16	367,416	Yes	Unknown

* Coach-specified lineup constraint

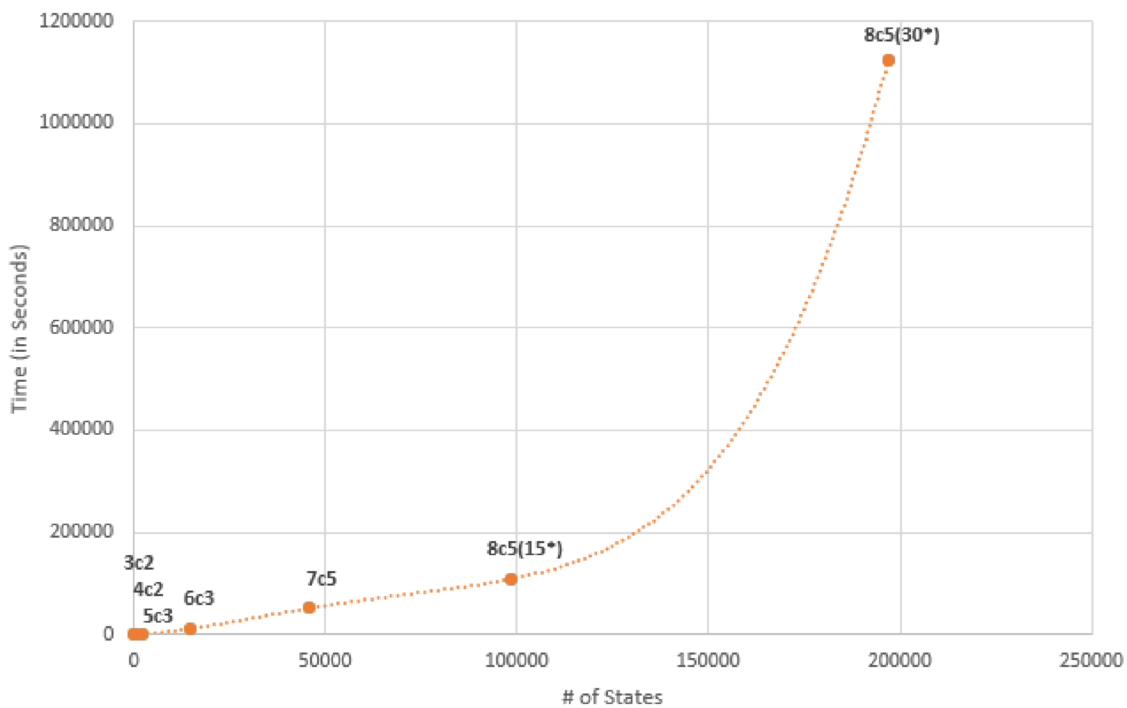

Figure 23: Computing Time for Each Scenario

5.2. Toy Problem (3c2)

To demonstrate how the ADP model works conceptually, I solve a small toy problem first so that one can more easily understand a larger, real-world problem. For the toy problem, I determine the optimal substitution strategy for a 3-person team (m = 3) with two players on the court (c = 2) and one player on the bench (b = 1). That implies that there will be three possible lineups (3 choose 2). Using only three endurance levels (l = 3) means that each one of the three possible lineups has 27 states (3^3= 27 states). 27 states per lineup means that there are 81 total states (3*27 = 81). Since the coach is limited to no more than one player swap at a time, there are only four actions (three of which are feasible from any state) the coach can take: (1) replace no one, (2) replace player 1, (3) replace player 2, and (4) replace player 3.

I am going to follow the same limitations, constraints, and assumptions that were listed in the Scope section.

5.2.1 Toy Problem (3c2 – Variant #1)

First, assume that the endurance tests results for all three players are given in Table 13. From these test results, it is clear that Player 1 decays the slowest and recovers the fastest. Then, Player 2's decay rate is the next slowest, and his recovery rate is the next fastest. Finally, Player 3 has the fastest decay rate and the slowest recovery rate.

Table 13: Endurance Test Decay/Recovery Results for Players 1, 2, & 3 (3c2 – Variant#1)

	Player 1 (Minutes)	Player 2 (Minutes)	Player 3 (Minutes)
Test #1 (14.0 mph)	2	1	1
Test #2 (12.6 mph)	4	3	2
Test #3 (11.2 mph)	6	5	2.86
Test #4 (9.8 mph)	12	8	5

	Player 1 (Minutes)	Player 2 (Minutes)	Player 3 (Minutes)
Test #1 Recovery	7	9	14
Test #2 Recovery	6	8	13
Test #3 Recovery	6	7.5	12
Test #4 Recovery	5	7.5	11
Average Recovery	6	8	12.5

From these endurance test results, decay and recovery matrices (not shown) can be derived in the same way that is described in Section 4.3. Also, assume that the opponent is equally likely to use all four defenses. See Table 14 below. With this information, we can calculate the overall decay matrix (not shown) in the same way that is explained Section 4.3.1.

Table 14: Opponent's Likelihood of Using Each Defense

Opponent's Defense	% of Time Defense Used
Full-Court Pressure	25%
Half-Court Pressure	25%
Man-to-Man	25%
Zone	25%

Assume that we talk to the coaching staff and generate a value hierarchy (not shown) which gives value matrix found in Table 15. From this value matrix, it is clear that Player 1 is the best player, followed by Player 2, and then Player 3. Overall, Player 1 has the best value (skill) and best endurance. On the other hand, Player 3 not only has the lowest value scores, but also has the worst endurance. One would expect him to not play very often.

Table 15: Value Matrix for Players 1, 2, & 3 (3c2 – Variant #1)

	1	2	3
H	10	8	7
M	9	6	4
L	7	3	1

From this value matrix, we determine the contribution matrix, a portion of which is shown below in Table 16. Note that in the last column (Replace Player 3 Action) that there are all zeros (0). This is because it is an infeasible state. Player 3 is not in the lineup, thus Player 3 cannot be replaced.

Table 16: Portion of the Contribution Matrix (27 out of 81 Total States) (3c2 – Variant #1)

	Replace No One	Replace Player 1	Replace Player 2	Replace Player 3
1H,1H,0H	18	15	17	0
1H,1H,0M	18	12	14	0
1H,1H,0L	18	9	11	0
1H,1M,0H	16	13	17	0
1H,1M,0M	16	10	14	0
1H,1M,0L	16	7	11	0
1H,1L,0H	13	10	17	0
1H,1L,0M	13	7	14	0
1H,1L,0L	13	4	11	0
1M,1H,0H	17	15	16	0
1M,1H,0M	17	12	13	0
1M,1H,0L	17	9	10	0
1M,1M,0H	15	13	16	0
1M,1M,0M	15	10	13	0
1M,1M,0L	15	7	10	0
1M,1L,0H	12	10	16	0
1M,1L,0M	12	7	13	0
1M,1L,0L	12	4	10	0
1L,1H,0H	15	15	14	0
1L,1H,0M	15	12	11	0
1L,1H,0L	15	9	8	0
1L,1M,0H	13	13	14	0
1L,1M,0M	13	10	11	0
1L,1M,0L	13	7	8	0
1L,1L,0H	10	10	14	0
1L,1L,0M	10	7	11	0
1L,1L,0L	10	4	8	0

Solving Equation 4 by recursively computing Bellman's optimality equations (see Equation 5), convergence occurs (with a threshold of 0.01) after 89 iterations for synchronous ADP, which means it solves to optimality. Changing the gamma value from 0.9 to 0.0 to get the greedy policy vector causes immediate convergence, as it should. A portion of the two optimal policy vectors is shown below in Table 17. There are 12 differences between the policy vectors. (Four of the differences are highlighted in yellow and shown in Table 17.) This is quite an interesting finding, especially considering how the numbers are stacked against Player 3. One might think that making substitution decisions with the make-up of those three players would be relatively straight forward. However, approximately 15% (12 out of 81 states) of the ADP policy vector is different.

Table 17: Optimal Policy Vector Comparison between ADP and Greedy (3c2 – Variant #1)

State	Optimal Action (ADP)	Optimal Action (Greedy)
[0L,1L,1L]	Swap 1 and 3	Swap 1 and 3
[0L,1L,1M]	Swap 1 and 3	Swap 1 and 2
[0L,1L,1H]	Swap 1 and 2	Swap 1 and 2
[0L,1M,1L]	Swap 1 and 3	Swap 1 and 3
[0L,1M,1M]	Swap 1 and 3	Swap 1 and 3
[0L,1M,1H]	Swap 1 and 2	Swap 1 and 2
[0L,1H,1L]	Swap 1 and 3	Swap 1 and 3
[0L,1H,1M]	Swap 1 and 3	Swap 1 and 3
[0L,1H,1H]	Swap 1 and 3	Do Nothing
[0M,1L,1L]	Swap 1 and 3	Swap 1 and 3
[0M,1L,1M]	Swap 1 and 3	Swap 1 and 2
[0M,1L,1H]	Swap 1 and 2	Swap 1 and 2
[0M,1M,1L]	Swap 1 and 3	Swap 1 and 3
[0M,1M,1M]	Swap 1 and 3	Swap 1 and 3
[0M,1M,1H]	Swap 1 and 2	Swap 1 and 2
[0M,1H,1L]	Swap 1 and 3	Swap 1 and 3
[0M,1H,1M]	Swap 1 and 3	Swap 1 and 3
[0M,1H,1H]	Swap 1 and 3	Swap 1 and 3
[0H,1L,1L]	Swap 1 and 3	Swap 1 and 3
[0H,1L,1M]	Swap 1 and 3	Swap 1 and 2
[0H,1L,1H]	Swap 1 and 2	Swap 1 and 2
[0H,1M,1L]	Swap 1 and 3	Swap 1 and 3
[0H,1M,1M]	Swap 1 and 3	Swap 1 and 3
[0H,1M,1H]	Swap 1 and 2	Swap 1 and 2
[0H,1H,1L]	Swap 1 and 3	Swap 1 and 3
[0H,1H,1M]	Swap 1 and 3	Swap 1 and 3
[0H,1H,1H]	Swap 1 and 3	Swap 1 and 3

Using these two policy vectors, I simulate 25, 40-minute games for each strategy. (Seeded values are used for the defensive uncertainty, so that the comparison is fair.) The boxplots in Figure 24 show the cumulative post-decision state value of each game. It is clear to see that ADP policy vector performs better (when using the median, mode, or average as a comparison metric) than the greedy policy vector. See Figure 25 for descriptive statistics.

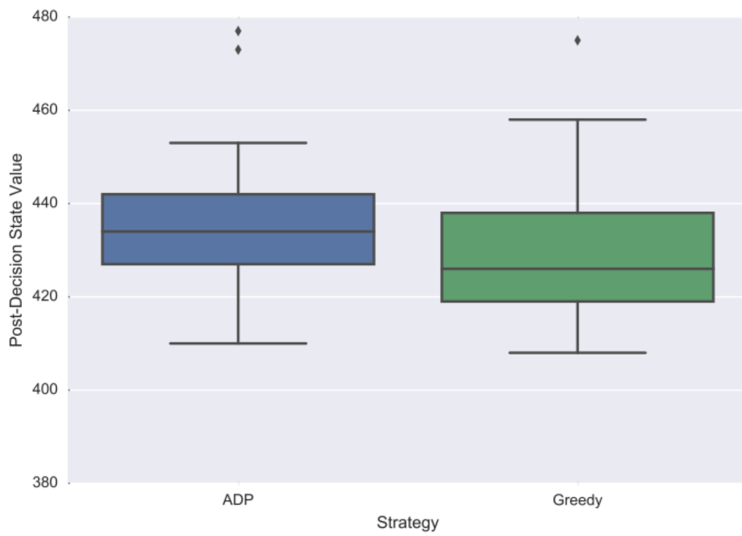

Figure 24: Boxplot Comparison of ADP vs. Greedy (3c2 – Variant#1, n = 25)

Post-Decision Value - ADP	
Mean	435.40
Standard Error	3.27
Median	434.00
Mode	434.00
Standard Deviation	16.34
Sample Variance	267.08
Kurtosis	1.14
Skewness	1.03
Range	67.00
Minimum	410.00
Maximum	477.00
Sum	10885.00
Count	25.00
Confidence Level(95.0%)	6.75

Post-Decision Value - Greedy	
Mean	430.52
Standard Error	3.37
Median	426.00
Mode	423.00
Standard Deviation	16.85
Sample Variance	283.76
Kurtosis	0.53
Skewness	0.95
Range	67.00
Minimum	408.00
Maximum	475.00
Sum	10763.00
Count	25.00
Confidence Level(95.0%)	6.95

Figure 25: Descriptive Statistics for ADP & Greedy Strategy (3c2 – Variant #1, n = 25)

To see how many times the greedy strategy outperformed the ADP strategy, I subtract the cumulative post-decision value of the ADP strategy from the greedy strategy of each seeded game. Thus, if the result is positive, then ADP beat greedy. If the result is negative, then greedy beat ADP. Greedy only beat ADP seven out of 25 times (28%). See Figure 26 below for complete results. Also, see Figure 27 to see the same results plotted as a percent change. The average percent change is +1.2%.

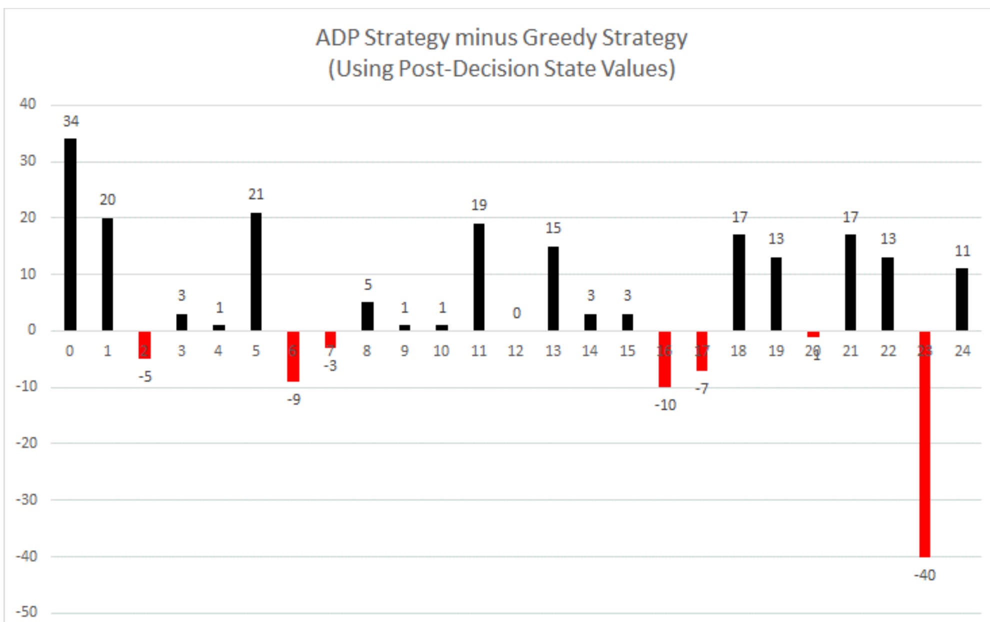

Figure 26: ADP minus Greedy (3c2 – Variant #1, n = 25)

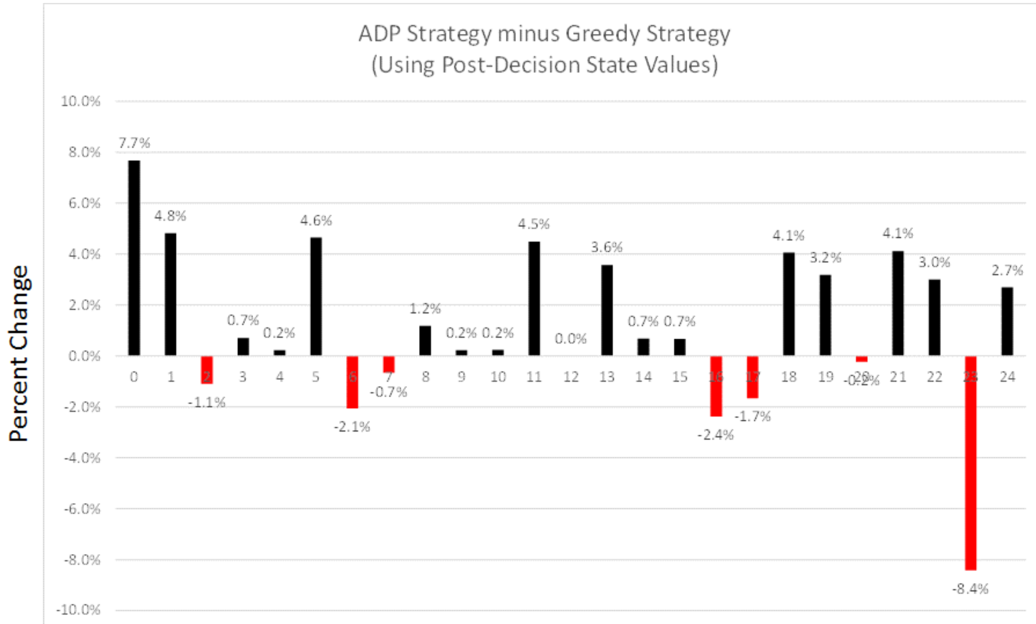

Figure 27: ADP minus Greedy – Plotted as a Percent Change (3c2 – Variant #1, n = 25)

5.2.2 Toy Problem (3c2 – Variant #2)

For Variant #2, I changed the value matrix (see Table 18) and the player times on each endurance test (see Table 19). In this scenario, Player 1 is still the best player on the team, but his decay rate is the fastest now and his recovery is in the middle. Player 2 is still the second best player, but his decay and recovery rate is the slowest. Finally Player 3 is still the worst player overall, but when he is at high endurance, his value is better than Player 1 at low endurance. Also, Player 3's decay rate is in the middle, and his recover rate is the fastest. The results for this scenario are below.

Table 18: Value Matrix for Players 1, 2, & 3 (3c2 – Variant #2)

	1	2	3
H	10	8	7
M	9	6	4
L	5	4	1

Table 19: Endurance Test Decay/Recovery Results for Players 1, 2, & 3 (3c2 – Variant #2)

	Player 1 (Minutes)	Player 2 (Minutes)	Player 3 (Minutes)
Test #1 (14.0 mph)	1	2	1
Test #2 (12.6 mph)	2	4	3
Test #3 (11.2 mph)	2.86	6	5
Test #4 (9.8 mph)	5	12	8

	Player 1 (Minutes)	Player 2 (Minutes)	Player 3 (Minutes)
Test #1 Recovery	9	14	7
Test #2 Recovery	8	13	6
Test #3 Recovery	7.5	12	6
Test #4 Recovery	7.5	11	5
Average Recovery	8	12.5	6

Using Bellman's equation, I solve for the ADP policy vector, which converges again after 89 iterations. Then, I change the gamma value from 0.9 to 0.0 to get the greedy policy vector. A portion of the two optimal policy vectors is shown below in Table 20. There were only 10 differences between the policy vectors. (Three of the differences are highlighted in yellow and shown in Table 20.) That means that approximately 12% (10 out of 81 states) of the ADP policy vector is different.

Table 20: Optimal Policy Vector Comparison between ADP and Greedy (3c2 – Variant #2)

State	Optimal Action (ADP)	Optimal Action (Greedy)
[0L,1L,1L]	Swap 1 and 3	Swap 1 and 3
[0L,1L,1M]	Swap 1 and 3	Swap 1 and 2
[0L,1L,1H]	Swap 1 and 2	Swap 1 and 2
[0L,1M,1L]	Swap 1 and 3	Swap 1 and 3
[0L,1M,1M]	Swap 1 and 3	Swap 1 and 3
[0L,1M,1H]	Swap 1 and 2	Do Nothing
[0L,1H,1L]	Swap 1 and 3	Swap 1 and 3
[0L,1H,1M]	Swap 1 and 3	Swap 1 and 3
[0L,1H,1H]	Do Nothing	Do Nothing
[0M,1L,1L]	Swap 1 and 3	Swap 1 and 3
[0M,1L,1M]	Swap 1 and 3	Swap 1 and 3
[0M,1L,1H]	Swap 1 and 2	Swap 1 and 2
[0M,1M,1L]	Swap 1 and 3	Swap 1 and 3
[0M,1M,1M]	Swap 1 and 3	Swap 1 and 3
[0M,1M,1H]	Swap 1 and 2	Swap 1 and 2
[0M,1H,1L]	Swap 1 and 3	Swap 1 and 3
[0M,1H,1M]	Swap 1 and 3	Swap 1 and 3
[0M,1H,1H]	Do Nothing	Swap 1 and 3
[0H,1L,1L]	Swap 1 and 3	Swap 1 and 3
[0H,1L,1M]	Swap 1 and 3	Swap 1 and 3
[0H,1L,1H]	Swap 1 and 2	Swap 1 and 2
[0H,1M,1L]	Swap 1 and 3	Swap 1 and 3
[0H,1M,1M]	Swap 1 and 3	Swap 1 and 3
[0H,1M,1H]	Swap 1 and 2	Swap 1 and 2
[0H,1H,1L]	Swap 1 and 3	Swap 1 and 3
[0H,1H,1M]	Swap 1 and 3	Swap 1 and 3
[0H,1H,1H]	Swap 1 and 3	Swap 1 and 3

Using these two policy vectors from Scenario 3c2 – Variant #2, I again simulate 25, 40-minute games for each strategy. (Seeded values are used for the defensive uncertainty, so that the comparison is fair.) The boxplots in Figure 28 show the cumulative post-decision state value of each game. It is clear to see that ADP policy vector performs better (when using the median, mode, or average as a comparison metric) than the greedy policy vector. See Figure 29 for descriptive statistics.

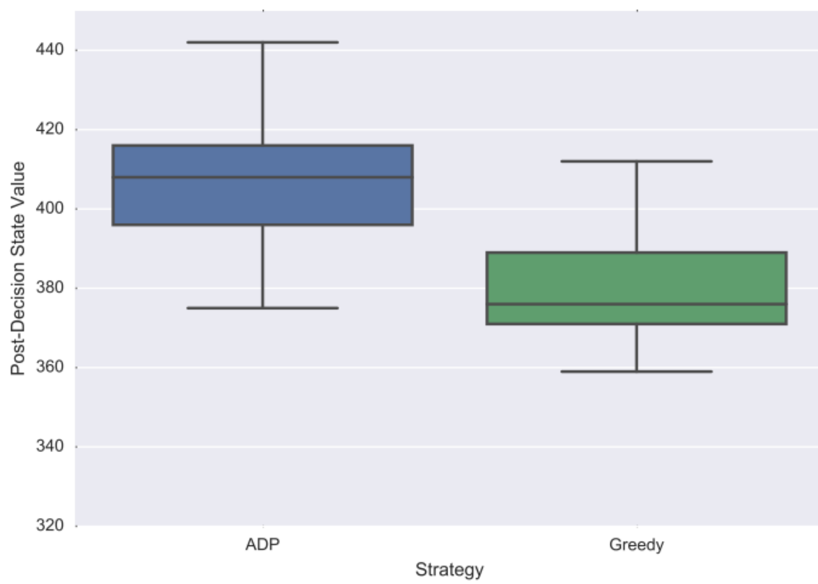

Figure 28: Boxplot Comparison of ADP vs. Greedy (3c2 – Variant #2, n = 25)

Post-Decision Value - ADP	
Mean	407.64
Standard Error	3.08
Median	408.00
Mode	396.00
Standard Deviation	15.38
Sample Variance	236.66
Kurtosis	0.19
Skewness	0.08
Range	67.00
Minimum	375.00
Maximum	442.00
Sum	10191.00
Count	25.00
Confidence Level(95.0%)	6.35

Post-Decision Value - Greedy	
Mean	379.72
Standard Error	2.70
Median	376.00
Mode	380.00
Standard Deviation	13.50
Sample Variance	182.21
Kurtosis	0.17
Skewness	0.80
Range	53.00
Minimum	359.00
Maximum	412.00
Sum	9493.00
Count	25.00
Confidence Level(95.0%)	5.57

Figure 29: Descriptive Statistics for ADP & Greedy Strategy (3c2 – Variant #2, n = 25)

To see how many times the greedy strategy outperforms the ADP strategy in Variant #2, I subtract the cumulative post-decision value of the ADP strategy from the greedy strategy of each seeded game. The greedy strategy never beat ADP in the 25 simulated games. In fact, only three of the games are even close. See Figure 30 below for complete results. Also, see Figure 31 to see the same results plotted as a percent change. The average percent change is now +7.4 (instead of +1.2%).

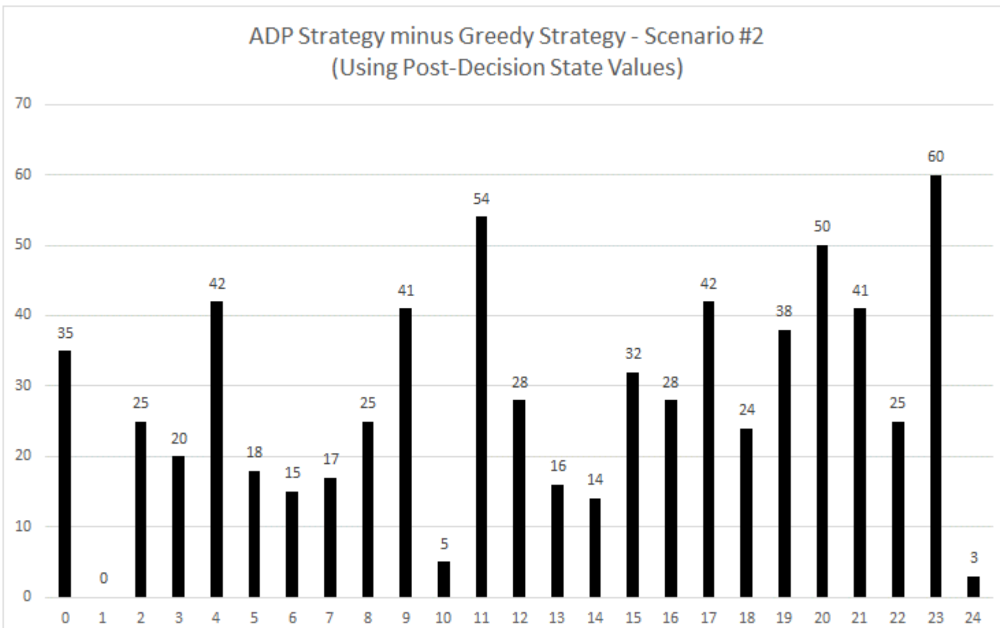

Figure 30: ADP minus Greedy (3c2 – Variant #2, n = 25)

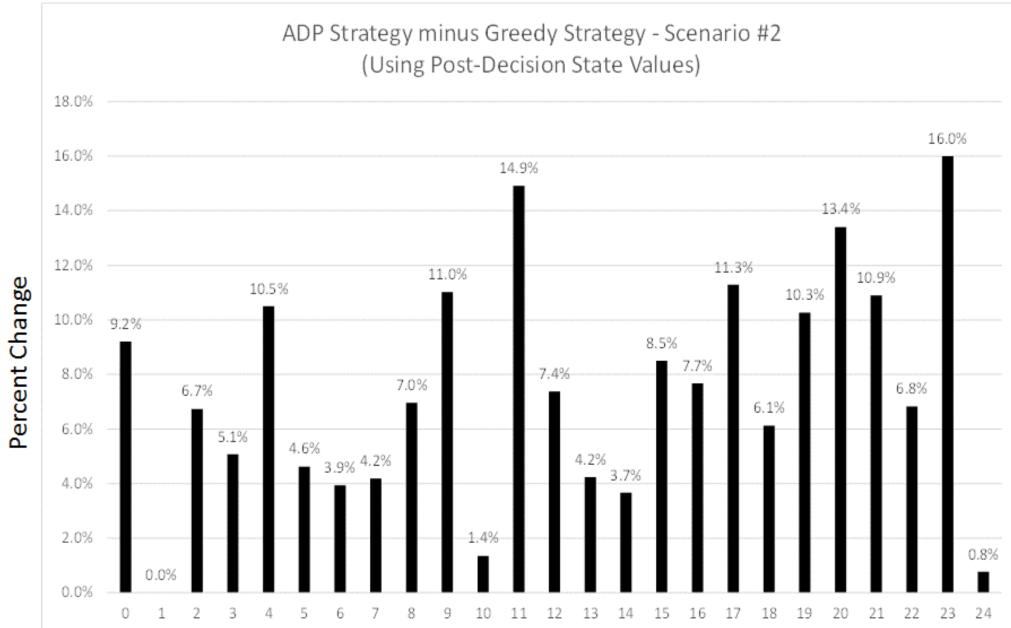

Figure 31: ADP minus Greedy – Plotted as a Percent Change (3c2 – Variant #2, n = 25)

5.2.3 Toy Problem (3c2 – Variant #2a)

Because simulating 25 games might be too small of a sample size, 1000 games (using the two policy vectors from Variant #2) for each strategy is simulated. The results in Figure 32 and Figure 33 are below. The blue and green vertical lines in Figure 32 represent each strategy's median. The ADP strategy's median (blue line) is 408 and the greedy strategy's median (green line) is 378. Figure 33 shows that the ADP strategy significantly beats the greedy strategy. (The number of occurrences where the greedy strategy beat the ADP strategy are shown to the left of the red line.) The greedy strategy did better than the ADP strategy less than two percent of the time. In other words, the ADP strategy outperforms the greedy strategy over 98% of the time. It is clear to see that the ADP strategy is better than the greedy strategy.

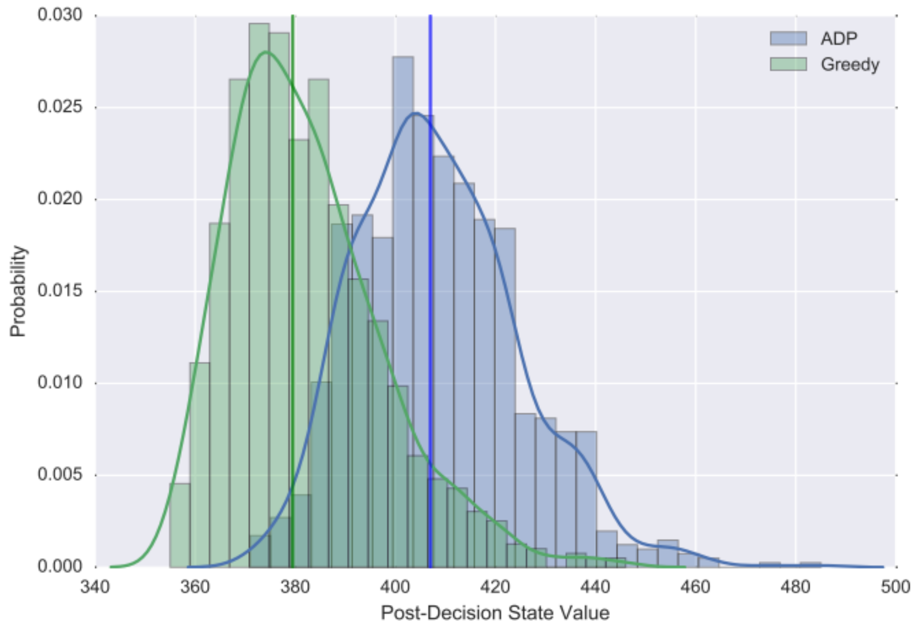

Figure 32: Density Plot Comparison of ADP vs. Greedy (3c2 – Variant #2a, n = 1000)

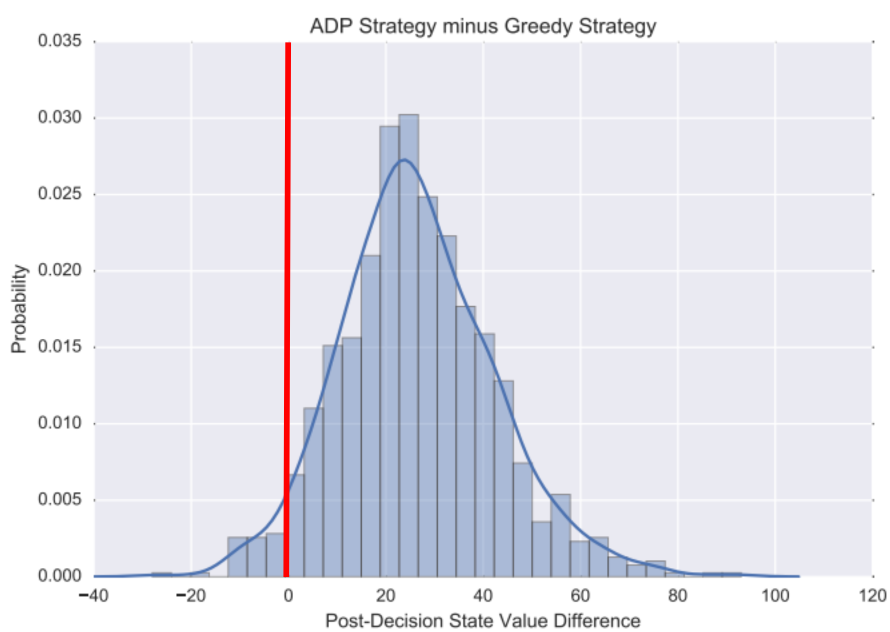

Figure 33: ADP minus Greedy (3c2 – Variant #2a, n = 1000)

5.3. Scenario 8c5(30*)

Given an 8-person team (*m* = 8) with five players on the court (*c* = 5) and three player on the bench (*b* = 3) implies that there will be 56 possible lineups (8 choose 5 equals 56). However, only the top 30 lineups are used in an effort to constrain the problem. (See discussion in Constraints section.) Using only three endurance levels (*l* = 3) means that each one of the 30 possible lineups would have 6,561 states (3^8= 6,561). 6,561 states per lineup means that there are 196,830 total states (30*6,561 = 196,830). Since the coach is limited to no more than one player swap at a time, there are only 16 feasible actions the coach can take. For example, from state [1H,1H,1H,1H,1H,0H,0H,0H], the 16 actions are to: (1) swap Player 1 with Player 6, (2) swap 1 with 7, (3) 1 with 8, (4) 2 with 6, (5) 2 with 7, (6) 2 with 8, (7) 3 with 6, (8) 3 with 7, (9) 3 with 8, (10) 4 with 6, (11) 4 with 7, (12) 4 with 8, (13) 5 with 6, (14) 5 with 7, (15) 5 with 8, or (16) do nothing.

I am going to follow the same limitations, constraints, and assumptions that were listed in the Scope section. Also, I am going to use the data that was collected and described in Chapter Four.

5.3.1 ADP vs. Greedy vs. Rotation (Scenario 8c5(30*))

From the endurance test results found in the Section 4.4, it is clear that Players 1, 3, 4, and 7 decay the slowest out of the eight players. Players 1, 3, 6, and 7 recover much faster than the other four players. Player 1, 3, and 7 are in the top four for both decay and recovery. However, only player 7 is in the top four for skill. Please see Table 21. It is obvious to see that player 7 will be utilized quite frequently, regardless of strategy, due to his skill set combined with his great fitness level. But determining the makeup of the four remaining players, particularly when using the ADP strategy, is murky at best.

Table 21: Rank Order of Each Player for Skill, Decay, and Recovery

Player #	Skill	Decay	Recovery
1	6	1	3
2	2	6	6
3	5	3	1
4	7	4	5
5	8	7	7
6	3	5	2
7	1	2	4
8	4	8	8

Since the opponent is equally likely to use all four defenses (see Table 10), the overall decay matrix (not shown) can be calculated in the same way that is explained in the Section 4.3.1.

From the value matrix (Table 7), we determine the contribution matrix. This contribution matrix is 196,830 states by 16 feasible actions. For example, for state [1H,1H,1H,1H,1H,0H,0H,0H], the contribution for the 'do nothing' decision is the calculated by summing the individual total values of players 1-5 at the associated endurance level (0.460+0.748+0.572+0.136+0.054 = 1.970).

Solving Equation 4 by recursively computing Bellman's optimality equations (see Equation 5), convergence occurs (with a convergence threshold of ≤ 0.01[9]) after 10.1+

[9] When solving Bellman's equation, the max value for each and every state will eventually converge. Achieving a threshold of 0.01 means that the max value from one iteration to the next does not change by more than 0.01. Setting the threshold to 0.01 is sufficient because we expect to see converged values between 0 and 15. Once the max value for each and every state converges, then the decision (from the 16 feasible decisions) that gave that converged number is the optimal decision for that state.

million iterations for asynchronous ADP, which means it solves to optimality. By changing the gamma value from 0.9 to 0.0, I obtain the greedy policy vector after only 1.3+ million iterations. A portion of the two optimal policy vectors is shown below in Table 22. Almost 47% of these two policy vectors differ. (Six of the differences are highlighted in yellow and shown in Table 22.) Please note that the rotation strategy does not need a policy vector in order to determine its next action. The rotation action is already pre-determined and is explained in the Assumptions section.

Table 22: Optimal Policy Vector Comparison between ADP and Greedy (Scenario 8c5)

State	ADP Action	Greedy Action
1H,1H,0H,1H,0H,1M,1M,0H	Swap 3 and 6	Swap 1 and 8
1H,1H,0H,1H,0H,1M,1H,0L	Swap 3 and 6	Swap 3 and 6
1H,1H,0H,1H,0H,1M,1H,0M	Swap 3 and 6	Swap 3 and 6
1H,1H,0H,1H,0H,1M,1H,0H	Swap 3 and 6	Swap 1 and 8
1H,1H,0H,1H,0H,1H,1L,0L	Swap 7 and 8	Do Nothing
1H,1H,0H,1H,0H,1H,1L,0M	Swap 7 and 8	Swap 7 and 8
1H,1H,0H,1H,0H,1H,1L,0H	Swap 7 and 8	Swap 7 and 8
1H,1H,0H,1H,0H,1H,1M,0L	Swap 7 and 8	Do Nothing
1H,1H,0H,1H,0H,1H,1M,0M	Swap 7 and 8	Do Nothing
1H,1H,0H,1H,0H,1H,1M,0H	Swap 7 and 8	Swap 1 and 8
1H,1H,0H,1H,0H,1H,1H,0L	Do Nothing	Do Nothing
1H,1H,0H,1H,0H,1H,1H,0M	Do Nothing	Do Nothing
1H,1H,0H,1H,0H,1H,1H,0H	Swap 1 and 8	Swap 1 and 8

Using these two policy vectors and the rotation policy (specified in the Assumptions section), I simulate 1,000, 40-minute games for each strategy. (Seeded values are used for the defensive uncertainty, so that the comparison is fair.) The boxplots and density plots in Figure 34 and Figure 35 show the cumulative post-decision state value of each game. It is clear to see that ADP policy performs better (when using the median as a comparison metric) than the rotation or greedy policy.

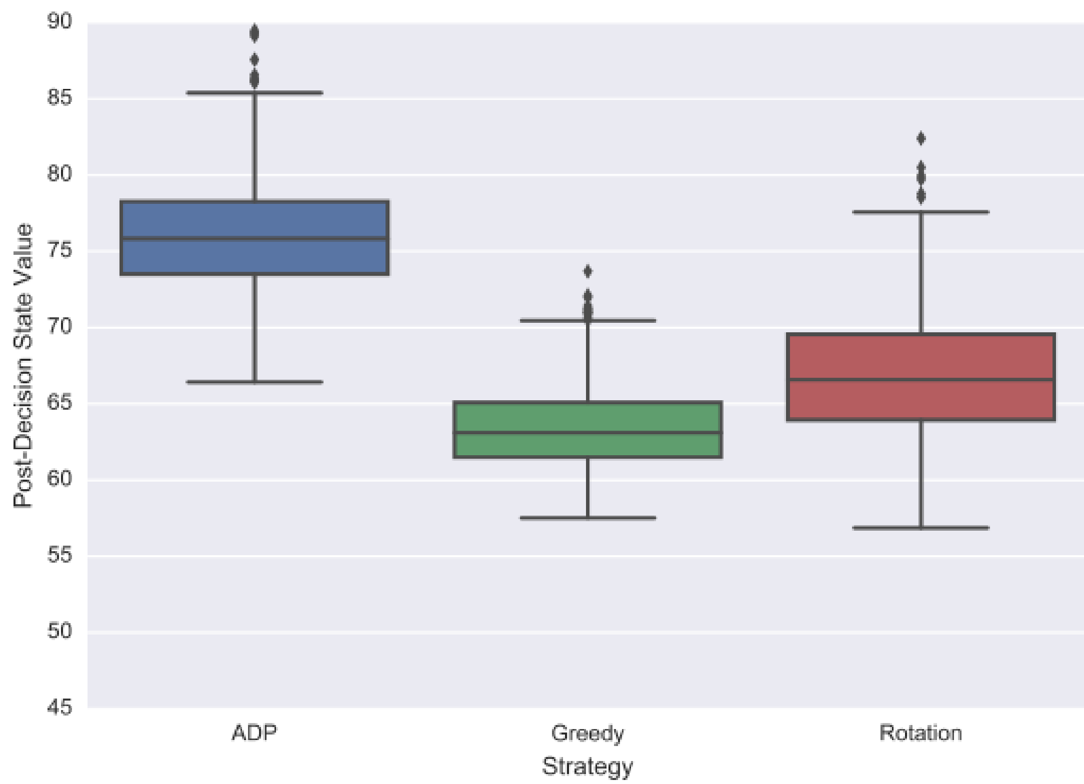

Figure 34: Boxplot Comparison of ADP vs. Greedy vs. Rotation (Scenario 8c5(30*), n = 1000)

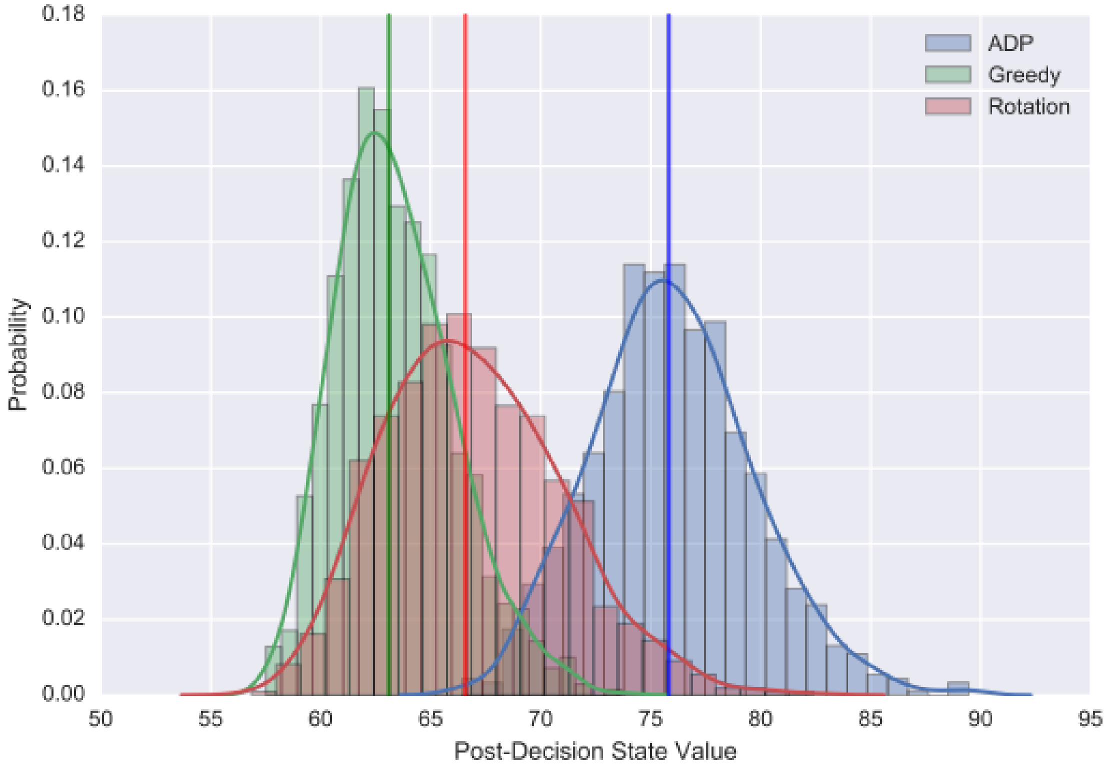

Figure 35: Density Plot Comparison of ADP vs. Greedy vs. Rotation (Scenario 8c5(30*), n = 1000)

5.3.2 ADP vs. Rotation (Scenario 8c5(30*))

While it might appear from looking at the boxplot and density plot (Figure 34 and Figure 35) that the rotation strategy occasionally outperforms ADP, that notion is completely false. When I subtract the cumulative post-decision value of the ADP strategy from the rotation strategy of each seeded game, I find that ADP beats rotation every single time. In fact, ADP improves the cumulate post-decision state value between 2% and 34%. Please refer to the density plot (Figure 36) and cumulative density function (CDF) plot (Figure 37) that shows percent change between the ADP strategy and rotation strategy. Since all the percent changes are positive, that proves that rotation never beats

ADP. While the CDF is plotting the same data as the density plot, the CDF is better for showing that there is an 80% chance of ADP outperforming rotation by approximately 10%-34%.

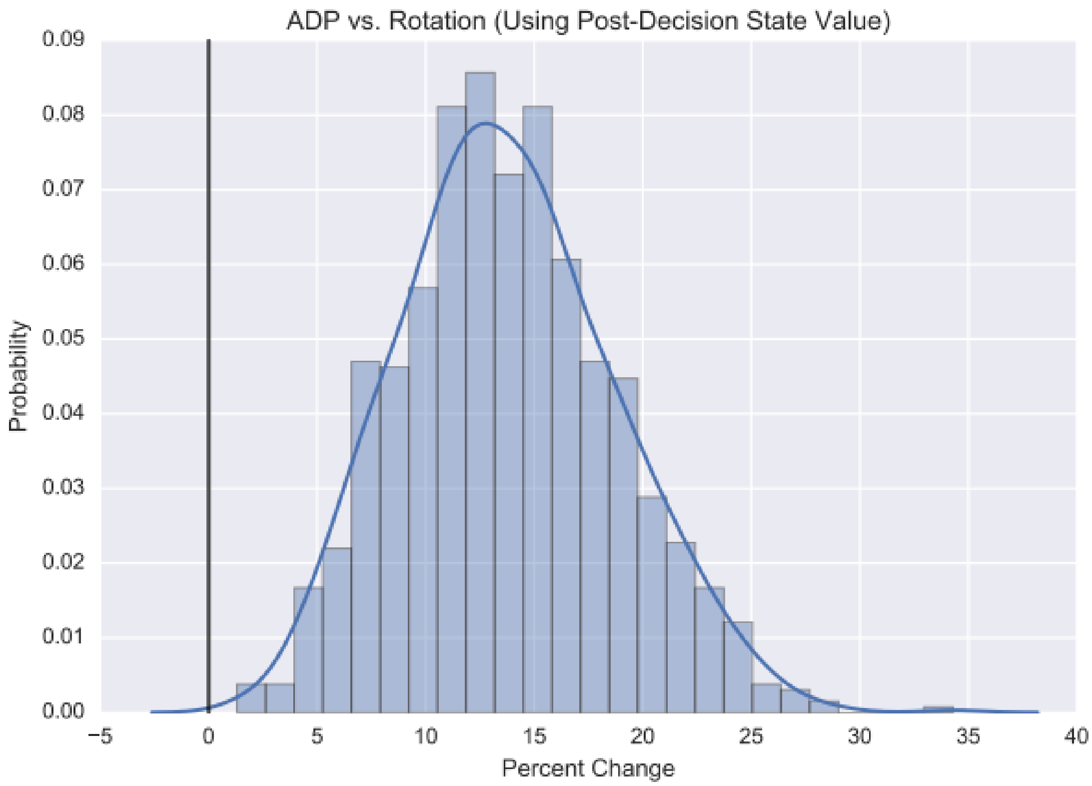

Figure 36: Density Plot of ADP minus Rotation – Plotted as a Percent Change (Scenario 8c5(30*), n = 1000)

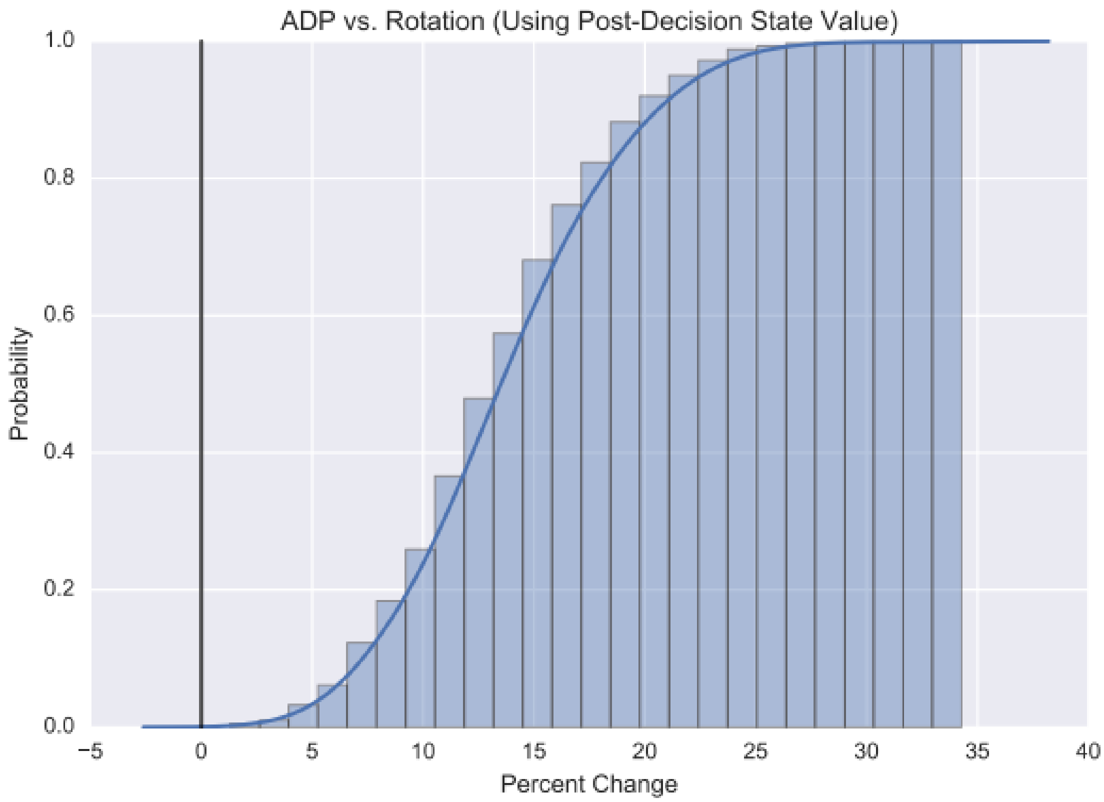

Figure 37: CDF Plot of ADP minus Rotation – Plotted as a Percent Change (Scenario 8c5(30*), n = 1000)

5.3.3 ADP vs. Greedy (Scenario 8c5(30*))

When I subtract the cumulative post-decision value of the ADP strategy from the greedy strategy of each seeded game, I find that ADP outperforms greedy one hundred percent of the time as well. In fact, ADP improved the cumulate post-decision state value between 8% and 34% when compared to greedy. Please refer to the density plot (Figure 38) and CDF plot (Figure 39) that shows percent change between the ADP strategy and greedy strategy. Since all the percent changes are positive, that shows that greedy never beats ADP. While the CDF is plotting the same data as the density plot, the CDF is better for showing that there is an 80% chance of ADP outperforming greedy by approximately 15%-34%.

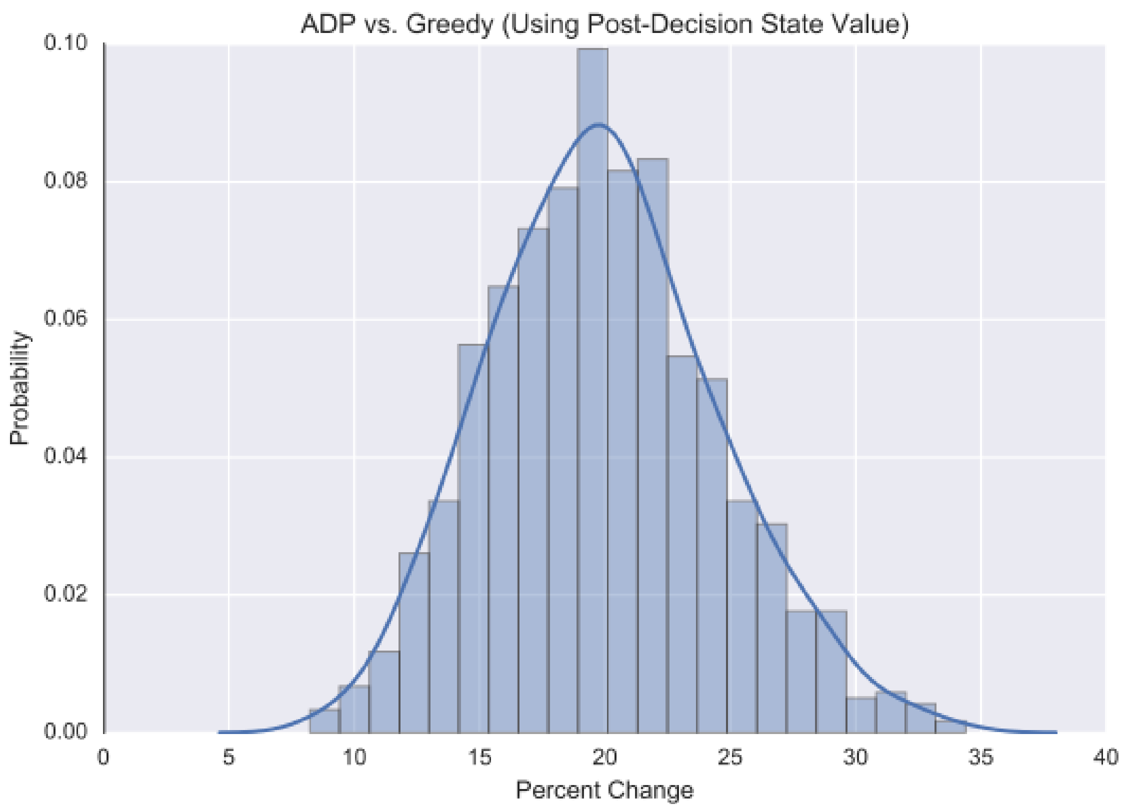

Figure 38: Density Plot of ADP minus Greedy – Plotted as a Percent Change (Scenario 8c5(30*), n = 1000)

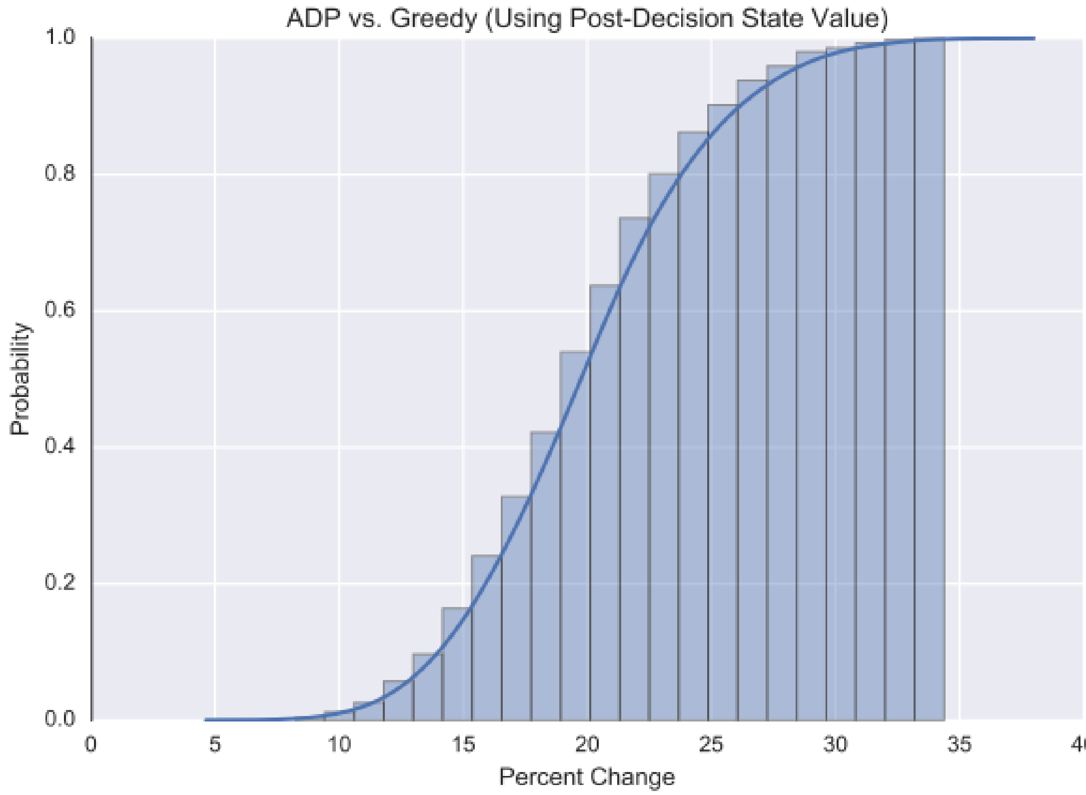

Figure 39: CDF Plot of ADP minus Greedy – Plotted as a Percent Change (Scenario 8c5(30*), n = 1000)

I want to understand what causes ADP to vary between a positive percent change of 8% and 34% when compared to greedy. What kind of scenarios produce the larger improvements? I hypothesize that larger differences between post-decision values of ADP and greedy correlate with less defensive intensity played against a team over 40 minutes. Since I assign the zone defense to a value of zero (0), man-to man to a value of one (1), half-court pressure to a value of two (2), and full-court pressure to a value of three (3), then the theoretical maximum value for defensive intensity for a 40-minute game is 120 (40*3). The theoretical minimum is zero (0) if zone is played the entire game. Now, reaching this theoretical minimum and maximum depend greatly on the historical defensive scheme probabilities (Table 10). Since the opponent's likelihood of using each defense is the same (25%), then the chance of getting a theoretical minimum or maximum is highly unlikely. But when I conduct a Monte Carlo simulation of 1,000 games, there will definitely be games that do not follow the distribution found in Table 10. Please refer to Figure 40 to see the scatterplot of the post-decision value difference of 1,000 seeded games against the cumulative defensive intensity of those seeded games. From the scatterplot, it is evident that the post-decision value difference is negatively correlated with cumulative defensive intensity. That means that, generally speaking, as the cumulative defensive intensity goes up, the post-decision value difference goes down.

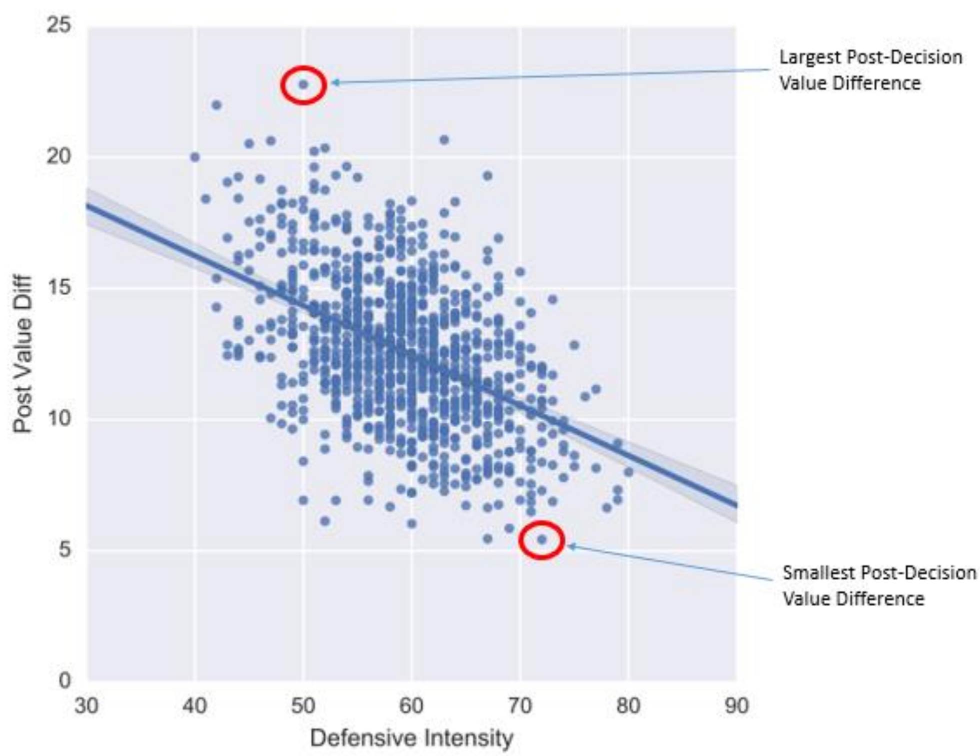

Figure 40: Post-Decision Value Difference vs. Cumulative Defensive Intensity (Scenario 8c5(30*), n = 1000)

While the negative correlation is interesting, I want to find out more about what is exactly happening in each game (who is in when and when are certain defenses played). So, I first decide to investigate the two games (ADP and greedy) that produced has the highest post-decision value difference (about 23). I use a combination plot of a heat map and two bar charts (please see Figure 41 and Figure 42). These combination plots detail: (1) who is in/out of the game each minute – if a player is in the game at a particular minute, a shaded rectangle is present, (2) the individual contribution of players in the lineup each minute – the darker the shading, the more value the player provides (3) the overall value that each lineup provides per minute – represented by the top bar chart, (4) the overall value that each player provides over the entire game – depicted by the bar chart on the right side, and (5) the defense played each minute – shown along the bottom of the combination plot. By comparing the details, valuable insights are made.

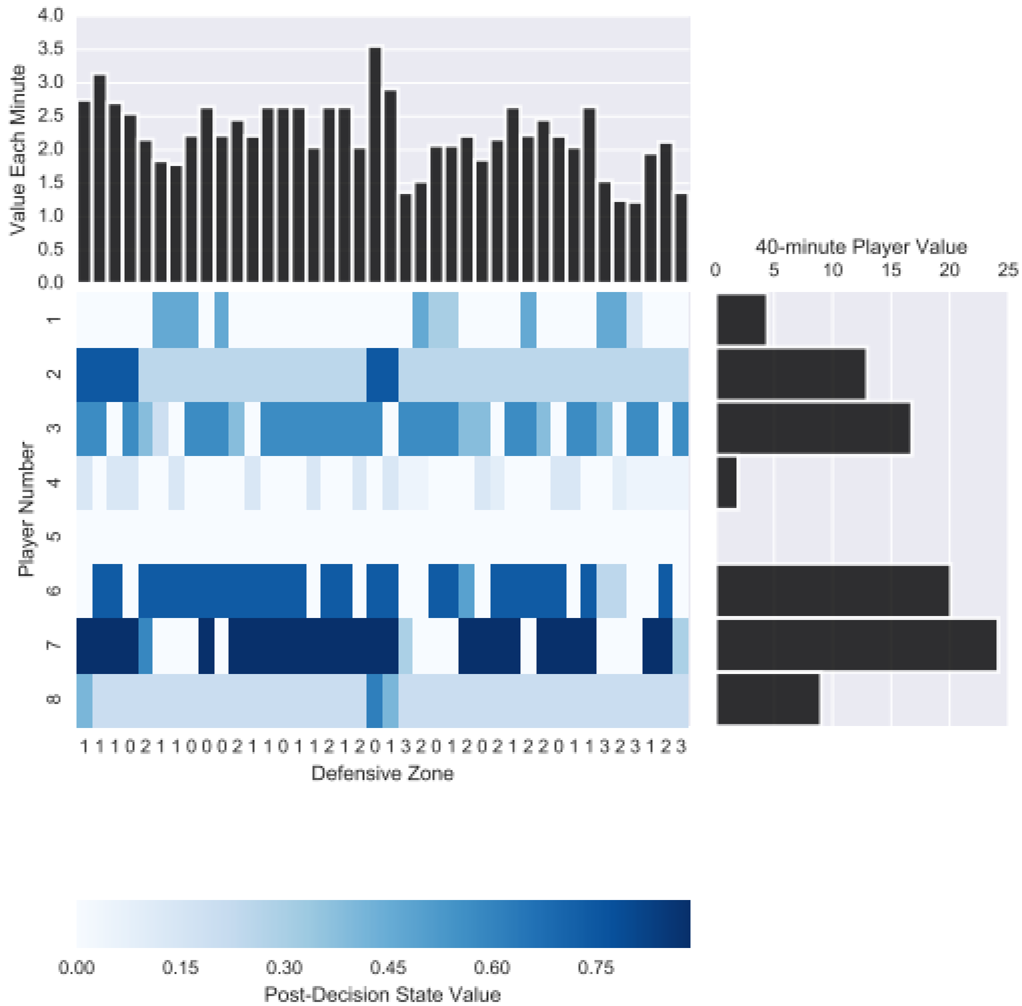

Figure 41: Single Game Using ADP Strategy – Largest Post-Decision Value Difference (Scenario 8c5(30*))

The expected cumulative defensive intensity for every game (when r = 40) is 60 (0*(40*.25) + 1*(40*.25) + 2*(40*.25) + 3*(40*.25) = 60). However, the cumulative defensive intensity from Figure 41/Figure 42 is 50. Only four full-court pressure defenses occur when 10 are expected (40*.25 = 10). Also, three of the four full-court pressure defenses occur in the last six minutes of the second half. Thus, this less intense defensive set for the first 34 minutes does not decay players as quickly. And when they finally do decay, there is sufficient time for the ADP strategy to rotate players in and out, which keeps their values higher throughout the game. The greedy strategy, when subjected to this same less intense defensive set, decays similarly for the first four minutes, but once the first half-court pressure defense hits in minute five, all players are significantly decayed. Once a player on the bench is slightly better than a player on the court, the greedy strategy calls for the immediate substitution. So, all players end up staying in the low endurance state or moving to the low endurance state within a few minutes. If they happen to recover to medium endurance when on the bench, they are immediately put back in the game and find themselves in the low endurance state in a minute or two. No player is ever able to recover to a high endurance state under the greedy strategy. Therefore, the ADP strategy accumulates more value over the greedy strategy for a large portion of the game (minutes 5-19 & 21-34).

The greedy strategy never puts player 4 in the game, while the ADP strategy puts player 4 in the game 19 times! Player 4, who is the seventh ranked skilled player, is coming in the game in order to allow others to recover to a higher endurance level.

Player 7 provides the most cumulative value over the entire game in both strategies. However, player 6 provides the next most overall value under the ADP strategy, followed closely by player 3. Under the greedy strategy, player 2 has the second most cumulative value, followed by player 6.

It is also interesting to note that under the greedy strategy, players 2, 3, 6 and 8 play the entire game, whereas with the ADP strategy, only players 2 and 8 play the entire game. Finally, player 5 is so bad that he does not play under either strategy.

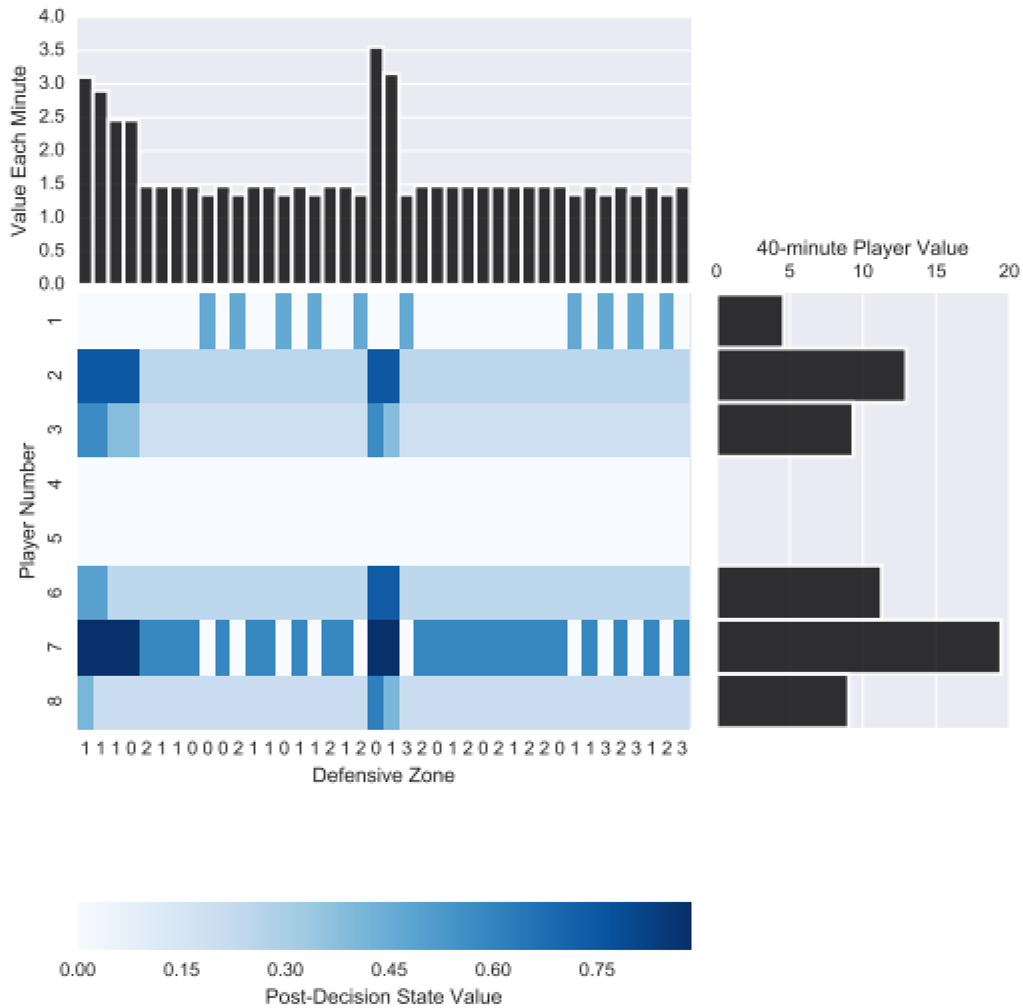

Figure 42: Single Game Using Greedy Strategy – Largest Post-Decision Value Difference (Scenario 8c5(30*))

Next, I investigate the two games (ADP and greedy) that produced has the lowest post-decision value difference (around 5). I again use a combination plot of a heat map and two bar charts (please see Figure 43 and Figure 44).

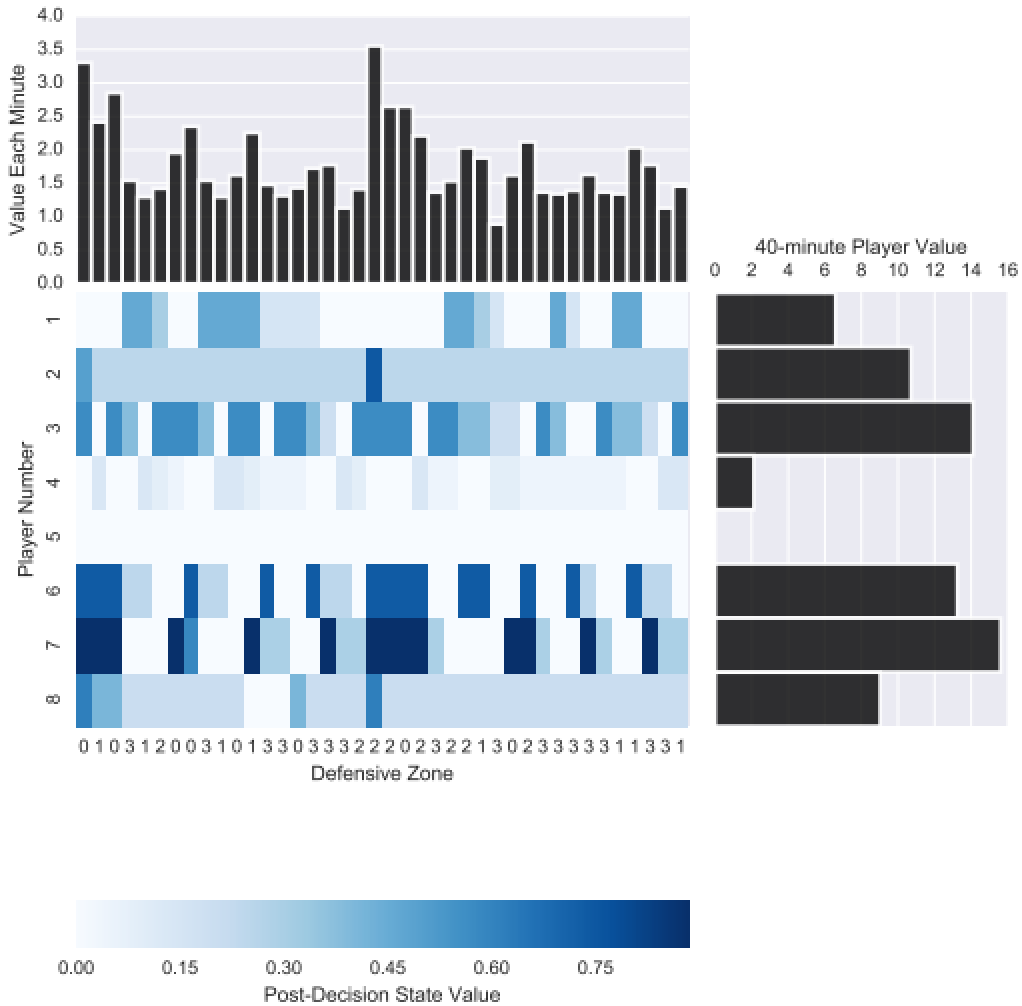

Figure 43: Single Game Using ADP Strategy – Smallest Post-Decision Value Difference (Scenario 8c5(30*))

The cumulative defensive intensity of the game depicted in Figure 43/Figure 44 is 72. Sixteen full-court pressure defenses occur, when only 10 are expected. The timing of these full-court pressure defenses occur throughout both halves. Players are decayed much quicker in this scenario due to the high intensity defenses. This results in less cumulative value for the ADP strategy. The greedy strategy yields results that are similar to the previous scenario (less intense defensive set). Players under the greedy strategy again find themselves in the low endurance state after the first high intensity defense (which occurred in minute four), and they basically remain there until halftime as the high intensity defenses (nine total) keep them fatigued. The second half yielded similar results. Thus, the ADP strategy is not able to accumulate as much value over the greedy strategy throughout the game.

Again, the greedy strategy never put player 4 in the game, while the ADP strategy put player 4 in the game 26 times! Player 4 (the seventh ranked skilled player) is still coming in the game in order to allow others to recover to a higher endurance level.

Player 7 again provides the most cumulative value over the entire game in both strategies. However, player 3 (instead of player 6) provides the next most overall value under the ADP strategy, followed closely by player 6 (instead of player 3). Under the greedy strategy, player 6 (instead of player 2) has the second most cumulative value, followed by player 2 (instead of player 6).

Players 2, 3, 6 and 8 again play the entire game under the greedy strategy, whereas with the ADP strategy, only player 2 (instead of players 2 and 8) plays the entire game. Finally, player 5 is still so bad that he does not play under either strategy.

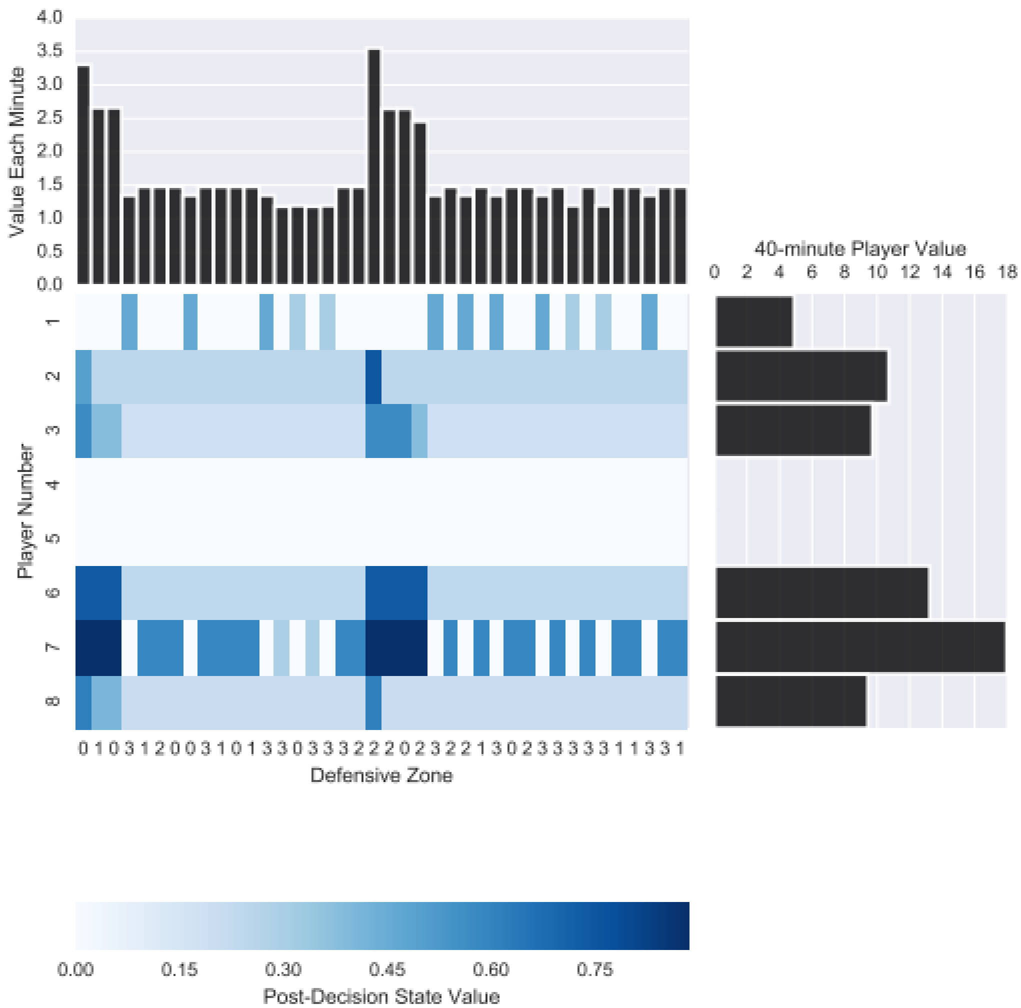

Figure 44: Single Game Using Greedy Strategy – Smallest Post-Decision Value Difference (Scenario 8c5(30*))

5.4 Discussion

The ADP model gives more value over the course of a game regardless of the size of the basketball team. But getting to the solution is not easy, cheap, or fast. It requires inputs from the team as well as (an) expensive computer(s) with lots of computing power. Finally, it takes patience as the computer solves for the optimal policy vector π^*over a period of days, not minutes! Since college teams play 2-3 games per week, the coach may have to invest in multiple computers so that each computer can be solving the optimal policy vector π^* for that week's opponents.

While the ADP strategy gives a 10%-20% (on average) improvement over the rotation and greedy strategies, one might wonder if that kind of improvement is worth the ADP investment of time, money, and effort. Well, in the 2017 NCAA tournament, approximately 30% of all games were decided by five points or less. Would those tight games be so close if the coach had the "right" players on the floor throughout the game? Using the ADP strategy to win a few close games (that might been have lost if another strategy was used) could easily affect the team's seeding in the conference tournament and/or translate into making the post-season tournament all together!

If one is not able to use the ADP model, the rotation strategy is decent substitute when compared to the greedy strategy. The reason is because the rotation strategy facilitates recovery of the players (which brings their value higher) whereas the greedy strategy depletes the resource (the player's value) as soon as it is available.

This research did confirm that using the greedy strategy at the end of the first half (last 2-4 minutes) is a viable option. The 15-minute halftime is more than sufficient to fully recover all players, regardless of their recovery rate. The only reason a coach may opt against using the greedy strategy at the end of the second half is because there is a chance the game could go into overtime, and the break between the end of regulation and overtime is only five minutes.

While ADP outperformed greedy and rotation every single time when $m = 8$, $c = 5$, and $b = 3$, ADP did even better when the defensive intensity was less than predicted.

Also, when the defensive intensity right after the start of the game (or right after halftime) was high, the ADP model did not perform as well. So, the take away for the coach should be to monitor the intensity of the defense. If the intensity is much higher than expected, then the coach should be aware that the ADP policy vector may not produce as favorable of results.

If a coach notices in the ADP policy vector that a player is never being recommended to enter the game (like player 5), then the coach should look closely at the data and see what can be improved. In player 5's case, he decays the second fastest and recovers the second slowest. So, the coach should focus efforts on getting player 5 in better shape. Player 5 is also the lowest ranked skilled player. Thus, player 5 must improve in the metrics that the coach values. Otherwise, player 5 will continue to sit the bench the entire game.

A coach, who is open to change, is able to enhance or supplement his/her current coaching strategy by gaining insights from this research. But the extra knowledge does not come without time, money, and effort. A coach must: (1) be willing to acquire and/or maintain heart rate monitors and high speed computers, (2) regularly test (at least monthly) the endurance/recovery of their players, (3) update the value matrix as more information becomes available, and (4) scout the defensive tendencies of each opponent. By doing these four things, the optimal policy vector π^* from the ADP model will be more accurate and meaningful, and hopefully lead to more wins.

CHAPTER SIX – RESEARCH CONTRIBUTIONS AND CONCLUSIONS

This research marks the first time approximate dynamic programming (ADP) is used to improve substitution strategies in the game of basketball. This is also the first time heart rate data (from four endurance/recovery tests conducted during practice) is utilized to serve as a proxy for decay and recovery rates in an actual basketball game. My ADP model takes these decay and recovery rates for each player and simulates endurance level changes based on whether they are playing on the court or recovering on the bench. These four endurance/recovery tests give coaches insights into their player's fitness levels that they did not have before. The results from the four endurance/recovery tests help bridge the gap between current substitution strategies of today (heuristics, player rotation, next best (greedy), go with your gut, etc.) and the substitution strategy of the future (when live, in-game HR monitoring is approved for wear by the NCAA during a college basketball game, thus allowing coaches to know precisely when a player is fatigued). See Figure 45.

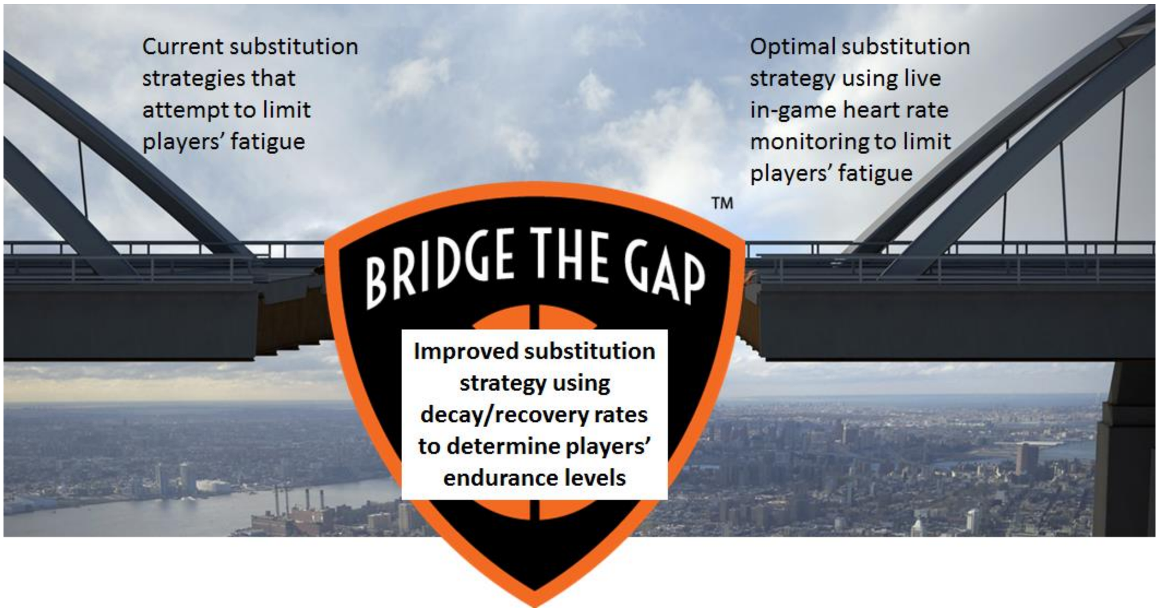

Figure 45: Bridge the Gap between Current and Future Substitution Strategies

Another research contribution stems from the qualitative value model that is created from discussions with the head coach and their coaching staff. The value hierarchy that is produced includes quantifiable metrics that help the coach measure the value of each player/lineup. The ADP model then uses the outputs from the value model to populate the contribution/reward matrix that is needed to solve Bellman's equation. This method of using value model results to populate the contribution/reward matrix has never been done before in ADP.

With the value model and endurance/recovery test inputs, the ADP model is complete. A policy vector π is solved for that prescribes what substitution, if any, a coach should make at the end of each minute of the game given the current state of the players (whether they are in/out and their current endurance levels). The resulting analysis shows that lineups generated from my ADP model provide more overall value over the course of an entire game when compared with a greedy or player rotation strategy.

A basketball team benefits from an analytic approach that prescribes an improved substitution policy. The ADP results aid a coach in their decision making, giving him/her more brain capacity to think about other important decisions he/she may have to make. Furthermore, this research helps mitigate the risk of leaving a player in too long and/or helps a coach decide between an ADP, greedy, or player rotation strategy at the end of each half. Until HR monitors are approved for wear during an actual game, this research potentially redefines how coaches approach their substitution decisions.

CHAPTER SEVEN – FUTURE WORK

This chapter includes two sections: future research and parallel research. As was mentioned in Chapter Two, very little formal research is currently being done in the area of substitution strategies for basketball. However, as the interest continues to grow, so will the research. The future research section points out potential research areas to further this work. The parallel research section discusses possible fields where this methodology can be used.

7.1 Future Research

There is a lot of potential to further this research. Considering that some college basketball teams play upwards of 12 players, the state space explodes even further. A '12 choose 5' scenario (with only three endurance levels) has over 420 million total states. Just imagine how many computations would have to be completed for each action-specific TPM. (There would be 36 action-specific TPMs.). Maybe there is a better approach for solving for the optimal substitution strategy, given extremely large state spaces.

Until HR monitors are approved for wear in a basketball game, there will be a need to model decay and recovery rates of players. One area of future research could be challenging the assumption that decay and recovery rates are linear. Another area of future research could explore the decay rates between more than three endurance levels.

Also, additional research could help validate or refute the assumption that decay and recovery rates are able to be modeled by Markov chains.

Many aspects of a player's effectiveness are tied to the human condition. How well a player slept last night, the quality of the food they eat, whether or not their relationships are going well are just some examples that can influence a player's effectiveness. Some of these are qualitative or subjective, and do not easily translate into quantifiable terms. Further research into how these areas impact decay and recovery rates as well as overall player value may be warranted in the future.

There may be a better way to model the differences in defensive intensity and how that correlates with a player's decay rate. If one could get access to data from the player tracking cameras that the NBA uses, then there is a possibility that the intensity of player movements could translate into more accurate decay rates.

In closing, it is clear that this topic could benefit from additional research. Until HR monitors are approved for in-game use, the need for additional research in this area will be required to bridge the current knowledge gap. If coaches are more informed and able to make better substitution decisions, then it will likely result in better outcomes (i.e. more wins).

7.2 Parallel Research

There are numerous fields where one could use this methodology to tackle similar "substitution" problems. As a former U.S. Army CH-47D Chinook helicopter pilot, I can certainly see parallels for this research in aviation. Specifically, determining the fatigue and recovery rates of individual pilots and crew members could make flying safer.

Currently, aviation uses rule sets to govern crew rest as well as the length of missions. For example, Army Aviation requires 12 hours of crew rest between missions, and missions can be no longer than eight hours in length, unless approved by a higher command. But what if a pilot fatigues in five hours, due to uncertain events like taking hostile fire from the enemy or an engine failure? What if the pilot is not fully rested (recovered) by the 12 hour mark? Putting or keeping fatigued pilots in the cockpit puts the mission at risk. It certainly would be nice to know when the optimal "substitution" should occur.

Another parallel field could be an infantry squad or platoon that is conducting a three-day movement to contact. Each member of the squad or platoon is fatiguing and recovering at different rates. The rate at which they fatigue is dependent on the load they are carrying, their fitness level, amount of rest they are getting, etc. Thus, a squad or platoon's efficiency in movement is only as good as their weakest link. And if the squad or platoon leader push their weakest members too far without stopping to rest/recover, then they may end up moving even slower or not accomplish the mission. If the mission cannot be delayed, then making a "substitution" might make sense.

Hospitals could also benefit from this research. Doctor and nurses fatigue and recover at distinct rates. If I was about to be operated on, I would definitely want a medical team that was well rested and alert. Hospitals and, in particular, emergency rooms can be stressful, and thus cause premature fatigue. Therefore, knowing when to "substitute" a member of the medical staff could save lives and money (if a lawsuit is avoided).

It is evident from these three examples that there are parallels for this research in many different fields. If live, physiological monitoring is not practicable or possible, then using the methodology detailed in this dissertation could be very useful. Substituting fatigued individuals in real-life situations, outside of sports, could decrease the riskiness of a mission or even save a life!

APPENDIX A (METRICS AND THEIR FORMULAS)

Common player performance/effectiveness metrics and their associated formulas are found below.

Table 23: Player Performance/Effectiveness Metrics and their Formulas

Player Performance/Effectiveness Metric	Forumula
Effective Field Goal Percentage (EFG)	EFG = (all FGs made + 0.5(3pt FGs Made)) / (all FGs attempted)
Turnovers Committed per Possession (TPP)	TPP = # of TOs per game / Total possessions per game
Offensive Rebounding Percentage (ORP)	ORP = # of offensive rebounds / # of missed shots
Free Throw Rate (FTR)	FTR = foul shots made / FGs attempted
Opponent's Effective Field Goal Percentage (OEFG)	OEFG = (all opponent's FG made + 0.5(3pt FGs made by opponent)) / all FGs attempted by opponent)
Defensive Turnovers Committed per Possession (DTTP)	DTPP = # of opponent's TOs per game / Total opponent's possessions per game
Defensive Rebounding Percentage (DRP)	DRP = # of defensive rebounds / # of opponent's missed shots
Opponent's Free Throw Rate (OFTR)	OFTR = foul shots made by opponent / FGs attempted by opponent
NBA Efficiency Rating / Game	Efficiency per game = (points per game) + (rebounds per game) + (assists per game) + (steals per game) + (blocked shots per game) - (turnovers per game) - (missed FGs per game) - (missed FTs per game)
Player Efficiency Rating (PER)	Very complex
Game Score Rating	game score = (points x 1) + (FGM x 0.4) + (FGA x -0.7) + ((FTA-FTM) x -0.4) + (OREB x 0.7) + (DREB x 0.3) + (STL x 1) + (AST x 0.7) + BLK x 0.7) + (PF x -0.4) + (TO x -1)
Win Scores	This formula is a close approximation to Win Scores: Win Scores = points + rebounds + steals + 0.5(assists) + 0.5(blocked shots) - FGA - TOs - 0.5(FTA) - 0.5(PFs)
+/- Ratings	Team points scored while player is in minus opponent team points scored while player is in
Adjusted +/- Ratings	Needs Excel Solver
WINVAL Points Rating	points rating = offensive rating - defensive rating
Expected Possession Value - Added	Unknown and very complex

REFERENCES

Achten, Juul and Asker E. Jeukendrup (2003) "Heart Rate Monitoring: Applications and Limitations." Sports Med, 33:7, 517-538.

Adkins, Clint, Steven Bain, Edward Dreyer, Robert A Starkey (2007) Basketball Drills, Plays, and Strategies: A Comprehensive Resource for Coaches. Cincinnati, Ohio: Betterway Books.

Attali, Yigal (2013) "Perceived Hotness Affects Behavior of Basketball Players and Coaches." Psychological Science, XX:X, 1-6, DOI: 10.1177/0956797612468452.

Balakrishna, P. (2009) "Scalable Approximate Dynamic Programming Models with Applications in Air Transportation." George Mason University. Retrieved from http://mars.gmu.edu/handle/1920/5619?show=full

Bangsbo, Jens, F. Marcello-Iaia, and Peter Krustrup (2008) "The Yo-Yo Intermittent Recovery Test: A Useful Tool for Evaluation of Physical Performance in Intermittent Sports." Sports Med, 38:1, 37-51.

Bellman, R. (2003). Dynamic Programming. Mineola, NY: Dover Publications.

Birge, J.R. and F. Louveaux (2011) Introduction to Stochastic Programming. Second Edition. New York: Springer Publications.

Bourgase, Brock (2016) "Substitution Strategies." Accessed online, 13 December 2016. http://www.bourgase.com/coaching/team-building/substitution-strategies/

Brown, Bruce (2015) "A Self-Substitution System for Basketball." Accessed online, 13 December 2016. https://www.nfhs.org/articles/a-self-substitution-system-for-basketball/

Clay, D. C., and K.E. Clay (2014) "Player Rotation, On-Court Performance and Game Outcomes in NCAA Men's Basketball." International Journal of Performance Analysis in Sport, 14:2, 606-619.

Denardo, Eric V. (1982) Dynamic Programming: Models and Applications. Englewood Cliffs, NJ: Prentice-Hall Publications.

Fagan, Kate (2010) "Jordan's Substitution Pattern Comes into Question." The Philadelphia Inquirer, January 7, 2010. Accessed online, 06 Dec 2016. http://www.lexisnexis.com/hottopics/lnacademic/?shr=t&csi=247189&sr=HEADLINE(%22Jordan's%20substitution%20pattern%20comes%20into%20question.%22)%20and%20date%20is%202010

Gels, James (2016) "Basketball Coaching…Game Strategy" Accessed online, 13 December 2016. http://www.coachesclipboard.net/GameCoaching.html

Goldsberry, Kirk (2014) "Databall." Accessed online, 06 February 2014. http://grantland.com/features/expected-value-possession-nba-analytics/

Goldsberry, Kirk (2014) "Behind DataBall: A Discussion on the Methodology of Expected Possession Value." Accessed online, 07 February 2014. http://grantland.com/the-triangle/behind-databall-a-discussion-on-the-methodology-of-expected-points-value/

Gómez, Miguel-Ángel, Roberto Silva, Alberto Lorenzo, Rasa Kreivyte & Jaime Sampaio (2016) "Exploring the effects of substituting basketball players in high-level teams." Journal of Sports Sciences, DOI: 10.1080/02640414.2016.1161217.

Gosavi, A. (2003) Simulation-Based Optimization: Parametric Optimization Techniques and Reinforcement Learning. Kluwer Academic Publishers.

Kirkwood, Craig W. (1997) Strategic Decision Making: Multiobjective Decision Analysis with Spreadsheets. Belmont, CA: Duxbury Press.

Kloppenburg, Bob, Tom Newell, and Ernie Woods (2016) "Making Sound Substitutions." Accessed online, 13 December 2016. http://hooptactics.net/premium/coach/gametime/substituting.php

Kuehn, Joseph (2016) "Accounting for Complementary Skill Sets When Evaluating NBA Players' Values to a Specific Team." MIT Sloan Sports Analytics Conference Paper.

Mah, Cheri, Kenneth Mah, Eric Kezirian, and William Dement (2011) "The Effects of Sleep Extension on the Athletic Performance of Collegiate Basketball Players." Sleep, 34:7, 943-950. DOI: 10.5665/SLEEP.1132.

Mohr, Magni, Peter Krustrup, and Jens Bangsbo (2003) "Match Performance of High-Standard Soccer Players with Special Reference to Development of Fatigue." Journal of Sports Sciences, 21:7, 519-528, DOI: 10.1080/0264041031000071182.

Mohr, Magni, Peter Krustrup, and Jens Bangsbo (2005) "Fatigue in Soccer: A Brief Review." Journal of Sports Sciences, 23:6, 593-599, DOI: 10.1080/02640410400021286.

Morey, C. (2005) "Constraints, Limitations, and Assumptions Guide." Defense Technical Information Center, 20-2005.

Myers, Brett R (2012) "A Proposed Decision Rule for the Timing of Soccer Substitutions." Journal for Quantitative Analysis in Sports, DOI: 10.1515/1559-0410.1349.

NBA Stuffer.com (2016) "Player Evaluation Metrics." Accessed online. http://www.nbastuffer.com/component/option,com_glossary/func,display/Itemid,90/catid,42/

Parnell, G. S. (2010) Decision Making in Systems Engineering and Management, Second Edition (eds G. S. Parnell, P. J. Driscoll and D. L. Henderson). Hoboken, NJ: John Wiley & Sons, Inc.

Powell, Warren B (2011) Approximate Dynamic Programming: Solving the Curses of Dimensionality. Second Edition. Hoboken, NJ: John Wiley & Sons, Inc.

Reilly, T (1997) "Energetics of High-Intensity Exercise (Soccer) with Particular Reference to Fatigue." Journal of Sports Sciences, 15:3, 257-263, DOI: 10.1080/026404197367263.

Riches, (2016) "Four Ways to think about Substitutions Strategically." Accessed online, 13 December 2016. http://functionalbasketballcoaching.com/four-ways-to-think-about-substitutions-strategically/

Sackrowitz, Harold (2004) "Dynamic Programming and Time Related Strategies in Sports with an Application to Onside Kicks." Economics, Management and Optimization in Sports, 31-41, DOI: 10.1007/978-3-540-24734-0_3.

Sartini, Ken (2009) "Do You Sub to Keep your Players Fresh and Out of Foul Trouble?" Accessed online, 06 December 2016. https://www.breakthroughbasketball.com/coaching/game-strategy-subbing.html

Song, Alex, Thomas Severini, Ravi Allada (2016) "How Jet Lag Impairs Major League Baseball Performance." Proceedings of the National Academy of Sciences of the United States of America, 114:6, 1407-1412, DOI: 10.1073/pnas.1608847114.

Wilkes, Glenn (2016) "Mastering the Art of Substitutions." Accessed online, 06 December 2016. http://www.basketballsbest.com/articles/Art-ArtofSubstitutions.htm

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

BIOGRAPHY

David W. Hughes graduated from Baltimore Lutheran High School, Towson, Maryland, in 1997. He received his Bachelor of Science from the United States Military Academy in 2001. He flew CH-47D Chinook helicopters for the U.S. Army from 2002-2007, deploying twice (Iraq in 2003-2004, Afghanistan in 2005). He then commanded at Fort Riley, KS before obtaining his Master of Science from Stanford University in 2009. After a three year teaching assignment in the Department of Systems Engineering at the United States Military Academy, he was assigned to the Center for Army Analysis (CAA) at Fort Belvoir, VA for two years. During that assignment, he was deployed overseas to do Operations Research/Systems Analyst (ORSA) work for NATO Special Operations Component Command – Afghanistan (NSOCC-A). Shortly after that, he was selected to return to West Point in the summer of 2017 to teach in the Department of Systems Engineering, following the completion of his PhD at George Mason University. Dave is married to the former Meghan White of Cockeysville, Maryland and has three children: Alyssa, JD, and Kelly.