PROGRAM

September 26, 2015

Harvard University
Science Center, Lecture Halls C and D
1 Oxford Street
Cambridge, Massachusetts 02138

2015 New England Symposium on Statistics in Sports

Breakfast and Registration: Foyer area

8:00am – 9:15am: Coffee, tea, pastries

Welcome Address

9:15am – 9:30am: Mark Glickman and Scott Evans

Morning Session: Lecture Hall C

9:30am – 10:00am: Stan A. Veuger, American Enterprise Institute
"Football under pressure: Assessing malfeasance in Deflategate"

10:00am – 10:30am: Carl Morris, Harvard University
"Comparing sports performances using counts and rates"

10:30am – 11:00am: Hannah Frick, University College, London
"Monitoring fatigue – How do physical status, wellness and training load relate?"

Break: Foyer area

11:00am – 11:30am: Coffee and tea

Late-morning Parallel Sessions

11:30am – 1:00pm: Lecture Halls C and D

Lecture Hall C

11:30am – 12:00pm: Dan Cervone, New York University
"NBA court realty"

12:00pm – 12:30pm: Jeremias Engelmann, Consultant to ESPN
"Estimating a player's influence on his teammates' boxscore statistics using a modified RAPM framework"

12:30pm – 1:00pm: Weihua An, Indiana University
"Analysis of the match networks of NBA teams"

Lecture Hall D

11:30am – 12:00pm: Sam Robertson, Victoria University and Western Bulldogs Football Club
"A method to assess the influence of individual player performance distribution on match outcome in team sports"

12:00pm – 12:30pm: Julien Guyon, Bloomberg L.P. and Columbia University
"Rethinking the FIFA World Cup final draw"

12:30pm – 1:00pm: Pete Philipson, Northumbria University
"A comparison of sporting heroes: Bayesian modeling of test match cricketers"

Lunch break: Foyer area

1:00pm – 2:00pm: Sandwiches, beverages, snacks

Afternoon Parallel Sessions

2:00pm – 3:30pm: Lecture Halls C and D

Lecture Hall C

2:00pm – 2:30pm: Michael Lopez, Skidmore College
"Predicting coin flips: Using resampling and hierarchical models to help untangle the NHL's shootout"

2:30pm – 3:00pm: Andrew Hoegh, Virginia Tech
"Nearest-neighbor matchup effects: Predicting March Madness"

3:00pm – 3:30pm: Mark E. Glickman, Boston University
"A stochastic rank ordered logit model for rating multi-competitor games and sports"

Lecture Hall D

2:00pm – 2:15pm: Sameer Deshpande, The Wharton School of the University of Pennsylvania
"A Bayesian hierarchical model to estimate the framing ability of major league baseball catchers"

2:15pm – 2:30pm: Justin Manjourides, Northeastern University
"Developing advances (or any) statistics for beep baseball, an adaptive form of baseball for blind and visually impaired athletes"

2:30pm – 2:45pm: Stephanie Kovalchik, RAND Corporation
"Is there a Pythagorean theorem for winning in tennis?"

2:45pm – 3:00pm: Frank Silva, University of North Carolina
"Contextualized goals in soccer"

3:00pm – 3:15pm: Shinwoo Kang, Seattle Humane Society
"Triangle partitioning and linear optimization of forward lines"

3:15pm – 3:30pm: Steven Mintz, Texas A&M University
"Exploring the effectiveness of NBA play types with synergy possession data"

Poster Session: Foyer area

3:30pm – 5:00pm: With snacks and beverages

Panel Discussion: Lecture Hall C

5:00pm – 6:30pm: *"Hockey Analytics in the NHL*

Moderator: Michael Schuckers – St. Lawrence University

Panelist: Brian Macdonald – Florida Panthers
Panelist: Michael Peterson – Tampa Bay Lightning
Panelist: Andrew C. Thomas – University of Florida, Consultant to the Minnesota Wild

Post-NESSIS Get-Together

7:00pm – 9:00pm: John Harvard's Brewery & Ale House
Appetizers provided; Cash bar only.

Directions to *John Harvard's Brewery & Ale House*

Walking directions: Exit the Harvard Science Center by the doors near Lecture Halls C and D (or slightly worse at the Oxford Street exit as shown in the map on the right), and go through the gates into Harvard Yard. Walk straight through Harvard yard to one of the gate exits onto Mass Ave. Cross the street, and walk towards Dunster Street. John Harvard's Brewery & Ale House is on the right, and less than one block down on Dunster Street.

John Harvard's Brewery
33 Dunster Street
Cambridge, MA 02138
(617) 868-3585
https://www.johnharvards.com/locations/cambridge-ma/

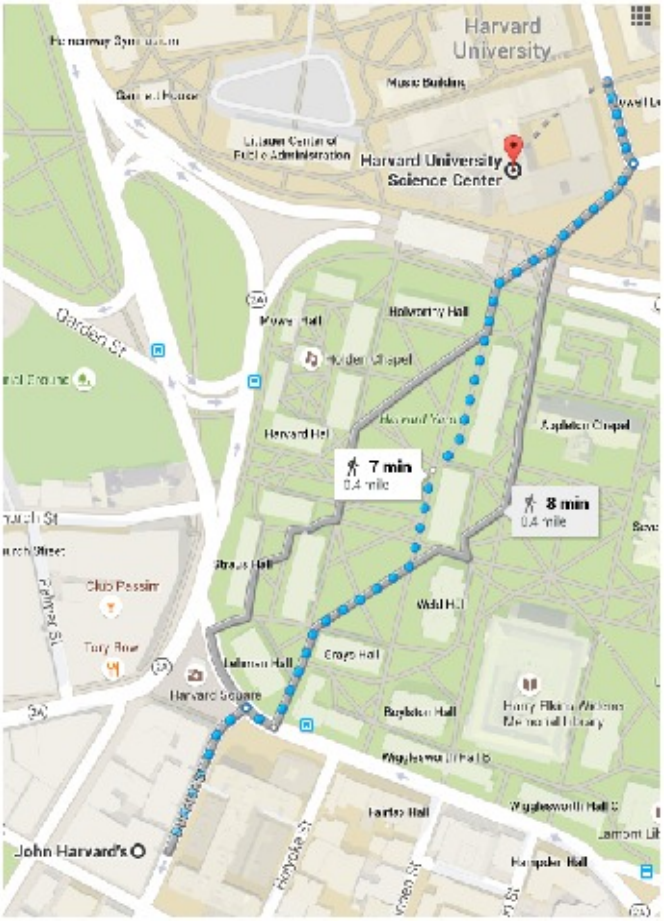

Symposium Co-Chairs:

Mark E. Glickman, Boston University

Scott R. Evans, Harvard University

Sponsors:

- Boston Chapter of the American Statistical Association (`http://www.amstat.org/chapters/boston/`)
- Section on Statistics in Sports of the American Statistical Association (`http://www.amstat.org/sections/sis/`)
- Harvard University Department of Statistics (`http://www.stat.harvard.edu/`)
- DeGruyter (`http://www.degruyter.com/`)
- IOS Press (`http://www.iospress.nl/`)
- ESPN Stats & Info (`http://espn.go.com/blog/statsinfo/`)
- DataRobot (`http://www.datarobot.com/`)
- The NBA (`http://stats.nba.com/`)

Wifi access:
Guest wifi access will be available for conference participants in the Science Center. Please follow these instructions:

- With your laptop or mobile device's wifi turned on, open a browser and choose "Harvard University" Wireless Network (SSID). Do not choose "Guest."
- Open a browser and go to any public site (e.g., `google.com`). You will be redirected to a network registration portal. The process will check your device. *Make sure you do not have pop-ups blocked.*
- You should receive an option to agree to rules and responsibilities. After agreeing, the connection process will be completed. You may be prompted to reboot your device/machine at this point, after which you will have guest access for the rest of the day.

Acknowledgments: We wish to thank everyone who helped to make NESSIS possible. We would also like to thank Kevin Rader, Sowmya Rao, Tom Lane, Dick Evans, Iram Farooq, Betsey Cogswell, Madeleine Straubel, Maureen Stanton, Jeffrey Myers, and Lisa Mukherjee for their parts in helping with the symposium.

Oral Presentation Abstracts

ANALYSIS OF THE MATCH NETWORKS OF NBA TEAMS

An, Weihua[†]

Indiana University, Bloomington, IN, USA

[†] E-mail: *weihuaan@indiana.edu*

In this study, I demonstrate how to use the latest advances in social network analysis to analyze NBA match records in order to better understand the NBA teams' performance and interactions. Based on the match records, I construct a network for each of the 2006-2015 NBA seasons. In each network a tie indicates a win-lose relationship between each pair of teams in a given season. I first examine the evolution of the networks across time, with a focus on significant structural changes in transitivity, core-periphery, etc. Then I use network centrality measures (e.g., indegree and eigenvector centrality) to predict the teams' standings in playoffs. Last, I provide in-depth analyses of the latest match records in 2014-2015 season through newly developed exponential random graph models (ERGMs). The ERGMs model a series of covariate effects on a team's winning odds, including not only the team's play style (e.g., average number of three-points, rebounds, and passes), but also its social characteristics (e.g., age, salary, number of all-stars, franchise legacy, and city size). Probably more importantly, the ERGMs also account for endogenous network formation processes such as transitivity (A beats B, B beats C, and so A is more likely to beat C) and "winners-take-all" (A beats many teams and so A is more likely to beat others). Hence, the ERGMs help to separate the two distinct mechanisms and provide more robust estimates of their contributions to team performance.

NBA COURT REALTY

Cervone, Dan[†]

New York University, New York, NY, USA

[†] E-mail: *dcervone@gmail.com*

During every instant of a basketball possession, offensive and defensive players occupy and control different regions of the basketball court. As the possession unfolds, player movements can be collectively viewed as real estate transactions. Offensive players exchange their previously occupied court space for new areas, presumably in order to increase the combined value of the court space they occupy. Defensive players behave analogously. Combined with the possession outcome, such court space transactions allow us to infer the price or value of any point on the basketball court from NBA optical tracking data.

Quantifying the NBA court real estate market enables new insights and metrics for both offense and defense. For instance, we can measure players' off-ball impact on offense by calculating the value of the space freed up for their teammates to control. For analyzing defense, we can quantify how effectively different teams (and different lineups within teams) contain the offense within low-value regions of the court.

A BAYESIAN HIERARCHICAL MODEL TO ESTIMATE THE FRAMING ABILITY OF MAJOR LEAGUE BASEBALL CATCHERS

Deshpande, Sameer K.[†]; Wyner, Abraham

The Wharton School of the University of Pennsylvania, Philadelphia PA, USA

[†] E-mail: *dsameer@wharton.upenn.edu*

Since its widespread adoption in 2008, Pitch FX technology has provided an enormous amount of high-resolution data on pitch location, trajectory, and velocity. Recently, there has been growing interest in quantifying a catcher's ability to "frame" a pitch (i.e. increase the chance that a pitch is called as a strike) and, more broadly, to determine how much credit (or blame) the pitcher, catcher, batter, and umpire deserve for each call. We introduce a Bayesian hierarchical model to estimate the probability of a called strike, adjusting for the pitch participants, pitch location, and contextual information like the count. This enables us to measure a catcher's ability to frame pitches and to identify systematic biases in umpires' calls. We translate our estimated effects into runs saved or gained across a season using a novel run expectancy framework.

ESTIMATING A PLAYERS' INFLUENCE ON HIS TEAMMATES' BOXSCORE STATISTICS USING A MODIFIED RAPM FRAMEWORK

Engelmann, Jeremias[†]

Consultant to ESPN.com

[†] E-mail: *jeremias.engelmann@gmail.com*

'Adjusted Plus Minus' and its' variants are useful tools to estimate a player's impact on his and the opponent teams' points per possession, effective FG% (eFG%), rebounding, turnovers and Free Throw rate. What these tools can't tell us is whether a player raises (or hurts) his team's e.g. eFG% through his own shot-making ability, or by making teammates shoot better (or worse). I introduce a method that, through changes in the regression design matrix and further refinement of existing Bayesian techniques, will allow us to estimate a player's impact on his teammates' BoxScore

statistics. This will further allow us to answer whether a player is "stealing rebounds" away from his teammates, whether (and how much) "stars" make it easier for their teammates to hit their shots, and to better predict a player's BoxScore statistics when he switches teams.

MONITORING FATIGUE – HOW DO PHYSICAL STATUS, WELLNESS AND TRAINING LOAD RELATE?

Frick, Hannah†; Kosmidis, Ioannis

University College, London, United Kingdom

† E-mail: *h.frick@ucl.ac.uk*

This talk aims to capture fatigue in high-performance athletes through physical status and wellness monitoring and relate fatigue to training load. The rich amount of data gathered by GB Hockey on their athletes covers on-pitch training load assessed via GPS tracking and heart rate monitors, physical status assessed via fitness and strength tests as well as self-rated wellness. Fatigue indicators such as motivation, feeling of soreness and freshness, working and resting heart rates, drop jump measure as well as power production and absorption are analysed via log-linear mixed models to describe their pairwise association structure. The on-pitch training load is summarised in training distribution profiles (Kosmidis & Passfield, 2015) for speed and heart rate. Rather than summarising a training session in a set of scalar quantities (e.g., quantiles), a distribution profile captures the time spent training above a certain (speed/heart rate) threshold in a function of speed and heart rate, respectively. A functional principal components analysis of the heart rate profiles suggests that the profiles can be characterised in terms of three factors including duration and intensity of the training session. This talk will also be looking at modelling various aspects of fatigue through heart rate and speed profiles in functional regression models.

A STOCHASTIC RANK ORDERED LOGIT MODEL FOR RATING MULTI-COMPETITOR GAMES AND SPORTS

Glickman, Mark†[1]; Hennessy, Jonathan[2]

[1]*Boston University, Boston, MA, USA;* [2]*Houston Rockets, Houston, TX, USA*

† E-mail: *mg@bu.edu*

Many games and sports, including races, involve outcomes in which competitors are rank ordered. In some sports, competitors may play in multiple events over long periods of time, and it is natural to assume that their abilities change over time. We propose a Bayesian state-space framework for rank ordered logit models to rate competitor abilities over time from the results of multi-competitor games. Our approach assumes competitors' performances follow independent extreme

value distributions, with each competitor's ability evolving over time as a Gaussian random walk. The model accounts for the possibility of ties, an occurrence that is not atypical in races in which some of the competitors may not finish and therefore tie for last place. Inference can be performed through Markov chain Monte Carlo (MCMC) simulation from the posterior distribution. We also develop a filtering algorithm that is an approximation to the full Bayesian computations. The approximate Bayesian filter can be used for updating competitor abilities on an ongoing basis. We demonstrate our approach to measuring abilities of 268 women from the results of women's Alpine downhill skiing competitions recorded over the period 2002-2013.

RETHINKING THE FIFA WORLD CUP FINAL DRAW

Guyon, Julien[†1,2]

[1] *Bloomberg L.P., New York, NY, USA;* [2] *Columbia University, New York, NY, USA*

[†] E-mail: *julien.guyon.1977@gmail.com*

The soccer World Cup is the most popular sporting event in the world, even more widely viewed and followed than the Olympic Games. In this talk, we critically examine a number of flaws in the current procedure for the final draw of this tournament: imbalance (the eight groups are not of the same competitive level), unfairness (some teams have a greater chance than others of ending up in a tough group), and uneven distribution (all the possible outcomes of the draw are not equally likely). These flaws result from the way FIFA, the sport's world governing body, has decided to enforce the geographic constraints that they put on the draw. We explain how, by building eight pots by level organized in an S-curve, and drawing first a continental distribution of the groups and then the teams, we can enforce the geographic constraints without sacrificing balance, fairness, and even distribution. As a result, we describe a new tractable draw procedure that produces eight balanced and geographically diverse groups, is fair to all teams, and gives equally likely outcomes. Monte Carlo simulations demonstrate the superiority of the proposed method over the current system in terms of balance, fairness, and even distribution.

NEAREST-NEIGHBOR MATCHUP EFFECTS: PREDICTING MARCH MADNESS

Hoegh, Andrew[†]; Carzolio, Marcos; Crandell, Ian; Hu, Xinran; Roberts, Lucas; Song, Yuhyun; Leman, Scotland C

Virginia Tech, Blacksburg, VA, USA

[†] E-mail: *ahoegh@vt.edu*

NCAA tournament prediction competitions have long been a rite of passage for spring. Recently, variations on traditional bracket formats have increased the use of data-driven predictive models.

Most predictive models are based entirely on measures of overall team strength, resulting in the unintended "transitive property." These models are therefore unable to capture specific matchup tendencies. We introduce our novel nearest-neighbor matchup effects framework, which presents a flexible way to account for team characteristics above and beyond team strength that may influence game outcomes. In particular we develop a general framework that couples a model predicting a point spread with a clustering procedure that borrows strength from games similar to a current matchup. This results in a model capable of issuing predictions controlling for team strength and that capture specific matchup characteristics.

TRIANGLE PARTITIONING AND LINEAR OPTIMIZATION OF FORWARD LINES

Kang, Shinwoo[†]

Seattle Humane Society, Bellevue, WA, USA

[†] E-mail: *charlestreykang@gmail.com*

One important objective of an ice hockey coach is to produce optimal sets of four forward lines and three defensive pairings. A popular method of choice for evaluating chemistry is the "With or Without You" (WOWY) analysis on two or more players. Given enough data for each combination of players, determining the optimal defensive pairings is simple with WOWY analysis, as only 20 different combinations of pairings need to be compared. However, when optimizing forward lines, the WOWY analysis becomes a difficult task. Because many players do not play together frequently, the standard errors on any sets of three players become too large to distinguish.

This paper proposes a rigorous method for optimizing line combinations by representing each forward line as a triangle, where the vertices are players, and edges describe the binary relationships between player-pairings according to some evaluation metric. The evaluation metric can be specified by the user (e.g. Corsi For Percentage, Fenwick Against per 60, etc), such that the edge weights are, e.g., CF% when the pair of players are both on the ice. Then, partitioning 12 forwards into 4 vertex-disjoint triangles produces 369,600 unique configurations. By performing linear optimization on each configuration of triangles with predetermined boundaries on Time on Ice, one can compute the forward line combinations that produce the highest feasible CF% for the team. The method of triangle partitioning and linear optimization has many possibilities in improving on-ice coaching decisions, such as determining optimal match up scenarios.

IS THERE A PYTHAGOREAN THEOREM FOR WINNING IN TENNIS?

Kovalchik, Stephanie[†]

RAND Corporation, Santa Monica, CA, USA

[†] E-mail: *skovalch@rand.org*

Bill James' discovery of a Pythagorean formula for win expectation in baseball has been a useful resource to analysts and coaches for over 30 years. Extensions of the Pythagorean model have been developed for all of the major professional team sports but none of the individual individual sports. The present paper attempts to address this gap by deriving a Pythagorean model for win production in tennis. Using performance data for the top 100 male singles players between 2004 and 2014, this study shows that, among the most commonly reported performance statistics, a model of breakpoints won provides the closest approximation to the Pythagorean formula, explaining 80% of variation in season win percentage with an error of +/- 2.5 matches for a 50-match season. The inclusion of relative strengths in total points won and tiebreak points won each had a modest improvement in the performance of the model. The relatively simple formulas relating these performance metrics to tennis win production have a wide range of potential uses for player performance evaluation and match forecasting.

PREDICTING COIN FLIPS: USING RESAMPLING AND HIERARCHICAL MODELS TO HELP UNTANGLE THE NHL'S SHOOTOUT

Lopez, Michael[†1]; Schuckers, Michael[2]

[1]*Skidmore College, Saratoga Springs, NY, USA;* [2]*St. Lawrence University, Canton, NY, USA*

[†] E-mail: *mlopez1@skidmore.edu*

Roughly 15% of regular season National Hockey League games are decided by a shootout, and the resulting allocation of points impacts playoff races each season. But despite interest from fans, players, and league officials, little is known about the variability in shootout skill within either shooters or goaltenders. We attempt to fill that void, presenting both generalized linear mixed model and Bayesian hierarchical model frameworks to model shootout outcomes. Results suggest (i) small but significant talent gaps between shooters, (ii) marginal differences in performance among netminders, and (iii) that there are few, if any, notable predictors of player success after accounting for individual talent. We also provide a resampling strategy to highlight a selection bias with respect to shooter assignment, in which coaches choose their most skilled offensive players early in shootout rounds and are less likely to select players with poor past performances. Finally, given that per-shot data for shootouts does not currently exist in a single location for public use, we provide both our data and source code for other researchers interested in studying shootout outcomes.

DEVELOPING ADVANCES (OR ANY) STATISTICS FOR BEEP BASEBALL, AN ADAPTIVE FORM OF BASEBALL FOR BLIND AND VISUALLY IMPAIRED ATHLETES

Manjourides, Justin[†]

Northeastern University, Boston, MA, USA

[†] E-mail: *j.manjourides@neu.edu*

Beepball is an adaptive form of baseball for blind and visually impaired athletes, played by over 35 teams across the United States. Despite more than a 30-year history, little has been done to advance the collected and reported statistics beyond batting averages for hitters, strikeout percentages for pitchers, and number of outs made for fielders. These statistics are reported and ranked within a yearly double-elimination tournament, further biasing these measures of performance, as better teams tend to play better teams as the tournament progresses, and weather plays a big factor in the outcome of games.

Consider a form of baseball where the pitcher is on the same team as the batter, meaning strikeouts are negative outcomes for both players. A hit equals a run, meaning there is no base running, beyond successfully reaching First Base. Defensively, an out is made if the fielding player successfully gains control of a ball in play, prior to the runner reaching First Base. These basic differences make the adaptation of Sabermetrics to beepball challenging.

In this talk I will introduce the sport, highlighting differences between beepball and baseball. Current player, pitcher, and team statistics will be evaluated. I will propose advanced metrics to better describe and evaluate player performance. Few current advanced statistics can be used directly, however concepts such as "defense independence" and "field adjustments" can be applied and should help determine who the best players, pitchers, and teams are in a given tournament, and help compare performance year-to-year.

EXPLORING THE EFFECTIVENESS OF NBA PLAY TYPES WITH SYNERGY POSSESSION DATA

Mintz, Steven[†]

Texas A&M University, College Station, TX, USA

[†] E-mail: *Mintzsj@gmail.com*

Basketball data, to this point, primarily capture end results of plays. Historical records of rebounds or fouls explain how possessions may end or change. However they do little to characterize the microcosmic events that eventually lead to these outcomes. Furthermore, many subsequent models or statistics deemed advanced are reliant on these traditional markers in time as their basis

of construction. Synergy Sports Technology implements a proprietary possession-tagging system that provides a granular level of data. Typically, these player tendencies or decisions leading to possession-finality are significantly reduced. There are many more pieces to a possession's story that Synergy wonderfully tracks. Rather than simply accepting Al Jefferson's career scoring average is 17.0 ppg, we can contextualize his dominance on the right block in the post, turning over his right shoulder.

This paper treats a collection of a game's possessions in the form of a time series. Clustering techniques are also implemented to understand the associative properties of similar play types. Results identify the, particularly effective, play types as well as contributions of players during possessions. By analyzing the meta-events of certain possessions, we form a richer comprehension of team play. Additionally, overlaying this work with more traditional statistics and tracking data illuminates underlying tendencies within particular possessions.

COMPARING SPORTS PERFORMANCES USING COUNTS AND RATES

Morris, Carl N.[†]

Harvard University, Cambridge, MA, USA

[†] E-mail: *morris@stat.harvard.edu*

Talent evaluations that only use total successes but ignore the player's total opportunities are simple to understand. Fans like them also because totals provide entertainment value when players pursue previous record totals or when their totals approach special round numbers. However, by failing to account for opportunities, e.g. via rates that relate successes to opportunities, totals used alone can distort evaluations and awards. Totals always increase, even when rates of success fall to substandard levels, and that can encourage a leading athlete to play even when playing poorly.

A case of considerable current media interest is Pete Rose who, barring his gambling issues that have prevented his consideration for baseball's Hall-of-Fame (HOF), would be a shoo-in for that honor. That's especially due to his holding MLB's career hit record. By continuing to play while a substandard hitter, Rose in 1984 managed to Ty Cobb's once seemingly untouchable record of 4189 career hits at age 45. Those underperforming years could drag Rose's overall career batting record down to a marginal level for HOF entry.

We review these issues in baseball by considering better statistical measures to improve evaluations that, among other differences, account for the different number of opportunities a player has. These measures include runs per game formulas, Markov models, and hierarchical models that address regression-toward-the-mean issues.

A COMPARISON OF SPORTING HEROES: BAYESIAN MODELLING OF TEST MATCH CRICKETERS

Philipson, Pete[†]

Northumbria University, Newcastle upon Tyne, UK

[†] E-mail: *pete.philipson@northumbria.ac.uk*

In this work the contentious problem, and eternal bar room debate, of comparing sportsmen whose careers took place in different eras is addressed. The application here is to Test match cricket, encompassing both batsmen and bowlers using data from the first Test in 1877 onwards. Direct comparisons are compromised by changes to the game itself over time, whether this is due to an expanding talent pool, fundamental changes to rules and equipment or other factors. The overlapping nature of careers is exploited to form a bridge from past to present. The overall aim is to compare all players simultaneously, rather than just relative to their contemporaries.

An additive log-linear model that incorporates year-, age- and game-specific components is used to allow the innate ability of an individual to be identified. A Bayesian approach is adopted and the posterior distribution for model parameters is determined by using Markov Chain Monte Carlo (MCMC) methods with random walks. We use this posterior distribution to construct a table of leading batsmen and bowlers via their predictive distributions.

A METHOD TO ASSESS THE INFLUENCE OF INDIVIDUAL PLAYER PERFORMANCE DISTRIBUTION ON MATCH OUTCOME IN TEAM SPORTS

Robertson, Sam[†1,2]; Gupta, Ritu[3]; McIntosh, Sam[1,2]

[1] *Victoria University, Melbourne, Australia;* [2] *Western Bulldogs Football Club, Melbourne, Australia;* [3] *Curtin University, Perth, Australia*

[†] E-mail: *sam.robertson@vu.edu.au*

This study proposes a method to determine whether the distribution of individual player performances can be modelled to explain match outcome in team sports. Player-recorded values (converted to a percentage of team total) in 11 commonly-reported performance indicators were obtained for all regular season matches played during the 2014 Australian Football League season, with team totals also recorded. Multiple features relating to heuristically determined percentiles for each performance indicator were then extracted for each team and match, along with the outcome (Win/Loss). A generalised estimating equations model comprising eight key features was developed, explaining match outcome at a median accuracy of 63.9% accuracy under 10-fold cross validation. Lower 75th, 90th and 95th percentile player contributions to team goal totals and higher 25th and

50th contributions to team disposals were the features most strongly linked with a positive match outcome. The model provides evidence in support of strategies which aim to obtain an even spread of team goal kickers in order to facilitate winning a match. The analysis methodology used in this investigation can be used to quantify the importance of individual contributions to overall team performance in a range of team sports.

CONTEXTUALIZED GOALS IN SOCCER

Silva, Frank[†]

University of North Carolina – Chapel Hill, Chapel Hill, NC, USA

[†] E-mail: *frank.silva@unc.edu*

Work has been done in soccer and other sports to evaluate the goals/points a player has scored beyond simply counting the total amount scored. As Stephen Pettigrew has recently argued, points and goals are not created equal. In soccer, for instance, a goal scored at the 90th minute or later of a tied game is more critical than a goal scored at the same time in a 4-0 game. In this research, I introduce a method to value goals that takes into account the context of the game, which includes the time remaining, number of players on the field, home/away information, gambling data, and the difficulty of the shot. One feature of this method is a win-probability model, which I use to assign goals' values according to how much they increase the probability of a team winning a game. I also incorporate Expected Goal Value data to provide more context to goals based on the degree of difficulty of the goal scored. My research adds a new level of context to similar work in two particular ways. First, by incorporating gambling data, I can take into account team skill levels. Second, by including Expected Goal Value data, I can better distinguish players who score easy shots from those who score more difficult ones. This analysis can help teams, journalists, and the general public make effective comparisons between players who have scored a similar number of total goals in a season.

FOOTBALL UNDER PRESSURE: ASSESSING MALFEASANCE IN DEFLATEGATE

Veuger, Stan A.[†]

American Enterprise Institute, Washington, DC, USA

[†] E-mail: *Stan.Veuger@aei.org*

After the AFC Championship game between the New England Patriots and the Indianapolis Colts, the Patriots were accused of deflating their footballs to gain an unfair advantage over their opponents. A subsequent investigation by the NFL led to the publication of a report named after Ted Wells, its main author. The Wells report's central conclusion was that the Patriots and their

quarterback, Tom Brady, were at least generally aware of what was deemed to be probably illicit behavior by some of the Patriots employees responsible for football preparation. Team and player were then penalized with fines, draft pick losses, and suspensions by NFL Commissioner Roger Goodell.

The Wells report rapidly became the subject of a wide range of critical responses from the Patriots organization as well as disinterested third parties from a variety of disciplinary backgrounds. You are looking an example of exactly that. This presentation focus on the statistical analysis in the report and finds fault with the set of hypotheses it tests, the way in which it tests them, the robustness of its test results, and the conclusions it draws from its tests. It also highlights problems with the quality of the data used in the report, and sketches more appropriate interpretations of the evidence presented to the NFL.

Poster Presentation Abstracts

A MIXTURE MODELING APPROACH TO BASEBALL PERFORMANCE PREDICTION

Andelin, Kyle; Kaplan, Sam; Hutchinson, Brian†

Western Washington University, Bellingham, WA, USA

† E-mail: *Brian.Hutchinson@wwu.edu*

There is clear evidence that a player's true ability varies over long periods of time (e.g. several seasons), but also evidence that it does on a shorter time scale (e.g. weeks). This evolution of ability complicates the already challenging task of extrapolating from past to future performance. In this work we take a mixture modeling approach to projecting performance, and use the interpolation weights to gain insight into the trade-off between robustness (larger samples) and recency (relevant sample). We estimate discrete distributions over per-game offensive statistics for each of: strikeouts, walks, singles, doubles, triples and home runs. Specifically, for each player and statistic, we estimate three distributions: one trained on games from the prior year ("Prior"), one on games from the first K weeks of the current season ("Early") and one trained on games from the most recent $(T - K)$ weeks of the current season ("Recent"). We find the optimal three-way interpolation per statistic so that each player's three previous time periods best predicts his performance in the next M weeks. Using data from the 2000-2013 seasons, with $T = 17$, we find an interesting observation: for large M, there is no significant difference in the mixture weights between Early and Recent models of the same length, but when $M = 1$, the Recent model has significantly higher mixture weights than Early. This results provides some new evidence for the "hot hand" phenomenon; i.e. that recent performance is more predictive than earlier performance.

IS PATIENCE A VIRTUE?: STUDYING THE EFFECTS OF PITCHES PER PLATE APPEARANCE IN THE MLB

Balton, Tim; Maher, Tyler†; Rose, Brian

Tufts University, Medford, MA, USA

† E-mail: *tylermaher@comcast.net*

This study investigated the role that pitches per plate appearance (P/PA) plays in the offensive performance of major league baseball players, with the overarching goal of determining whether patience at the plate increases offense. Several studies have tackled this subject recently with mixed results, generally finding that P/PA has little to no effect on team and player success. Therefore, our hypothesis was that although announcers, coaches, and general managers advocate waiting for one's pitch, there is not much of a relationship between pitches per plate appearance and offense because the benefits and drawbacks of taking pitches essentially cancel each other out. For example, P/PA would be expected to correlate with positive outcomes like walks and home runs, but would also likely lead to strikeouts and have an inverse relationship with batting average. Using league and team data going back to 1988, we examined similar measures of offensive performance one would expect to correlate with P/PA in the hopes of determining the extent to which P/PA influences offense. Our preliminary findings were that there is a negligible relationship between P/PA and offense, thus confirming our hypothesis that the perceived value of patience far outweighs its actual value. In fact, recent data suggests that patience may actually be a detriment to offense, suggesting more data and further investigation is needed.

MAN IN THE MIDDLE: OPTIMAL DEFENSIVE STRATEGIES AND DISRUPTING PASSING NETWORKS IN SOCCER

Bojinov, Iavor†; Bornn, Luke

Harvard University, Cambridge, MA, USA;

† E-mail: *bojinov@fas.harvard.edu*

Soccer, the most watched sport in the world, is a dynamic game where a team's success relies on both team strategy and individual player contributions. Previous statistical studies have used basic summary statistics such as number of goals per minute, assists in a season, and average number of clean sheets. These statistics, however, can never completely summarize the characteristics and effectiveness of a particular playing strategy. Fortunately, with modern optical tracking techniques, companies like Prozone can collate high resolution player location data with low resolution events data.

This work utilizes the Prozone English Barclays Premier League season 2012/13 tracking data in order to develop statistical analysis for understanding and comparing different defensive strategies

in soccer. More specifically, we use a Bayesian hierarchical spatiotemporal model for the real time evolution of passing networks. Passing is a cardinal soccer skill and key factor in strategy development; it helps the team to keep the ball in their possession, move it across the field, and outmaneuver the opposing team in order to score a goal. Passing networks, therefore, provide a useful characterization of a team's tactics and performance. In this work, we discover when and how teams, as well as individual players, disrupt these passing networks for the opposing team.

PLAYER DEVELOPMENT IN THE NBA

Chase, Harrison†; Ver Steeg, Nathaniel; Smith, Daniel

Harvard University, Cambridge, MA, USA

† E-mail: *hchase@college.harvard.edu*

Our study attempts to analyze some possible influences on player development in the NBA, focusing on postseason experience and veteran leadership. To do this we first created an aging curve using the delta method looking at BPM (Box Plus/Minus), a measure of player skill. We then attempted to control for survivor bias by creating an adjusted aging curve. We then regressed the difference between age-estimated progression and actual progression on a variety of factors, finding a significant relationship between this and postseason experience, after controlling for other factors. These other factors included a player's BPM in the first year, their minutes played in the first year, the draft pick number of a player, and interaction between these variables. We selected these variables using a stepwise procedure with AIC as the criterion. Veteran leadership, measured in several forms, was not significant. We performed a cross-validation procedure to compare the selected model against various others selected using different criteria (for example, BIC), and the one selected using AIC performed the best. Finally, we looked at the residuals for both teams and coaches, to see who is associated with the most player development.

DATA MINING MAJOR LEAGUE BASEBALL'S PACE OF PLAY PROBLEM

Crowley, Aaron†; He, Zhuolin; Blair, RH

State University of New York at Buffalo, Buffalo, NY, USA

† E-mail: *acc34@buffalo.edu*

Longer Major League Baseball (MLB) games and declining popularity have recently been a hot topic in sports journalism. MLB addressed this in the 2015 season by implementing new rules that quicken the pace of play by limiting the amount of time between each half-inning and pitching change. We assess the problem empirically by analyzing all nine inning regular season MLB games played between the 2000 and 2014 seasons. We perform feature selection and produce interpretable

models utilizing two popular data mining techniques, the lasso, and principal component regression in order to identify the components of baseball games that add the most time. The number of pitchers used and the amount of batters walked were found to prolong MLB games the most. We conclude our analysis by suggesting potential rule changes which are designed to significantly decrease the length of games and in hopes of increasing MLB popularity.

A LOGISTIC REGRESSION APPROACH TO PREDICTING WHO WILL MAKE THE NBA PLAYOFFS

Elmore, Ryan[†]

University of Denver, Denver, CO, USA

[†] E-mail: *rtelmore@gmail.com*

At any point in the NBA season, it is easy to find a prediction for the likelihood of any team making the playoffs, see e.g. John Hollinger's NBA Playoff Odds on espn.com. A major shortcoming of the various predictions by "experts" is that they rarely, if ever, provide a measure of uncertainty surrounding their results. In this talk, I will introduce a logistic regression model for predicting the probability of making the playoffs for each team in the NBA at any point in the season. In addition, this methodology provides yet another power ranking and perhaps, more importantly, we introduce a Monte Carlo-based estimate of the uncertainty in our predictions. The model is based on every NBA game since the 2002 season. This work was performed in consultation with Tommy Balcetis, Manager of Basketball Analytics at Denver Nuggets.

SAMPLING FROM THE 9,223,372,036,854,775,808 POSSIBLE BRACKETS IN THE NCAA MEN'S BASKETBALL TOURNAMENT USING THE POWER MODEL

Jacobson, Sheldon H.[†], Khatbi, Arash, King, Douglas M.

University of Illinois at Urbana-Champaign, Champaign, IL, USA

[†] E-mail: *shj@illinois.edu*

The National Collegiate Athletic Association's (NCAA) annual men's Division I college basketball championship tournament, commonly referred to as March Madness, draws significant national attention by both avid sports fans and the general population. A popular activity is attempting to pick the winners of all 63 games (excluding the so-called First Four). This presentation proposes the Power model for sampling from all the possible brackets (9,223,372,036,854,775,808 in total), in the correct proportion of how seeds typically perform in advancing through the tournament. The model uses historical data based on how seeds have propagated through the bracket, and estimates

the necessary model parameters based on such data. The results using this model are compared to the results for past tournaments, to assess how well the sampling is done. Practical implications of the Power model are discussed. Note that the model does not attempt to predict who will win the games, but rather, focuses on sampling across brackets based on historical results, with the goal of designing a basket of brackets that contain the appropriate mix of upsets and favorites advancing. The Power model is publicly available on the website, *bracketodds.cs.illinois.edu*.

SURVIVOR FUNCTIONS, HAZARD FUNCTIONS, AND MEDIAN LIFETIMES FOR OLYMPIC SPONSORSHIPS

Jensen, Jonathan A.[†]

Merrimack College, North Andover, MA, USA

[†] E-mail: *jensenjo@merrimack.edu*

This study features the first application of event history analysis (EHA; i.e., survival analysis) approaches to describe and illustrate the duration of sport marketing partnerships (i.e., sponsorships). The specific context utilized is one of the more recognized and influential sponsorship programs, the TOP (i.e., The Olympic Partners) sponsorship program, which began in 1985 and continues to this day. Various aspects of the EHA methodology, including hazard functions, survivor (and alternatively, failure) functions, and median lifetimes for these global sponsorships will be reviewed and depicted graphically. The study seeks to not only illustrate a novel application of the EHA methodology in a new context, but assist those who buy and sell global sponsorships in both forecasting and revenue projection efforts by increasing the accuracy with which the duration of such partnerships are analyzed.

RECURRENT NEURAL NETWORK MODELING OF BASEBALL PERFORMANCE

Kaplan, Sam; Andelin, Kyle; Hutchinson, Brian[†]

Western Washington University, Bellingham, WA, USA

[†] E-mail: *Brian.Hutchinson@wwu.edu*

Recurrent neural networks (RNNs) have recently demonstrated impressive success on a wide range of time series modeling tasks, in application areas from language to music. Here we consider their use in modeling batter performance over the course of a season: given the counts of a given offensive statistic (e.g. hits) in each game up to some point in the season, we predict the batter's performance (as a distribution over outcomes) in the next game. The offensive statistics we consider are strikeouts, walks, singles, doubles, triples, home runs and hits. Unlike a standard feed-forward

neural network, the RNN maintains a hidden state across time steps (by treating the previous hidden state as an additional input), allowing the model to adjust to short and long term trends in the batter's performance. For our experiments, we consider the 477 batters from the 2013 season who play 25 or more games. We evaluate absolute error between expected and actual performance one game at a time on the second half of the season for a set of held out players (having fed the RNN the first half), and find that our model outperforms two baselines (league average and season running average) on all statistics, by an average of 0.045 absolute (14.5% relative reduction in mean absolute error). Future extensions include incorporating key contextual variables, including opposing pitcher and stadium, and modeling at finer granularity (e.g. plate appearance).

ASSIST QUALITY: MEASURING THE TRUE VALUE OF AN ASSIST

Lee, Alexander[†]

University of Toronto, Toronto, Ontario, Canada

[†] E-mail: *Aleejsj@gmail.com*

On the offensive end, many of the statistics recorded are very straightforward: number of shots taken, number of shots made and now through Sports VU data we have precise measurements of shot distance. However, one offensive statistic that has been left in the past is the assist. One of the reasons is its subjective nature. The NBA defines: "an assist is awarded only if, in the judgment of the statistician, the last player's pass contributed directly to a made basket." The subjectivity of each home court's statistician is at the core to recording assists, and in fact, there have been several pieces that demonstrate the home-court effect on how assists are given out. But even if there was a perfectly uniform judgment system, the assist is still a binary statistic for something much more complex. Some assists lead to wide open layups while some lead to contested fadeaway three-pointers. There is clearly a difference in the difficulty of the shots, but based on the timing of the pass, both passes may be recorded as equal assists. We attempt to capture a quantitative metric for assessing assists by rewarding the quality of shots that result from them. Using player tracking and statistical modeling, we can uncover a deeper understanding of assists. Our method provides a way to sift through players who are truly the better passers in the league compared to players who are simply rewarded for having teammates who hit difficult shots.

THE AUTOMATED GENERAL MANAGER AN UNBIASED, BACKTESTED ALGORITHMIC SYSTEM FOR DRAFTS, TRADES, AND FREE AGENCY THAT OUTPERFORMS HUMAN FRONT OFFICES

Maymin, Philip Z.[†]

Vantage Sports, Seattle, WA, USA

[†] E-mail: *philip@maymin.com*

I introduce an automated system and interactive tools for NBA teams to better decide who to draft, who to trade for, and who to sign as free agents. This automated general manager can serve either as an expert-system replacement or a complement to a team's front office, and also as a calibrating benchmark to compare against actual team building performance. Backtested over the past ten years, the automated GM outperforms every single team, and by substantial margins that often represent a major portion of the team's market value. From draft decisions alone, the average team lost about $130,000,000 worth of on-court productivity relative to what they could have had with the automated GM; this shortfall represents about a quarter of the average franchise value. Historically the automated GM's choices would have produced about twice as much as the human choices actually did: approximately one extra win per year per draft pick. The system is calibrated using an innovative extension of traditional machine learning methods, applied to a uniquely broad historical database that incorporates both quantitative and qualitative evaluations, in a way that avoids possible survivorship bias, and for a variety of performance metrics; it is thus robust, comprehensive, realistic, and does not overfit information from the future. I provide virtually all of the interactive tools supporting this paper, including backtesting results, projections, scenario analysis, and more, online, for free, at nbagm.pm.

DEEP NEURAL NETWORKS FOR SPORTING EVENT OUTCOME PREDICTION

Moorhead, Jacob; Baldwin, Elijah; Hutchinson, Brian[†]

Western Washington University, Bellingham, WA, USA

[†] E-mail: *Brian.Hutchinson@wwu.edu*

Deep learning has enjoyed a strong resurgence over the past decade, due to advances in both algorithms and computational resources. Deep neural networks (DNNs), which have multiple hidden layers, provide state-of-the-art performance for many challenging tasks, including speech and image recognition. Here we use them to predict the outcome of baseball (MLB) and college football games. Part of the challenge of this task is the sparsity of the data: each team plays different opponents and under different circumstances (e.g. in different stadiums). Our approach takes as input the

identities of the two teams playing and predicts the difference in points scored by the two teams. Even with this very limited source of information, our models give relatively strong performance. For college football matches, the correct winner on a set of 84 held out games from 2013 is chosen 80.5% of the time (17.5 root mean squared error), beating random guessing (50% accuracy) and picking the team with the better season point differential (69.0%). For baseball, the correct winner of a set of 244 heldout games from 2013 is chosen 58.3% of the time, again beating random guessing and the season point differential baseline (54.5%). As a side effect, our model learns low dimensional representations of teams through a linear "bottleneck" layer, which is used to cluster similar teams. A logical extension to our work is to incorporate more contextual information (e.g. starting pitchers, stadium), which should further improve prediction performance.

COMPETITION FROM OTHER SPORTS HURTS TV RATINGS: HOW TO SHIFT LEAGUE CALENDARS TO OPTIMIZE VIEWERSHIP

Pagels, Jim[†]

Columbia University, New York, NY, USA

[†] E-mail: *jimpagels@gmail.com*

Television is becoming an increasingly critical revenue stream in the sports industry, as media rights deals in the five major North American sports (NBA, MLB, NHL, NFL, and college football) continue to escalate by huge rates every time they are up for renewal. However, games frequently overlap with each other on the calendar, and this competition for viewership often has substantial negative effects on ratings. This paper attempts to isolate the effects of overlap from each sport, examine how that competition hurts viewership in each league, and quantify the value lost due to that overlap. We find that competition can have very damaging effects for TV viewership for every sport, most notably the NHL, and these losses can significantly diminish the value of a network's investments in sports programming. In most cases, this overlap is entirely avoidable with some relatively unobtrusive season calendar shifts.

IS THE PRICE RIGHT? AN ANALYSIS OF NATIONAL FOOTBALL LEAGUE TICKET MARKETS RELATIVE TO THE GREAT RECESSION

Singer, Ben[†]

Columbia University, New York, NY, USA

[†] E-mail: *bs2831@columbia.edu*

The Great Recession (December 2007 through June 2009) in the United States not only incurred widespread economic turmoil, but also was one of the few times during the last decade when National Football League (NFL) attendance figures fell below almost full attendance. However, NFL ticket prices never decreased on average during this time. This study takes advantage of this natural experiment to analyze macroeconomic, team performance, and other factors that influence NFL ticket markets before, during, and after the recession. Data for 30 teams from 2004 through 2014 are used in a linear regression model that accounts both for team-based fixed effects and autocorrelation in the errors. The results show that the unemployment rate in a team's metropolitan area is a strong predictor of that team's attendance; that the mean annual wage in a team's metropolitan area and the opening of a new stadium are strongly predictive of the price of non-premium tickets; and that the relative ticket price increase associated with a new stadium is greater for non-premium tickets than for premium tickets. Additionally, after accounting for team fixed effects, there is no evidence that team performance impacts either ticket demand or team financial health. These findings provide insight into how the recession influenced NFL ticket demand and into how certain team and macroeconomic factors may predict future NFL ticket demand. Furthermore, this study lays the foundation for investigating to what degree NFL team ticket pricing strategies were economically rational during this period.

SURVIVAL IN THE NBA: AN ANALYSIS OF COLLEGIATE EXPERIENCE AND CAREER DURATION

Smith, Daniel J.[†]

New York University, New York, NY, USA

[†] E-mail: *dan.smith@nyu.edu*

The NBA's annual draft underscores a salient issue for amateur players and team front offices: the value of collegiate playing experience. One decision problem for a collegiate player is the number of years to commit to NCAA basketball that maximizes his time in the NBA. In examining the effect of collegiate experience on NBA career duration, I hypothesize that professional career length is increasing in collegiate career length.

This paper analyzes data from the 2000-2013 NBA drafts. The sample comprises all NCAA players entering those drafts ($n = 1189$). Survival analysis of time to first event is conducted using Kaplan-Meier estimation and accelerated failure-time Weibull regression. The analysis is extended to include all players from those draft years ($n = 5814$).

From the preliminary analysis of the sample of 1189 collegiate players, comparison of the univariate Kaplan-Meier functions suggests that professional career length is decreasing in collegiate experience; the Peto-Peto-Prentice test further indicates differences in duration based on collegiate career length. In multivariate analysis, the Weibull estimates provide no support for the effect of collegiate experience; however, variables related to player draft position and performance are suggested as determinants. The findings remain unchanged across the sample of 5814 draft entrants.

Contrary to the alternative hypothesis, there is no evidence to suggest that collegiate career length affects NBA career length. Of the implications for the draft strategies of amateur players and general managers the most important is that investments in collegiate experience may not be optimal.

FUNCTIONAL DATA APPROACH TO LONGITUDINAL MODELING IN THE NATIONAL HOCKEY LEAGUE

Valente, Matthew J.†; MacKinnon, David P.; Suk, Hye Won

Arizona State University, Tempe, AZ, USA

† E-mail: *m.valente@asu.edu*

Functional Data Analytic (FDA) techniques were applied to National Hockey League (NHL) end-of-season team-level data across 13 seasons. FDA techniques allow for non-parametric modeling of longitudinal data that is assumed to come from a smooth underlying curve. The purpose of this project was to demonstrate the application of FDA to NHL team-level data, assess the predictive validity of two popular advanced statistics used in the NHL, and to investigate patterns of change in these advanced statistics across the 13 NHL seasons. All data were retrieved from NHL.com and analyzed using a functional data analytic approach. In particular, a series of concurrent functional linear models were conducted to assess the best predictor of total points and patterns of change across these 13 seasons. It was hypothesized that advanced statistics such as the Fenwick and PDO would be strong predictors of total points earned by NHL teams and would account for non-overlapping proportions of variance. A concurrent functional linear model with Fenwick and PDO scores accounted for a large amount of variance in total points at the end of each season. Across the 13 seasons analyzed, the combination of Fenwick and PDO scores accounted for 65 – 87% of the variance in total points. These predictors also exhibited change in their relation with total points across the 13 seasons. That is, the strength of the relation between Fenwick and total points and PDO and total points varied across seasons. Implications and limitations are discussed.

THE IMPACT OF USING AN ANCHORED PUTTER ON WINNING AND PUTTING PERFORMANCE

Yoon, Sukjoon†[1]; Petrick, James F[1]; Kim, Jun Woo[2]

[1] *Texas A&M University, College Station, TX, USA;* [2] *Arcadia University, Glenside, PA, USA*

† E-mail: *sukjoon@tamu.edu*

On May 21st 2013 the main golf governing bodies, the Royal and Ancient (R&A) and the United States Golf Association (USGA), announced that putters anchored by the body will be banned on the professional golf tours (new rule 14-1b), starting on January 1st, 2016. While the issue

surrounding this rule was raised again after four players who used an anchored putter had won the six PGA tour major championships during the 2011-2012 seasons (Golf Digest, 2013), there is no empirical evidence that anchored putters give golfers an unfair advantage. Therefore, this study attempted to examine the direct impact of using anchored putters on career success (Study 1) and putting performance (Study 2). Censored normal regressions were employed to examine the differences between PGA Tour golfers who used anchored versus traditional putters. The sample included the population of PGA players ($n = 735$) on the tour from 2009-2014. Results revealed that players who used the traditional putter were more likely to win, averaged more prize money and were more likely to finish in the top 10 and top 25. The disadvantage of using the anchored putter was consistent across all four dependent variables (the number of winning championships, total prize money, appearance on top 10 & top 25). Additionally, players using anchored putters averaged 42% fewer one putts per round than those who used traditional putters. While the current study was unable to control for putting skill, these results suggest anchored putter users do not have an unfair advantage.