The Analysis of Serve Decisions in Tennis using Bayesian Hierarchical Models

Peter Tea and Tim B. Swartz *

Abstract

Anticipating an opponent's serve is a salient skill in tennis: a skill that undoubtedly requires hours of deliberate study to properly hone. Awareness of one's own serve tendencies is equally as important, and helps maintain unpredictable serve patterns that keep the returner unbalanced. This paper investigates intended serve direction with Bayesian hierarchical models applied on an extensive, and now publicly available data source of professional tennis players at Roland Garros. We find discernible differences between men's and women's tennis, and between individual players. General serve tendencies such as the preference of serving towards the Body on second serve and on high pressure points are revealed.

Keywords: Bayesian multinomial logistic regression, ball tracking data, Roland Garros, Roger Federer.

*P. Tea is an MSc candidate and T. Swartz is a Professor, Department of Statistics and Actuarial Science, Simon Fraser University, 8888 University Drive, Burnaby BC, Canada V5A1S6. Both Tea and Swartz have been partially supported by the Natural Sciences and Engineering Research Council of Canada. The Canadian Statistical Sciences Institute (CANSSI) Collaborative Research Team in Sports Analytics has also partially supported the research. The authors thank three anonymous reviewers whose comments have improved the manuscript.

1 INTRODUCTION

All points in tennis are initiated with a serve. The different tactics underlying each serve are driven by a player's dynamic risk tolerance, and explain why we observe a rich variety of shots differing in direction, speed and spin. Given the foray of potential serves one may encounter, especially at the professional level, serve anticipation is paramount in tennis strategy. This paper presents some candidate statistical models describing a player's intended serve decision. While there exist many important traits that can collectively describe a serve decision, this paper only considers intended *serve direction*, classified in three categories: Wide, Body and T (see Figure 1). Tactically, one must consider varying levels of faulting risk — risk of landing the ball outside the service box — when choosing a serve direction. Namely, serves aimed at the two corners (Wide and T) are *risky* and more likely to land out; by contrast, serves aimed towards the middle (Body) are *safe* and more likely to land in the service box.

From the returner's point-of-view, it is difficult to anticipate the upcoming serve direction, especially when under constrained reactionary time pressure. While research interest in serve anticipation is buoyant, the lack of accessible ball-tracking information has relegated most studies to ones of low sample size, or qualitative in nature (Reid, Whiteside, and Elliott 2011; Vernon, Farrow, and Reid 2018). In spite of these limitations, many studies highlight a list of salient variables in serve anticipation that are worthy of consideration. For instance, in a small sample-sized study of professional men's players, Reid, Whiteside, and Elliott (2011) suggest that spatio-temporal characteristics such as the ball position at serve impact are significantly different when players aim at the three distinct serve regions.

Beyond spatio-temporal features, player-specific traits are also deemed important. With Hawk-Eye tracking data, Loffing, Hagemann, and Strauss (2009) show that right- and left-handed servers exhibit different serve direction tendencies; in particular, left-handed servers frequently choose to slice serves out Wide on Advantage court. Paralleling server handedness, Loffing, Hagemann, and Strauss (2009) further suggest that returner handedness is also associated with serve direction, owing to tactical considerations. Specifically, on second serve, players are remarkably predictable and seem to repeatedly aim for the returner's backhand: a common strategy in deliberately targeting the returner's weakness and prompting a soft return shot. Serving strategy is also dictated by changes in faulting risk tolerance, which in turn is affected by match-state. Most notably, players commonly display changes in risk aversion based on serve number. Unierzyski and Wieczorek (2004) indicate that most first serves are aimed Wide and T, while Body serves are more commonly observed on second serve. These findings are not surprising, given that players typically feel emboldened on first serve, yet apprehensive of double-faulting on second serve. According to the rules of tennis, a double-fault results in a lost point to the server.

Another interesting assertion about anticipating serve direction is the ability to be conscious of a server's unique behavioural patterns. Vernon, Farrow, and Reid (2018) describe elite returners as those attuned to their opponent's serve tendencies, especially during pressure points. That is, servers may have specific serve preferences that manifest during tense situations in a match. Interestingly, Bailey and McGarrity (2012) corroborate this view and point out that during pressure situations, servers frequently hit serves to the returner's backhand. In terms of optimal strategy, servers may benefit from utilizing a more randomized serve direction pattern — a mixed strategy — that keeps the returner unbalanced and perhaps maximizes the expected number of serving points won. However, Walker and Wooders (2001) show that randomized serve strategies are uncommon in tennis and that player serve decisions are not serially independent; that is, players tend to switch from one action to another far *too* often from an optimal strategy point-of-view.

Collectively, the previous discussions suggest a rich library of candidate predictors to construct a serve direction model. The novelty of serve anticipation research presented in this paper is threefold: an application on an extensive and now publicly available serve location dataset; the joint investigation of several spatio-temporal, player-specific and match-state predictor variables; and the consideration of *intended* serve direction as the response variable. Modelling intended serve direction allows the inclusion of net-faults[1] and a complete account of all *attempted* serves within our dataset. Our goal is to predict what a server wants to do, and if we exclude the unsuccessful serves, there would be a bias in the dataset. Of course, unsuccessful serves may be viewed as serving errors, and hence unsuccessful serves whose intended bounce locations are imputed are subject to greater measurement error than successful serves. We acknowledge that imputing intended serve direction rather than consider it missing will artificially reduce the variability in the data and may also introduce bias.

The models considered in this paper are developed in a Bayesian framework: parameters are treated as random variables with distributions, and prior knowledge on parameters can be easily incorporated. This flexible approach facilitates intuitive probabilistic interpretations for parameter estimates and provides a systematic approach for prediction via the predictive distribution. The models are moderate dimensional and analyses of the models are carried out using Markov chain Monte Carlo (MCMC) methods. For a survey of some of the statistical work that has been done in sports analytics, see Albert, Glickman, Swartz and Koning (2017). Operations research is a natural venue for research in sports analytics. Recent examples of sports analytics in operation research include Cea et al. (2020), Friesl et al. (2020) and Nikolaidis (2015).

In Section 2, we begin with a description of the serve location dataset obtained from the 2019 and 2020 Roland Garros tournaments, and present the imputation methods for intended serve direction. We also provide an exploratory data analysis to guide our subsequent modelling. In Section 3, details of candidate Bayesian serve direction models are

[1]Serves that were impeded by the net and have no bounce locations.

provided. In Section 4, the Bayesian models are applied to both men's and women's data, along with a discussion of covariate effects and model fit diagnostics. We conclude with a discussion in Section 5. To date, data-driven solutions in tennis are scarce, and have stagnated from a lack of publicly available ball-tracking data. To our knowledge, this paper represents the first instance of fitting Bayesian models on extremely detailed tennis tracking data.

2 EXPLORATORY DATA ANALYSIS

Ball tracking data from the 2019 and 2020 Roland Garros tournaments has recently been made publicly available by CourtVision[2]: a product owned by Infosys. In total, the CourtVision data consists of 82 men's and 81 women's matches, amounting to 23,588 and 14,862 available serve observations, respectively. Only single's matches were considered. These matches involved 74 distinct male participants and 78 distinct female participants. Each serve observation is supplemented with both ball-tracking and match-state information. In terms of ball-tracking data, three-dimensional coordinates of the ball trajectory — including location at serve impact, location as the serve reaches the net, and serve bounce location — are provided. The (x,y,z) location coordinates refer to the longitudinal, lateral and vertical axes, respectively. Meanwhile, the match-state data include information like the current score, serve number and player IDs. Player handedness data was scraped elsewhere, from official tour websites.[3]

[2] An example: `https://www.infosys.com/roland-garros/match-centre-3d.html?matchId=SM001&year=2020&tournamentId=520&matchDate=2020-10-11`

[3] Men's: `https://www.atptour.com/en/players/` ; Women's: `https://www.wtatennis.com/players/`

2.1 Response Variable - Intended Serve Direction

Our response variable of interest — intended serve direction — has three response levels: Wide, Body and T. The visual demarcations of these three regions, along with other helpful references to tennis jargon referenced throughout this paper, are provided in Figure 1. We note that on the left-side of a full tennis court (not shown on Figure 1), the Advantage (Ad.) court sits above the Deuce court. During the serve motion, the server must stand behind the baseline and serve diagonally across court towards the returner. Specifically, the server must stand toward the right of the center-mark when serving on Deuce court and toward the left of the center-mark when serving on Ad. court. The serves alternate from Deuce court to Ad. court with each consecutive point until the completion of a game. Rules of tennis, including its scoring system, are available online at: `https://www.tenniscanada.com/tournaments/officiating/rules-of-the-court/`.

As part of data management, all (x,y,z) serve bounce locations landing inside the three serve regions were categorized as either Wide, Body or T. Faulted serves that missed the serve regions *long* (longitudinally beyond the service line) or *wide* (laterally beyond the singles sideline or in the non-intended court), were categorized with the intended serve region, identified as the serve region closest in distance to the serve bounce location. We further emphasize the server's *intention* by including imputed locations of balls obstructed by the net (i.e. "net faults"). Imputation was performed assuming a linear trajectory ball path, removing net obstruction. That is, using both the ball location at serve impact and the location as the served ball reached the net, we imputed where the ball would have landed assuming an unimpeded linear trajectory path. We note that this approximation may not be entirely valid, since players are known to impart top- and side-spin on served balls, as well as gravity that induce a non-linear ball trajectory. In both the men's and women's datasets, faults consisted of $\sim 30\%$ of serve observations; of all faults, $\sim 25\%$ are classified as net faults.

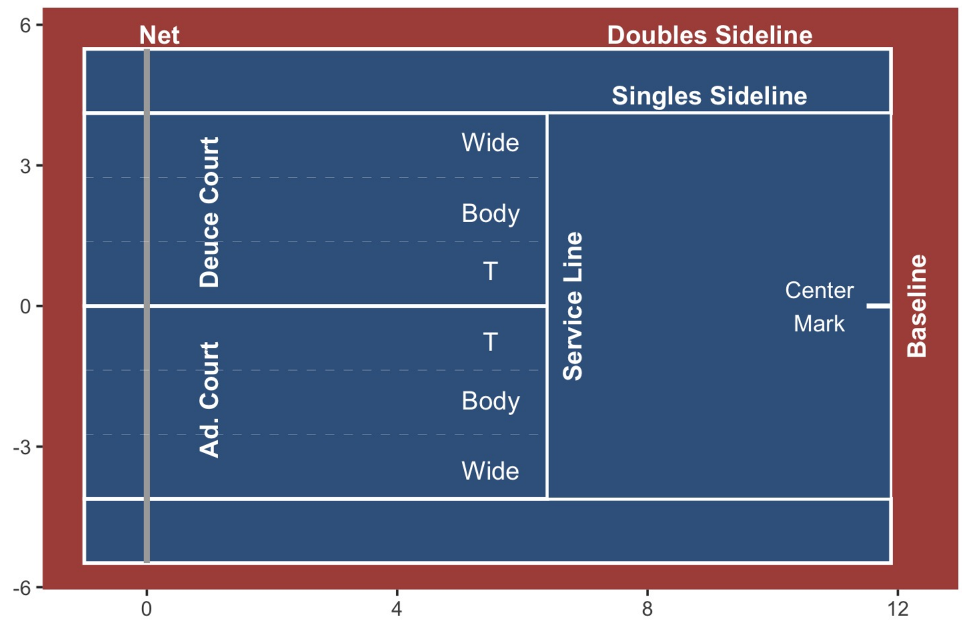

Figure 1: The anatomy of one-half of a symmetric tennis court. The axes are recorded in metres. For classifying serve direction, each service court was divided into three equally spaced zones: Wide, Body, and T.

2.2 Serve Speed

Serve speed, like serve direction, is intentional. And while it may seem logical to include speed in our serve-decision response classification,[4] we find this inclusion redundant given the consideration of other covariates. Specifically, Figure 2 illustrates a clear drop in speed from a player's first serve, to their second serve. Since serve number will be considered as a covariate in our analysis, and speed is important with respect to serve number, we purposefully omit serve speed as part of the response variable.

[4]For example, player-specific slow speeds vs. fast speeds.

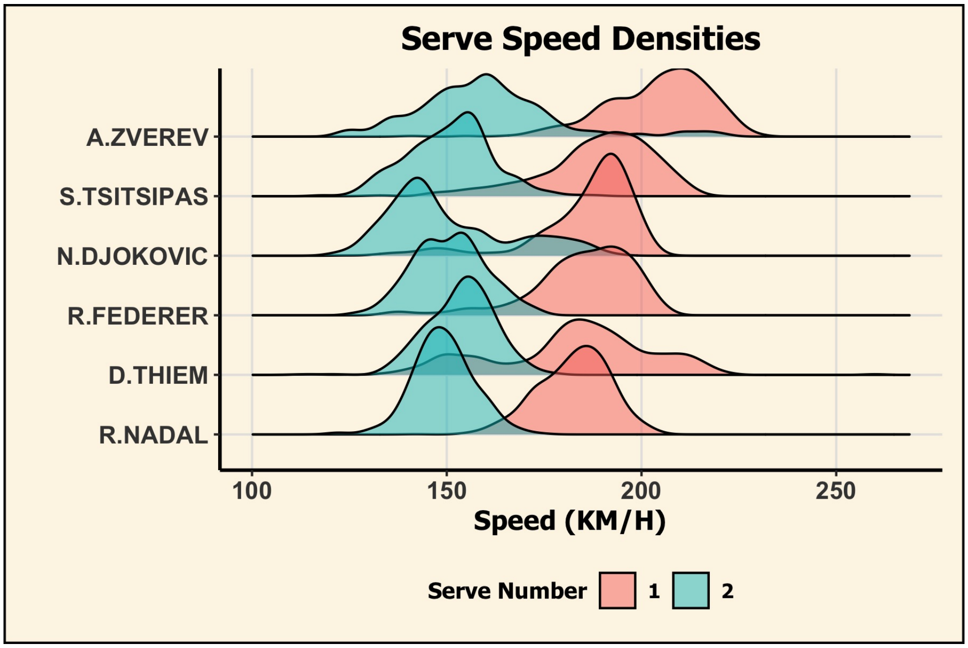

Figure 2: Serve speed densities on first and second serves for six men's players.

2.3 Covariates

In this section we produce graphics based on tennis knowledge that help inform covariates of interest to be included in our models. We emphasize certain relationships by plotting two-dimensional kernel density heatmaps, using the *kde2d* function from the *MASS* R package. Darker regions on these heatmaps represent areas with higher density of serve bounce locations. For convenience, densities were scaled such that binned areas with the highest density are given a value of 1.

With the Deuce court serve direction maps presented in Figure 3, one can observe the difference in serve direction tendencies between men's and women's tennis and between individuals. For example, the men aim their serves more towards the corners: Wide and T. Meanwhile, the women have a more conservative approach where the Body is also commonly targeted. The differences in serve direction profiles between genders is

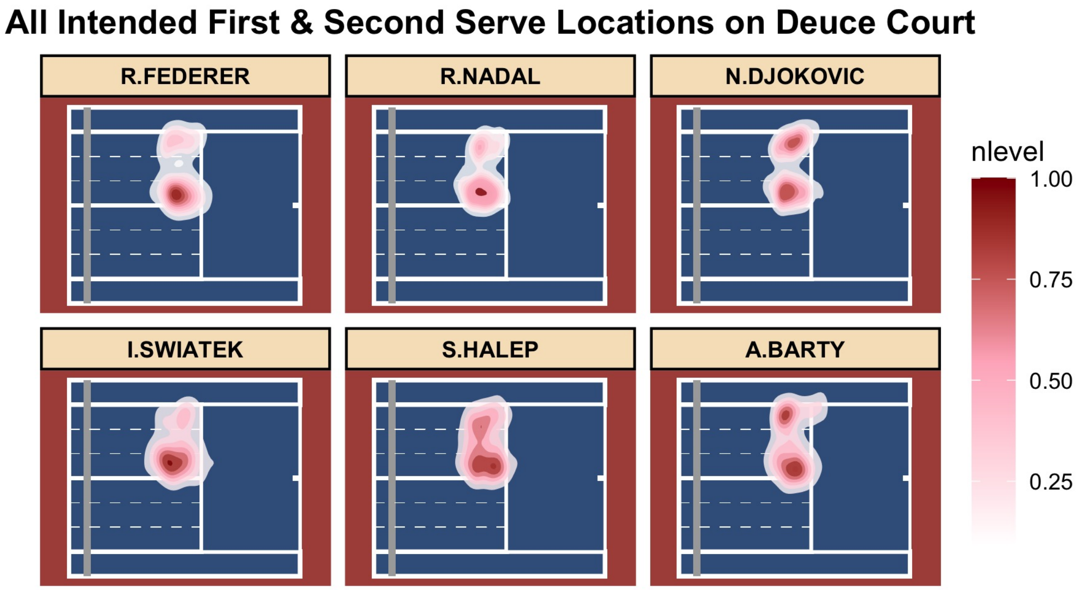

Figure 3: All first and second intended serve location densities for three men (top row) and three women (bottom row) corresponding to the Deuce court.

interesting, especially given the prevailing thought of the serve being more crucial in the men's game compared to the women's (Rothenberg 2017). Also, from Figure 3, both Nadal and Federer are inclined to serve towards T; conversely, Djokovic is more balanced towards both Wide *and* T. Similar patterns are also evident when considering serve direction densities on the Advantage court. Clearly, serve directions are server-dependent.

2.3.1 Match Pressure

It is a popular belief that some players have a supposed *go-to* serve they resort to in pressure situations, like a breakpoint. However, breakpoints do not represent all points with a sizable impact on match outcomes: for example, serving when down 15 - 30 in the fifth, match-deciding set is also a tense moment. Consequently, we consider *point importance* — a proxy for match pressure — which calculates a player's expected change in match win probability, depending on whether they win or lose the current point (Kovalchik 2017). If

P(a,b) represents the probability that the current server eventually wins the match, with a and b denoting the server's and returner's current scores respectively, then the point importance is calculated as P(a+1, b) – P(a, b+1), where a+1 is the change in score if the server wins the point. Match win probabilities are derived with recursive formulas unique to the tennis scoring system. Initially proposed by Morris (1977), and popularized by Kovalchik and Reid (2018), point importance is a numeric score lying in the interval [0,1], where 0 represents no importance and 1 represents maximal match importance.

To highlight point importance's impact on serve decision, we present a case-study involving Roger Federer. For convenience, point importance was arbitrarily binned such that all points above the 80th quantile are *High Importance*, while all other points are *Low Importance*. In total, Federer had 113 High Importance points and 442 Low Importance points. Presenting the point importance serve location densities in Figure 4(a), we notice that on important points, Federer typically serves toward T on Deuce Court, and opts for more Wide/Body serves on Advantage Court. By contrast, Federer is less predictable on his less important points. This case-study is, of course, only a sliver of point importance's impact on serve direction. For other players — Djokovic, for example — there is virtually no change in serve direction behaviour, as seen in Figure 4(b).[5]

2.3.2 Ball Position at Serve Impact

Reid, Whiteside, and Elliott (2011) suggest that the ball's lateral position at serve impact can inform players on the intended serve direction. Presumably, the further away a player strikes the serve from the center mark, the more angle is created that enables a Wide serve. From Figure 5, we see that the relationship between lateral displacement and serve direction is non-existent on first serve, but more apparent on second serve. That is, it is only on second serve where the greater the lateral displacement, the more likely the ball

[5]There may be some mixed strategy implications here with Djokovic seemingly more willing to randomise his serve direction options; at least, more so than Federer.

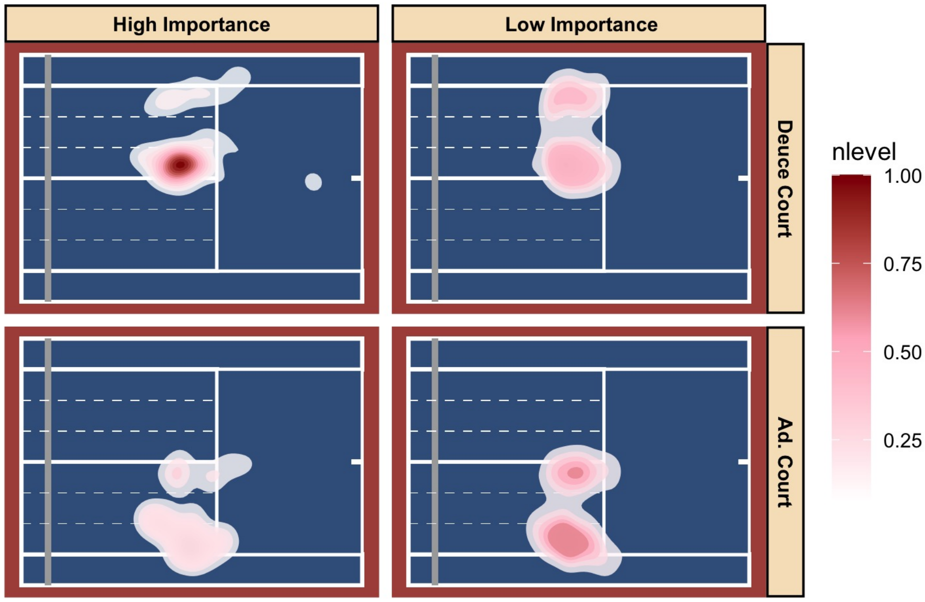

(a) Roger Federer.

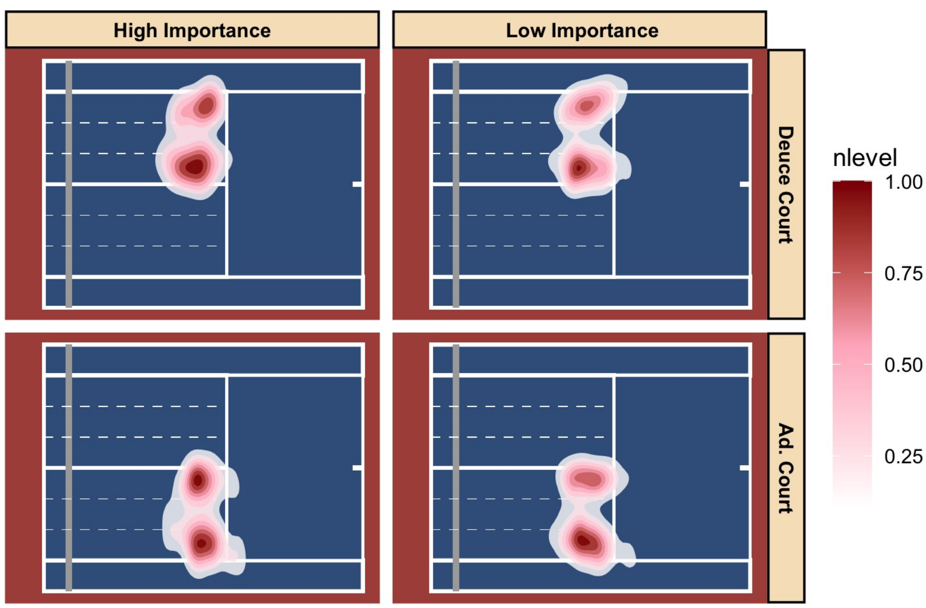

(b) Novak Djokovic.

Figure 4: Serve locations on important and less important points.

will land Wide.

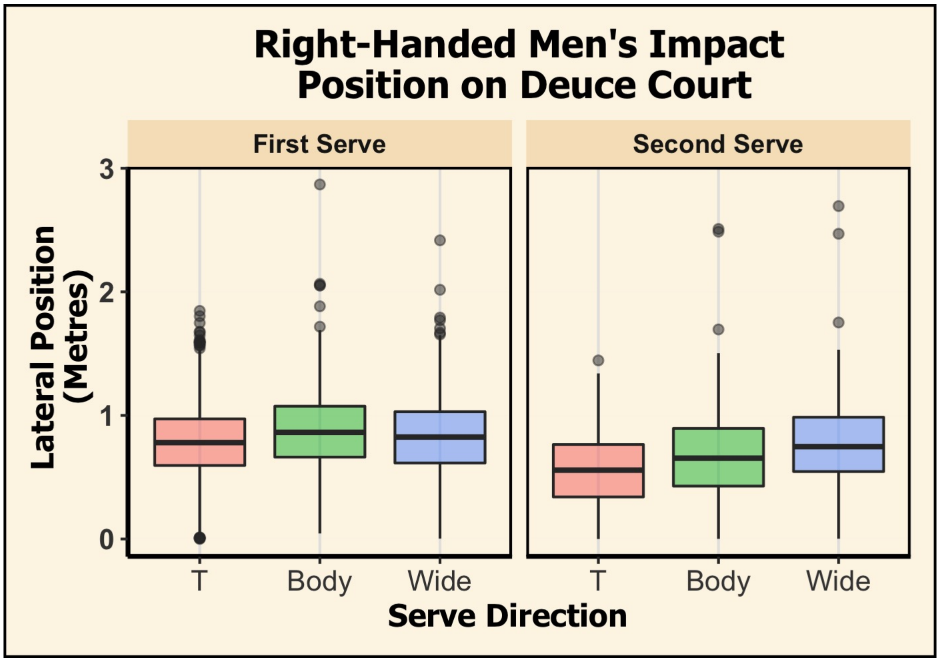

Figure 5: Ball lateral position densities at serve impact.

Interestingly, Reid, Whiteside, and Elliott (2011) also report that the longitudinal coordinate at serve impact is *not* significantly associated with serve direction. However, since their study featured only six right-handed participants and was conducted under non-match settings, we choose to include the longitudinal coordinate as part of our covariates of interest. Under real, live-match settings, some players opt for a serve-and-volley strategy,[6] which requires a serve impact location inside the Baseline. Under this unique strategy, players will likely choose Wide or T serves, over the Body serve. Other spatio-temporal features, like ball toss trajectory and player position on the baseline, are also deemed important. However, these features could not be extracted from our data, and will not be considered.

[6]Strategy where players follow their serve immediately towards the net.

2.3.3 Other Covariates

In addition to the covariates mentioned in this section, we will also consider serve decision lags in our potential models. According to Walker and Wooders (2001) player serve decisions are not serially independent. As such, in modelling a player's t^{th} serve decision, we will consider their previous serve decision at serve *t-1*.

Throughout our exploratory analysis, we find that certain serve direction patterns exist more prominently under interacting schemas of our covariates of interest. In particular, players overtly target the returner's backhand (BH) on second serve, and lefties are notorious for serving Wide on Advantage court. Therefore, in addition to main covariate effects, we also consider two interaction terms: one between serve number and returner's backhand location, and the other between court side and server handedness. Table 1 summarizes the full list of covariates.

Covariate	**Description**
Serve Number	0 if first serve; 1 if second serve.
Lateral Distance	Serve impact lateral distance from center mark.
Longitudinal Distance	Serve impact longitudinal distance from baseline.
Serve at $t-1$ is T	1 if T, 0 else.
Serve at $t-1$ is Wide	1 if Wide, 0 else.
Returner's BH Location	1 if BH location is T; 0 else.
Point Importance	[0,1] denoting match pressure.
Server Handedness	1 if left-handed; 0 else.
Court Side	1 if Advantage court; 0 else.
Interaction 1	Court Side x Server Handedness.
Interaction 2	Serve Number x Returner's BH Location.

Table 1: Candidate predictor variables.

3 METHODS

3.1 Bayesian Multinomial Logistic Regression

A multinomial logit (MNL) model with C unordered response categories requires the specification of C-1 logit equations, owing to identifiability concerns. In our application C=3, denoting the three intended serve regions. Let $Y_{it} \in (1, \ldots, C)$ represent an integer-indexed serve decision for server i ($i = 1, \ldots, m$) on their t^{th} serve number ($t = 1, \ldots, n_i$). To limit model complexity, we assume that serve decisions do not vary significantly match-to-match, and instead allow the subscript t to expand across all matches.

Let $\pi_{itr} = \text{Prob}(Y_{it} = r)$ denote the probability of player i choosing serve decision r on their t^{th} serve. Choosing the last level, C, as our *reference category*, the C-1 logit equations are written as

$$\eta_{itr} = \log\left(\frac{\pi_{itr}}{\pi_{itC}}\right), \qquad r = 1, \ldots, C-1. \tag{1}$$

With the C-1 logit equations in (1), we can then obtain the response category probabilities using the *softmax* function,

$$\pi_{itr} = \frac{\exp\{\eta_{itr}\}}{1 + \sum_{s=1}^{C-1} \exp\{\eta_{its}\}},$$

for r = 1, ..., C-1. For the reference category, we have

$$\pi_{itC} = \frac{1}{1 + \sum_{s=1}^{C-1} \exp\{\eta_{its}\}}.$$

In this paper, we consider three different Bayesian model structures for η_{itr}; each structure differs in their intercept component or inclusion of match covariates. We present these three model formulations in the following sections, which were largely inspired by Congdon (2020) and Liu (2015). All models are fit with STAN, which uses Hamiltonian Monte Carlo methods to generate posterior samples (Carpenter et al. 2017).

3.1.1 Baseline Common Intercept Model

The simplest structure we consider for η_{itr} is a common intercept model,

$$\eta_{itr} = \alpha_r \tag{2}$$

where α_r represents a common intercept for each r^{th} logit component. The model in (2) does not consider player-varying effects, nor does it consider match covariates. While this model structure is shallow, it will serve as a *baseline* comparison model against our more elaborate model structures. We impose a default prior on α_r as follows:

$$(\alpha_1, \ldots, \alpha_{C-1}) \sim N_{C-1}\big(0, \nu I_{C-1}\big), \tag{3}$$

where ν is a scalar and I_{C-1} is the identity matrix with $C-1$ rows and columns.

3.1.2 Player-Varying Intercept Model

The underlying tennis assumption associated with model (2) is that player strategies with respect to serve are centred about a common strategy. This makes sense from the point-of-view that there is a long history of tennis playing experiences which dictate standard and effective playing strategies. These ideas naturally suggest a hierarchical model where players can have different intercept parameters, drawn from a common distribution. We consider a player-varying intercept model where we allow intercepts to vary across server i as follows

$$\eta_{itr} = \alpha_{ir}. \tag{4}$$

In the above formulation (4), α_{ir} denotes player-level varying intercepts for each r^{th} logit component, and are given a multivariate normal prior,

$$(\alpha_{i1}, \ldots, \alpha_{i,C-1}) \sim N_{C-1}\big(\,(\alpha_1, \ldots, \alpha_{C-1})\,, \Sigma_{\alpha_i}\big).$$

To complete the hierarchical setup, we specify hyperprior distributions for the mean and covariance parameters. The mean parameter is given a weakly informative multivariate normal prior, similar to formulation (3). For the prior on the covariance matrix, we impose the same distribution for each server, i. A Cholesky factorization on Σ_{α_i} is implemented, and weakly informative priors are assigned: a half-Cauchy on the scale parameter and a computationally efficient Lewandowski-Kurowicka-Joe (LKJ) distribution originally proposed by Lewandowski et al. (2009) on the Cholesky factor of the correlation matrix (Stan Development Team 2021).

3.1.3 An Expanded Model

Covariates beyond player effects can also influence serve decisions. We expand our player-varying intercept model in (4) by including covariates,

$$\eta_{itr} = \alpha_{ir} + X_{it}^T \beta_r, \tag{5}$$

where α_{ir} again denotes player intercepts, X_{it} is a p x 1 covariate vector, and β_r is a p x 1 parameter vector to be estimated. For simplicity, we include only the subscript r to emphasize that a set of coefficients are obtained for each r^{th} response category; however, the components can also depend on player i or even the serve number index, t. The components of β_r are given vague multivariate normal priors with mean 0 and standard deviation $\nu_\beta = 3$, independently. We consider all eleven covariates listed in Table 1 for the Expanded Model.

3.2 Predictive Distributions

For a serve prediction, $\tilde{y}$, on a new match-state, $\tilde{x}$, the predictive distribution of $\tilde{y}$ is given by

$$p\left(\tilde{y} \mid \tilde{x}, x, y\right) = \int p(\tilde{y} \mid \tilde{x}, \theta) \cdot \pi(\theta \mid y, x) d\theta \tag{6}$$

where θ represents the MNL model parameters, y are the historical data and x are the corresponding historical match-states. To simulate from the predictive distribution in (6), we first draw parameter values $\theta^{(1)}, \ldots, \theta^{(M)}$ from the posterior distribution, $\pi(\theta \mid y, x)$. With the m^{th} drawn parameter value and a new match-state, $\tilde{x}$, we draw $\tilde{y}^{(m)}$ from the sampling distribution $p(\tilde{y} \mid \tilde{x}, \theta^{(m)})$. Repeating this procedure, we obtain a predictive sample $\tilde{y}^{(1)}, \ldots, \tilde{y}^{(M)}$ which can be used to address serve direction questions presented in Section 4.4.

4 RESULTS

All models were fit by modifying STAN code published by Koster and McElreath (2017). In our model fitting, we set the serve direction "Body" as the reference category. In total, 4,000 draws were obtained from each posterior distribution.

4.1 Model Comparison

To compare the candidate models, we compute Watanabe-Akaike Information Criterion (WAIC) where lower WAIC values indicate stronger model fit. The WAIC diagnostic estimates the expected log pointwise predictive density, which represents out-of-sample prediction error. WAIC includes a penalty term that accounts for model complexity. From Vehtari, Gelman, and Gabry (2017), we also compute WAIC standard errors to assess WAIC uncertainty. In Table 2, we provide pairwise differences in model WAIC along with standard errors (SE).

We note that for a pair of models, a large ΔWAIC relative to its SE reflects a discernible difference in model fit. From Table 2, the Expanded Model greatly outperforms its intercept model analogues for both men and women. It is also seen that the Player-Varying Intercept Model is preferable to the Common Intercept Model.

Model	WAIC	ΔWAIC (SE)
Men's		
Expanded Model	41,147.5	
Player-Varying Intercept	44,378.2	3230.7 (111.0)
Common Intercept	44,744.0	3596.5 (118.6)
Women's		
Expanded Model	28,281.0	
Player-Varying Intercept	30,009.0	1728 (86.6)
Common Intercept	30,587.4	2306.4 (98.2)

Table 2: Model comparisons with WAIC.

4.2 Player-Varying Intercepts

We explore player differences in serve direction behaviour with scatter plots (see Figure 6) of posterior T and Wide intercept means, obtained from the Player-Varying Intercept Model of Section 3.1.2. The means are transformed onto the probability scale. In general, the men have higher probabilities serving T and Wide than the women, which affirms the men's preference aiming towards the corners rather than the Body. This is apparent by noting that the points for men in Figure 6(a) lie further in the upper right quadrant than the points for women in Figure 6(b). However, a few women also habitually aim towards the corners — among them is 23 time Grand Slam champion, Serena Williams. With an estimated Prob(Wide) above 0.5, accompanied with a high Prob(T) of 0.4, returners are forced to cover more ground when bracing against Williams' formidable serve. Similar to Williams, both Shapovalov and Thiem in the men's game exhibit clear preferences aiming Wide; however, other servers such as DelPotro, Siniakova and Vondrousova clearly prefer aiming T over Wide. Preferences aiming T may be explained by net height; the net is six inches lower in the center, making it easier to serve accurately when aiming T compared to when aiming Wide. Interestingly, players like Djokovic and Kenin maintain a more balanced serve direction pattern in aiming both Wide and T nearly equally as often.

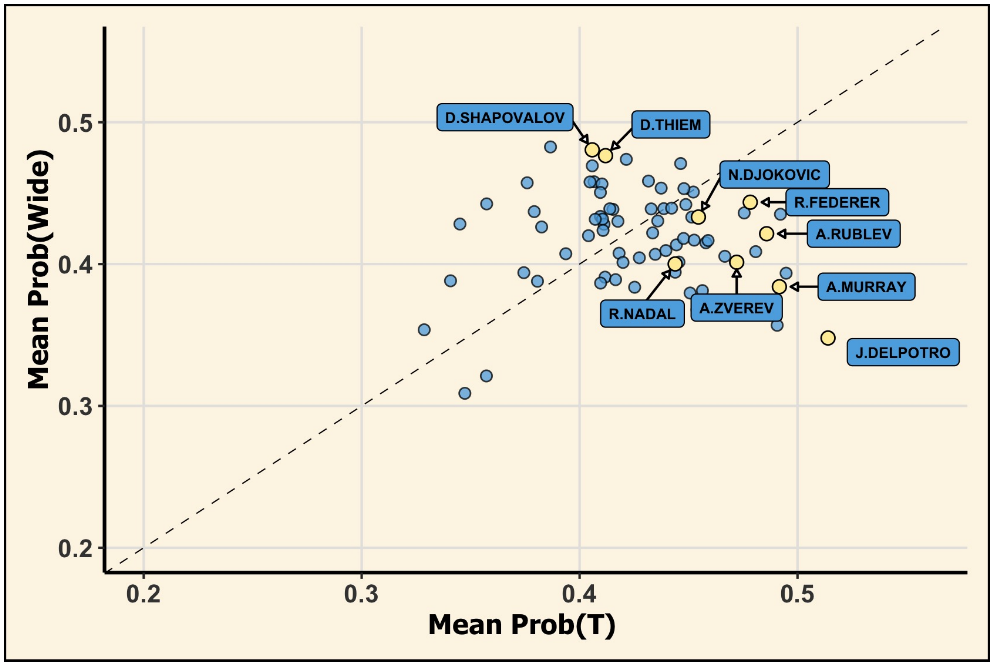

(a) Mens players.

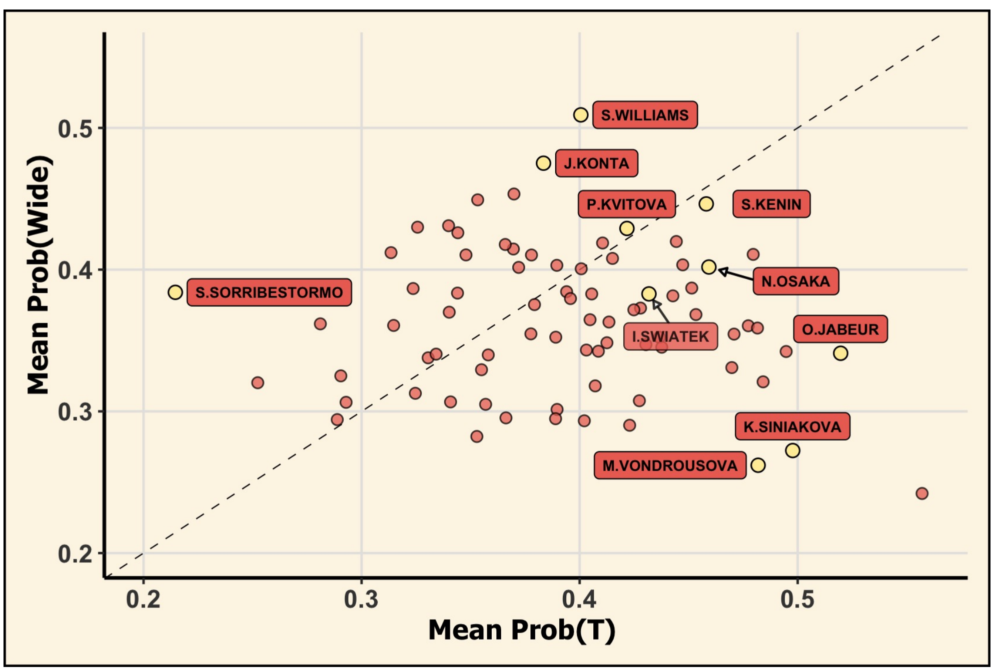

(b) Womens players.

Figure 6: Comparing player serve direction tendencies. The dashed-line represents equal probability serving Wide and T.

4.3 Covariate Effects

With the Expanded Model of Section 3.1.3, we investigate how covariates alter a player's intended serve direction. Posterior summaries of all eleven covariates are presented in Table 3. On average, the second serve greatly deters players from aiming towards the two corners. The same can be said of point importance where the greater the importance, the more likely a player will be conservative and aim towards the Body. Consistent with the literature review, higher lateral impact distances promote a Wide or Body serve over a T serve. Meanwhile, higher longitudinal impact distances are associated with serves aimed Wide or T — especially among women. From the interaction effects, left-handed servers often aim Wide on Ad. court, although the strength of this association is stronger for men compared to women. This difference may be explained by the disparity of available left-handed serve observations: the women do not have a southpaw as dominant as Nadal at Roland Garros. From the second interaction effect, servers deliberately target the returner's backhand on second serve. The remaining covariates have moderate effects on serve direction.

	T Logit		Wide Logit	
	Mean	SE	Mean	SE
Men's				
Serve Number	-2.17	0.07	-1.58	0.06
Lateral Distance	-0.61	0.06	0.16	0.06
Longitudinal Distance	0.25	0.09	0.14	0.08
Serve at $t-1$ is T	-0.07	0.07	-0.01	0.07
Serve at $t-1$ is Wide	-0.15	0.07	-0.20	0.07
Returner's BH Location	0.18	0.08	-0.07	0.08
Point Importance	-1.55	0.55	-0.75	0.53
Server Handedness	0.14	0.20	-0.15	0.21
Court Side	0.10	0.07	-0.04	0.07
Interaction 1	-0.38	0.14	0.36	0.13
Interaction 2	1.39	0.09	-0.10	0.09
Women's				
Serve Number	-1.45	0.08	-0.97	0.07
Lateral Distance	-0.82	0.08	0.68	0.08
Longitudinal Distance	1.75	0.17	1.14	0.17
Serve at $t-1$ is T	0.04	0.06	0.12	0.06
Serve at $t-1$ is Wide	0.05	0.06	0.01	0.06
Returner's BH Location	-0.03	0.10	-0.09	0.09
Point Importance	-0.95	0.51	-0.45	0.51
Server Handedness	0.26	0.24	-0.21	0.26
Court Side	0.03	0.09	0.00	0.09
Interaction 1	-0.48	0.17	0.18	0.18
Interaction 2	0.81	0.10	-0.44	0.10

Table 3: Posterior summaries of the parameters β_r corresponding to the Expanded Model (3.1.3).

4.4 Prediction Example: Federer

Following Section 3.2, we illustrate an application of predictive distributions with our Expanded Model. For a new match-state, we imagine the following scenario: Roger Federer trailing 0-1 in sets, 5-6 in games, and facing breakpoint against a right-handed returner. If Federer previously served T, which direction will he now serve? How would his predicted serve direction change on first serve compared to second serve? For this unique match scenario, we present Federer's predictive distributions in Figure 7. Consistent with Federer's temperament, on first serve the predictive probabilities reflect Federer's disdain for the Body direction and instead favour Wide and T. Meanwhile on second serve, there is a clear preference serving Wide, which incidentally would target the returner's backhand.

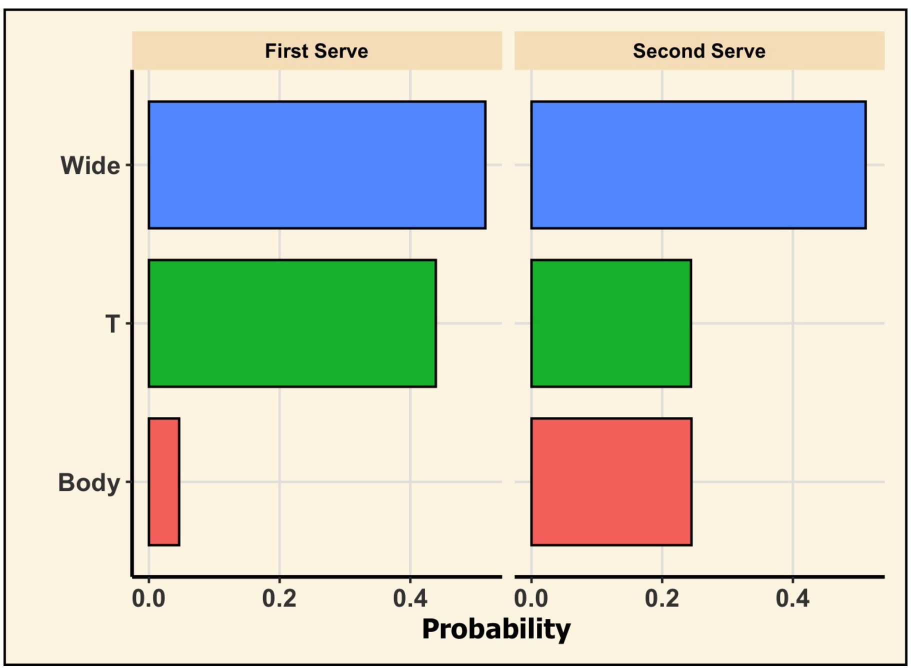

Figure 7: An example of Roger Federer's predictive probabilities. Probabilities were obtained form the proportion of predicted serve directions among the 4000 posterior draws.

We evaluate the predictive performance of the Expanded Model by fitting on a training set and calculating accuracy on the held-out test set. For this illustration, we randomly selected one of Roger Federer's matches as the test set, and fit the Expanded model on the remaining training data. The test set comprised of 126 Federer serve observations ($\sim 22\%$ of Federer's available data). Doing so, we obtained an overall accuracy of 64%.

5 DISCUSSION

Moving forward, more questions about serve direction can now be answered with the availability of ball-tracking data. For instance, what are the consequences of a player's choice about serve direction? For the last several years, both Nadal and Thiem have arguably been the premier clay court players in the men's game. But from Figure 6(a), both these players have drastically different serve direction patterns: Thiem prefers aiming Wide while Nadal prefers aiming T. There exists several different serve strategies ranging from unpredictable serve directions to blatant barrages aimed at the returner's backhand, but with unexplored success rates. Moreover, a further exploration of *why* player serve patterns are different can also be investigated. Vaverka and Cernosek (2013) have shown that body height is associated with serve speed, but could height also be associated with serve direction? At greater heights, a player has a higher margin of error to land their serve inside the serve box, which may promote more serves intended for the two corners. Unfortunately, we could not reliably assess the association between player height and intended serve direction since serve data from taller players like John Isner, Milos Raonic and Reilly Opelka were missing.

This paper implements a Bayesian framework analysing intended serve direction among professional tennis players. The framework considers a conglomerate of serve anticipation variables, ranging from spatio-temporal features to player-specific and match-state variables. We emphasize that the players in this study include top-ranked professionals who routinely get invited to Grand Slam tournaments. The inferences gained cannot be generalized to amateur players. Moreover, our framework assumes that players follow a consistent serve direction pattern across matches. However, we keep in mind that players can easily change their serve tendencies on a whim. The vagaries of player tendencies emphasize the importance of frequently updating predictive probabilities.

In general, the main difference in intended serve direction between the men's and women's game relates to risk: the men are more inclined to choose risky serves towards T and Wide, whereas the women blend in a notable amount of safe Body serves. Among all players, we observe a tendency to shrink towards more conservative Body serves during tense situations such as on second serve and on high pressure points. We find that players also frequently target the returner's backhand on second serve, and that left-handed servers distinctly prefer aiming Wide on Advantage court.

6 REFERENCES

Albert, J., Glickman, M., Swartz, T. and Koning, R. (2017). *Handbook of Statistical Methods and Analyses in Sports*. Chapman Hall/CRC Hand-books of Modern Statistical Methods. Boca Raton.

Bailey, B. and McGarrity, J. (2012). "The Effect of Pressure on Mixed-Strategy Play in Tennis: The Effect of Court Surface on Service Decisions". *International Journal of Business and Social Science*, Vol 3.

Carpenter, B., Gelman, A., Hoffman, M. D., Lee, D., Goodrich, B., Betancourt, M., Brubaker, M., Guo, J., Li, P. and Riddell, A. (2017). "Stan: A Probabilistic Programming Language". *Journal of Statistical Software*, Vol 76, 1-32.

Cea, S., Durán, G., Guajardo, M., Sauré, D., Siebert, J. and Zamorano, G. (2020). "An Analytics Approach to the FIFA Ranking Procedure and the World Cup Final Draw". *Annals of Operation Research*, Vol 286, 119–146.

Congdon, P. (2020). *Bayesian Hierarchical Models: With Applications using R*. 2nd ed. CRC Press, Taylor Francis Group. Chap. 8.

Friesl, M., Libich, J. and Stehlík, P. (2020). "Fixing Ice Hockey's Low Scoring Flip Side? Just Flip the Sides." *Annals of Operation Research*, Vol 292, 27–45.

Koster, J. and McElreath, R. (2017). "Multinomial Analysis of Behavior: Statistical Methods". *Behavioral Ecology and Sociobiology*, Vol 71, 1–14.

Kovalchik, S. (2017). "Quantifying Clutch Performance", retrieved May 4th, 2021 at `http://on-the-t.com/2015/12/27/quantifying-clutch-performance/`.

Kovalchik, S. and Reid, M. (2018). "Measuring Clutch Performance in Professional Tennis". *Italian Journal of Applied Statistics*, 255–268.

Lewandowski, D., Kurowicka, D. and Joe, H. (2009). "Generating Random Correlation Matrices Based on Vines and Extended Onion Method". *Journal of Multivariate Analysis*, Volume 100, Issue 9, 1989-2001.

Liu, X. (2015). *Methods and Applications of Longitudinal Data Analysis.* 1st ed. Elsevier Science Technology. Chap. 11.

Loffing, F., Hagemann, N. and Strauss, B. (2009). "The Serve in Professional Men's Tennis: Effects of Players' Handedness". *International Journal of Performance Analysis in Sport*, Vol 9, 255–274.

Morris, C. (1977). "The Most Important Points in Tennis". *Optimal Strategies in Sports*, Vol 5, 131–140.

Nikolaidis, Y. (2015). "Building a Basketball Game Strategy Through Statistical Analysis of Data." *Annals of Operation Research*, Vol 227, 137–159.

Reid, M., Whiteside, D. and Elliott, B. (2011). "Serving to Different Lo-cations: Set-up, Toss, and Racket Kinematics of the Professional TennisServe". *Sports Biomechanics*, Vol 10, 407–414.

Rothenberg, B. (Sept. 2017). "Filling a Weak Spot in Women's Tennis: The Serve". *The New York Times.*

Stan Development Team (2021). *Stan Modeling Language Users Guide and Reference Manual, Version 2.26*.

Unierzyski, P. and Wieczorek, A. (2004). "Comparison of Tactical Solutions and Game Patterns in the Finals of Two Grand Slam Tournaments in Tennis". *Science and Racket Sports III*, 200–205.

Vaverka, F. and Cernosek, M. (2013). "Association Between Body Height and Serve Speed in Elite Tennis Players." *Sports Biomechanics*, Vol 12, 30-37.

Vehtari, A., Gelman, A. and Gabry, J. (2017). "Practical Bayesian Model Evaluation using leave-one-out Cross-Validation and WAIC". *Statistics and Computing*, Vol 27, 1413–1432.

Vernon, G, Farrow, D. and Reid, M. (2018). "Returning Serve in Tennis: A Qualitative Examination of the Interaction of Anticipatory Information Sources Used by Professional Tennis Players". *Frontiers in Psychology*, Vol 9.

Walker, M. and Wooders, J. (2001). "Minimax Play at Wimbledon". *American Economic Review*, Vol 91, 1521–1538.