

Disequilibrium Play in Tennis*

Axel Anderson, *Georgetown University*[†]

Jeremy Rosen, *Topspin Shot Research*[‡]

John Rust, *Georgetown University*[§]

Kin-Ping Wong, *Digonex*[¶]

September 8, 2020

Abstract

Are the serves of the world's best tennis pros consistent with the theoretical prediction of Nash equilibrium in mixed strategies? We analyze their serve direction choices (to the returner's left, right or body) with data from an online database called the Match Charting Project. Using a new methodology, we test and decisively reject a key implication of a mixed strategy Nash equilibrium, namely, that the probability of winning a service game is the same for all serve directions. We also use dynamic programming (DP) to numerically solve for the best-response serve strategies to probability models of service game outcomes estimated for individual server-returner pairs, such as Novak Djokovic serving to Rafael Nadal. We show that for most elite pro servers, the DP serve strategy significantly increases their service game win probability compared to the mixed strategies they actually use, which we estimate using flexible reduced-form logit models. Stochastic simulations verify that our results are robust to estimation error.

Keywords tennis, games, Nash equilibrium, Minimax theorem, constant sum games, mixed strategies, dynamic directional games, binary Markov games, dynamic programming, structural estimation, muscle memory

JEL Codes: C61, C73, L21.

***Acknowledgements:** We thank Georgetown University for research support, and Rust gratefully acknowledges financial support provided by the Gallagher Family Chair in Economics, and helpful feedback from Pierre-Andre Chiappori, Peter Cramton, Drew Fudenberg, David Levine, Nicolas Legrand, Ignacio Palacios-Huerta, David Romer, Frank Vella, Mark Walker, and John Wooders.

[†]Department of Economics, Georgetown University, e-mail: aza@georgetown.edu

[‡]E-mail: jeremyr@topspinshotresearch.com

[§]Department of Economics, Georgetown University, e-mail: jr1393@georgetown.edu

[¶]Digonex, e-mail: kjw65@georgetown.edu

1 Introduction

Walker and Wooders (2001) (WW) analyzed 40 tennis “point games” from Grand Slam Tournaments, focusing on the server’s choice of first serve direction. They modeled first serves as a sequence of independent and identical simultaneous move games between the server and returner, where each has two possible decisions, left or right. They concluded that serve location choices are consistent with mixed strategy Nash equilibrium in their hypothesized static game. In particular, the server’s chance of winning a point is the same whether the serve is to the left or the right. Equality of win rates across serve directions has been confirmed in several follow up studies using additional data.¹ In contrast we find significant differences in win rates across serve locations and estimate that servers could significantly increase their chances of winning if they were to fully exploit these differences.² Our conclusions are based on analyses of the top professional tennis players such as Roger Federer, Rafael Nadal, and Novak Djokovic.

Our analysis differs from WW by considering three serve directions (left, right, and *body*) and modeling tennis a *dynamic game*. We allow body serves because tennis professionals believe they are important, see e.g. Rive and Williams (2011). Dynamics are relevant because the server’s strategy and the probability of winning the service game could depend on the score state in the presence of *muscle memory effects* — the possibility that a serve is more likely to be successful or a receiver is more effective in returning a serve hit to the same location as the previous serve.³ We show that muscle memory can induce serial correlation in serve locations even if play is in Nash equilibrium. Previous studies including WW have found serial correlation and interpreted it as evidence against Nash equilibrium.⁴

¹ Hsu, Huang, and Tang (2007) confirmed WW’s main conclusions using a slightly larger data set. Gauriot, Page, and Wooders (2018) used data from 3000 matches and nearly 500,000 serves and confirmed WW’s conclusions and noted that “the behavior in the field of more highly ranked (i.e., better) players conforms more closely to theory.”

² Klaassen and Magnus (2009) abstract from server direction and focus on the tradeoff between making a serve hard to return and faulting on the serve, considering both the first and second serves of a point. They reject the hypothesis that servers optimally solve this tradeoff, but find that “the estimated inefficiencies are not large.”

³ Habit-formation effects are well known in tennis. For example, Wiles (2006) also explores muscle memory effects, although he calls it a “timing variable.”

⁴ A notable exception is the thesis of Wiles (2006) who also pointed out that serial correlation could be consistent with Nash equilibrium. Klaassen and Magnus (2001) tested whether successive points in tennis are independent and identically distributed (*IID*) binary random variables using 481 Wimbledon matches containing nearly 90,000 points. They rejected the *IID* hypothesis, they found that “Deviations from iid are small, however, and hence the iid hypothesis will still provide a good approximation in many cases.” The study by Gauriot et al. (2018) analyzed 3000 matches and nearly 500,000 serves “resoundingly rejects the hypothesis that the direction of the serve is serially

The introduction of a third serve direction, combined with potentially state dependent serve direction probabilities, makes it far more challenging to test for equal win probabilities. Our analysis is based on an online database called the Match Charting Project (MCP) see Sackmann (2013). The MCP contains crowdsourced play-by-play data on professional tennis matches and records all three serve directions used in our analysis. Even after restricting to matches played on hard courts,⁵ we end up with roughly ten times as many serves per server-returner pair than WW use in their analysis.

We analyze the *service game* between a server and returner that ends when one wins at least 4 points and at least 2 more points than their opponent. At each serve, the server chooses the location, speed, and spin of the serve, while the returner allocates a fixed attention budget to the three serve locations. While our theory allows for rich strategy sets for the players, we only observe serve locations in our data. However we show that it is possible to analyze serve strategies as a single-agent dynamic programming (DP) problem, since in a Nash equilibrium of the dynamic game, the server’s strategy must constitute a best response to the returner’s strategy. Specifically, if each player only cares about winning or losing the service game, then the server’s probability of winning (value function) is a well-defined function of the current muscle memory state and the cumulative score in the current service game (score state).

We show that in all possible muscle memory and score states, all subgame perfect equilibria result in the same win probability for the server. Further, we prove there is a Markov Perfect Equilibrium (MPE) in which the server’s and returner’s strategies depend only on the muscle memory and score state.⁶ This allows us to define the *point outcome probabilities* (POPs) which are the probabilities that a serve to a given direction is in, as well as the probability the server wins the rally given the serve is in conditional on the current muscle memory and score state. These probabilities enable us to recast the game as a single-agent dynamic programming (DP) problem in which the server chooses serve locations to maximize his chance of winning the service game given the POPs.

independent.” (p. 1). We explore the theoretical relationship between the types of muscle memory effects and the implied sign of the serial correlation in serve location choices in Appendix B.

⁵ We do this to eliminate a potential source of heterogeneity that could confound our results, since grass and clay courts have different playing characteristics than hard courts.

⁶ While the underlying characteristics of the game do not directly depend on the current score, with muscle memory effects strategies do generally depend on the score state.

In a MPE, strategies must be mutual best responses. In particular, the server cannot increase his chance of winning the service game by changing his serve location strategy in any state of the game. In order to test this necessary condition, we need to estimate the POPs and the actual serve strategy used in the service game. However, our model has 298 muscle memory/score states, three serve directions and a fully unrestricted estimator of the serve strategy and the POPs would require 2512 parameters for each server-returner pair — far too many to estimate precisely given the size of our data set. In Section 3, under a testable assumption that serve strategies and POPs are stationary and Markovian (but not necessarily MPE strategies), we estimate flexible reduced-form parametric models of serves and the POPs that includes the unrestricted specification as a special case. We use the Akaike Information Criterion (AIC) to select a preferred specification with 44 parameters, 12 for the server’s strategy and 32 for the POPs, that balances the desire for flexibility against the danger of overfitting.⁷

Rather than separately testing for equal win probabilities across serve locations at each node in the game tree, we derive a new efficient omnibus Wald test of the hypothesis of equal win probabilities *across all possible states of the game, simultaneously*. This new test decisively rejects the hypothesis of equal win probabilities, even for elite pros such as Federer, Nadal, and Djokovic. In order to get a sense for the magnitude of the violations, we use DP to calculate best response serve strategies for individual server-returner pairs using estimated POPs to provide outcome probabilities for each point given the choice of serve direction. For all the elite pros we analyzed, the DP strategy significantly increases win probabilities relative to the mixed serve strategies implied by our reduced form estimates of their serve behavior. For example, we predict that by adopting the DP best response serve strategy Nadal can improve his chances of winning each service game vs. Djokovic from 71% to 91.5%, while Djokovic could improve his chance of winning vs. Nadal from 83% to 93.7%.⁸ Top tennis play does not constitute a MPE.

⁷ The first empirical analysis of tennis using statistical/probabilistic methods that we are aware of is by George (1973) who analyzed the decision of whether the serve should be strong (i.e. fast, more difficult to return but higher probability of faulting) versus weak (i.e. slow, easier to return but lower probability of faulting). The first formal DP analysis of tennis that we are aware of is by Norman (1985) who used DP to determine “whether to serve fast or slow on either or both serves at each in a game, and a simple policy is found” (p. 1985). Depending on the values of point outcome probabilities, Norman classifies three rules (F,F), (F,S) and (S,F) for whether the first and second serves in a point should be fast (F) or slow (S) to maximize the probability of winning the point.

⁸ Traditional game theory has little to say about “mental ability” since all players are equally rational and intelligent. In the context of our model, these increases in win rates result from a better mental approach to the game,

We do not advise elite pro servers to adopt our counterfactual best responses since they are pure strategies that the returner would likely learn and adapt to. However our approach is sufficient to test the hypothesis of Nash equilibrium play in tennis, since it leads to the strong prediction that *there exists no deviation strategy that strictly increase the server's win probability, taking the returner's strategy as given*. Indeed, the MPE hypothesis leads to a stronger implication, the *One Shot Deviation Principle*: there is no deviation at any stage of the dynamic game that results in a strict improvement in the server's win probability. Our analysis reveals many advantageous one shot deviations, and the DP strategy takes maximal advantage of all of them.

To gain insight into the reasons for suboptimal serve choices, we estimate three dynamic structural models of the directions chosen by the server involving increasing degrees of far-sightedness. The full DP model posits that the server uses backward induction to maximize the probability of winning the entire *service game*, which is effectively an infinite horizon problem because service games must be won by at least two points. The semi-DP model posits that the server solves a two period DP to maximize his probability of winning the current *point* taking into account the option value of a 2nd serve but ignoring the effect of winning or losing on the subsequent state of the service game. The fully myopic model posits that the server maximizes the probability of winning each *serve*, a completely static problem that ignores even the option value of the 2nd serve. The fully myopic model is typically rejected because we find very significant dynamics and differences in serve directions between first and second serves. In most cases the best fitting model is the semi-DP specification. Since the semi-DP specification is often the best fit (or nearly so), the suboptimal serve behavior we identify appears to be primarily driven by incorrect server beliefs (i.e. lack of rational expectations) of the strengths and weaknesses of the returner as captured in the POPs, rather than a failure to solve the full DP problem over the 36 non-terminal states of the game. In fact, the optimal serve strategies that we compute numerically by DP are typically rather simple and easy to describe verbally.

One potential shortcoming of our approach is that we only have *estimates* of the POPs rather than the *true* POPs. Estimation error in the POPs could result in spurious, upward biased, estimates of the win probability when we use these estimates to calculate a best response strategy

since the estimates assume the returners strategy and other aspects of the server's play are unchanged under the DP serve strategy, so relative physical ability is held constant.

instead of using the true POPs. To account for this, we derive an approximate probability distribution for the true POPs based on the observed data. We calculate the win probability for each of our three best response strategies (full DP, Semi-DP, and fully myopic) for a large random sample of POPs drawn from this distribution. This robustness exercise confirms our core finding: the full DP and semi-DP strategies result in a significant first order improvement in the distribution of win probabilities relative to our estimates of the serve strategies used by the elite pros we study.

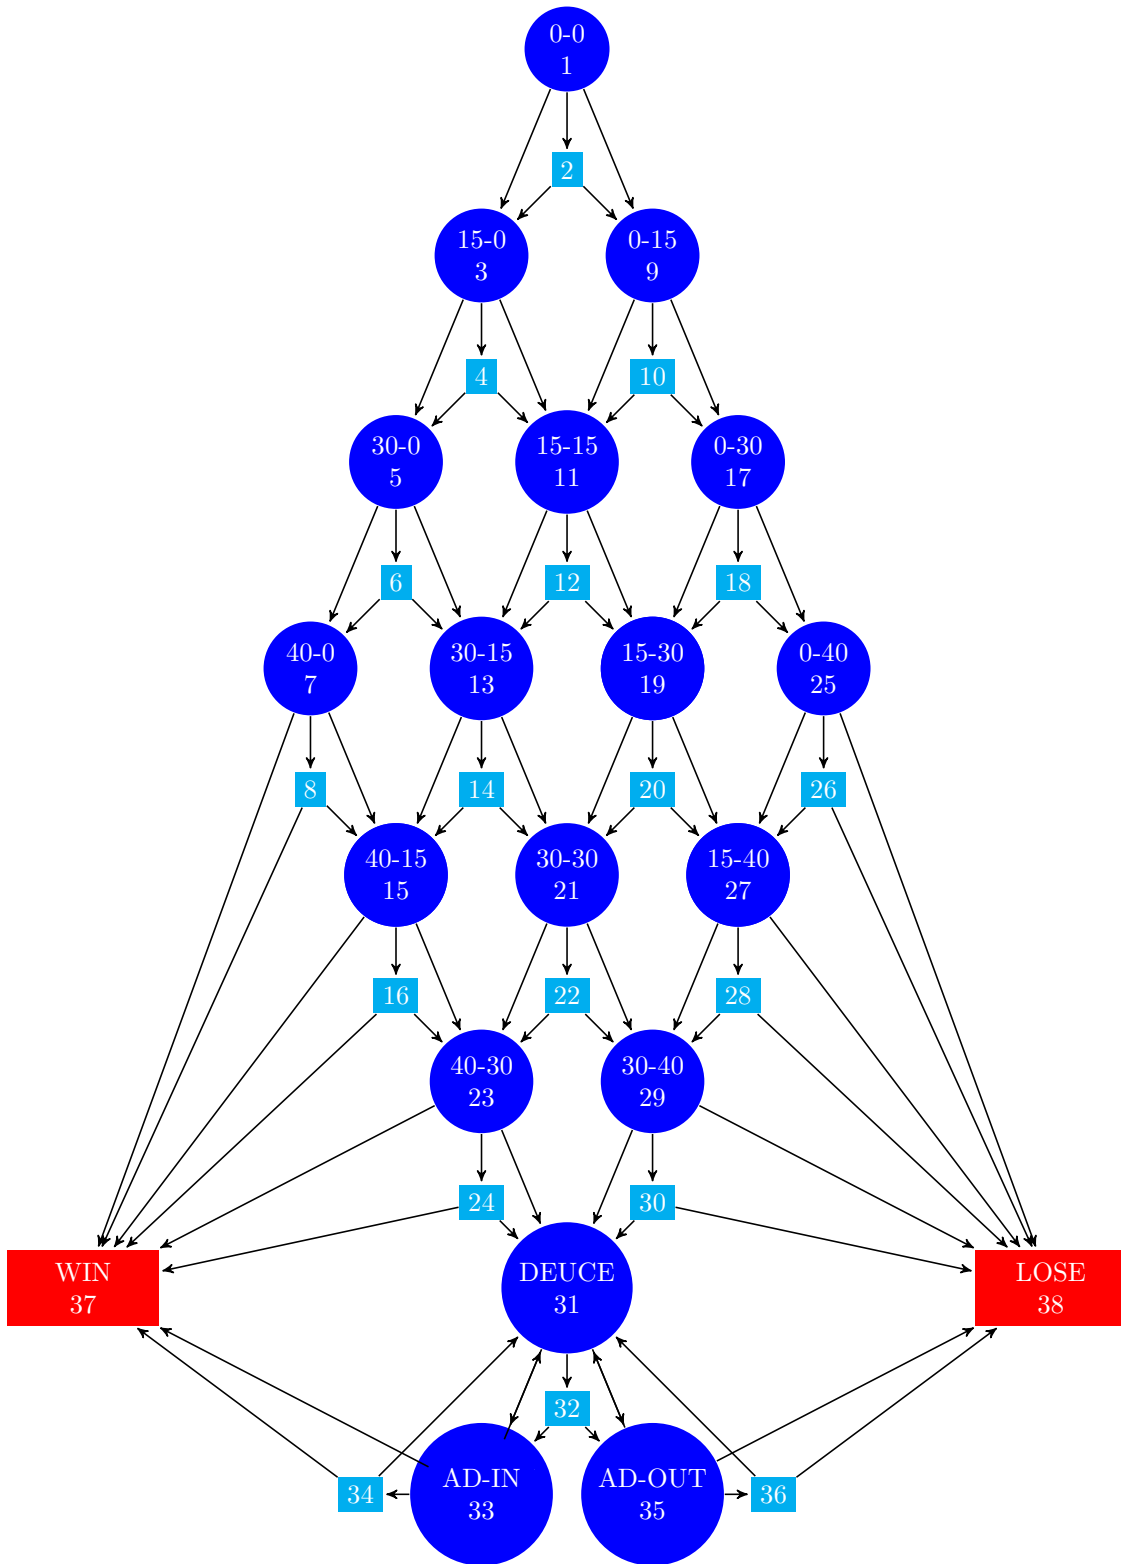
The paper is organized as follows. In section 2 we introduce our dynamic model of tennis and deduce the implied dynamic programming problem facing the server. In section 3 we summarize the key findings from our reduced-form empirical analysis of the MCP database, including our key finding: the strong rejection of the hypothesis of equal win probabilities for all serve directions. We also test stationarity across service games between fixed server-returner pairs in this section. In section 4 we present estimation results for the three structural models of tennis serve behavior discussed above, and perform the robustness test summarized in the last paragraph.

2 Modeling Tennis as a Dynamic Game

Tennis is two-player game between a *server* and a *returner* played in tournaments composed of *matches*. A match consists of a sequence of *sets*.⁹ A tennis set, in turn, is a sequence of *service games* where one of the two players is the server. The server is chosen by a flip of a coin in the first game, and the identity of the server alternates in each game thereafter. Typically, to win a set a player must win six games and be ahead by at least two games. Alternatively, if the score is tied six all, the set is decided by a tiebreak game in which the winner is the first to score seven points and be ahead by at least two. Each service game consists of a sequence of sub-games that are called *points*. A point consists of a first serve, plus an option for a second serve after a faulted, or missed, first serve. The service game ends when one of the players wins at least four points in total and at least two more points than their opponent.

⁹ Rules differ across tournaments, but often the player who wins the majority of 3 or 5 tennis sets wins the match.

Figure 1: Score states and transitions in the service game



2.1 Dynamic Theory of the Service Game

We use scalar x to track both the cumulative points scored by each player in the current service game and whether or not the server is making a first or second serve. Figure 1 is a directed graph of all the transitions for the *point-state* variable x within a service game. The circular nodes indicate first serves whereas the square nodes indicate second serves. The game starts in state $x = 1$ which corresponds to a first serve with conventional tennis score of $0 - 0$. If the server wins the point on the first serve, the point state transits to state $x = 3$, corresponding to a first serves at conventional score of $15 - 0$. If the server faults the first serve, the state transits to state $x = 2$, which is the second serve, and so forth. There are three possible transitions at every first serve node, and two possible transitions at all second serve nodes, and two absorbing states (i.e. terminal nodes where arrows only point in): the server wins the game ($x = 37$) or loses ($x = 38$).

For most nodes, the arrows connecting nodes are unidirectional, connecting to nodes in higher states x . However there are a subset of the nodes where the links connecting the nodes are bi-directional. These are the *end game states* or what is known in tennis as the *deuce end game*. The deuce endgame is reached by any path through the game tree where the players are tied after a total of 6 points have been won (by either player). At the state of Deuce ($x = 31$), the players are tied at $40-40$ and one of the players must win by two points in a row to win the overall game. If the server is ahead by 1 point in the deuce end game, the state ($x = 33$) is called “Advantage-in,” and when the server is behind by 1 point, ($x = 35$) it is called “Advantage-out”.

At every node, the server chooses the serve *type* $t = (s, d)$, where $d \in \{l, r, b\}$ indicates the *direction*: to the returner’s left l or right r , or directly into the returner’s body b , and $s \in \mathcal{S} \subset \mathbb{R}^2$ indicates the *speed* and *spin* of the serve (\mathcal{S} is non-empty, closed, and bounded). The returner, *anticipates* the direction choice of the server. Anticipation includes observable choices (ex. where to stand) and unobservable dimensions. We model anticipation with an attention vector, $(a^l, a^r, a^b) \geq 0$, where a^d denotes the attention the returner devotes to serve location d . We normalize the attention budget $a^l + a^r + a^b = 1$. We assume throughout that the serve direction choice weakly follows the choice of a . This captures the case in which a is a pure location choice, chosen strictly before the server chooses a direction, and the case in which a represents a

simultaneous pure mental choice of anticipation.¹⁰

The probability that a serve *lands in* ℓ (i.e. is not a fault) depends on the current state x and the serve type t , while the chance that the server wins the subsequent rally (conditional on serving in) ω depends on both the state x , the serve type t , and the attention vector a . In addition, we assume these probabilities depend on previous serve locations. For example, a server may have a lower chance of faulting a serve to the same location as the *previous* serve. Similarly, a returner may be better able to handle a serve to the same location as the previous serve. We term this the *muscle memory effect*, and model it by allowing ℓ and ω to depend on the location of the *two* previous first serves. We do this because first serves alternate between the deuce and ad courts. Also, we hypothesize that for first serves, the relevant previous serve direction is that of the previous first serve to the same court, whereas for second serves, the relevant direction is that of the faulted first serve. We initialize muscle memory to null, $m = (\emptyset, \emptyset)$, at the start of the service game, to $m \in \emptyset \times \{l, r, b\}$ after the first serve of the game, and $m \in \{l, r, b\}^2$, at any node in the game tree with at least two prior first serves. We assume that muscle memory is only updated after first serves at each point in the service game. If x is a first serve state (i.e. x is an odd number in the state numbering in figure 1), then $m = (d_{-2}, d_{-1})$ records the direction of the two previous first serves. If the serve is to direction d then muscle memory is updated to $m' = (d_{-1}, d)$ at the start of the next serve.

The key probabilities determining the score-state transitions in Figure 1 are $\ell(x, m, d, s)$ and $\omega(x, m, d, s, a)$, which we assume are continuous in (s, a) . We impose the following assumption for our theory and empirical analysis:

Assumption 1 (Stationarity I) *The functions ℓ and ω may vary across server-returner pairs, but do not vary over time (independent of (x, m)) or across service games.*

We assume each player’s objective is to win the service game,¹¹ normalizing the reward to winning the game to 1 and the reward to losing to 0. Since tennis cannot end in a draw, it is a

¹⁰ All results extend to a model in which the returner first chooses a subset of the unit triangle, and then chooses a specific element of this subset. This allows for the realistic case in which the physical location of the returner on the court constrains, but does not fully determine, the attention vector.

¹¹ It is not hard to show that it is WLOG to assume that players maximize chances of winning the service game. That is, that equilibrium values and strategies are identical to the case in which players maximize their chances of winning each match. This follows from our assumption that muscle memory resets at the start of each service game.

recursive constant sum game (as in Everett (1957)) which dramatically simplifies the analysis.

Use (σ_S, σ_R) to denote the server's and returner's strategies (perhaps mixed and arbitrarily history dependent) in the service game. Let $\mathcal{W}_S(x, m)$ be the set of probabilities that the server wins the game starting in state (x, m) induced by some pair of (not necessarily Markovian) sub-game perfect equilibrium strategies (σ_S^*, σ_R^*) for the server and returner. Appendix A proves:

Theorem 1 *All sub-games have a unique value (i.e. $\mathcal{W}_S(x, m)$ is a singleton), and there exists a Markov Perfect Equilibrium (MPE) in which strategies only depend on the current state (x, m) .*

2.2 Optimal Serve Strategies in the Induced Dynamic Program


Our empirical analysis uses Match Charting Project (MCP) data, which does not record speed, spin, or the location of the returner. To overcome this shortcoming, we use Theorem 1 to project any MPE into the induced dynamic programming (DP) problem facing a server choosing serve directions to maximize the chances of winning the service game. To do this, let $\rho(s|x, m)$ denote a Markov mixed strategy over the speed and spin vector $s \in \mathcal{S}$ for the server, and let $\alpha(a|x, m)$ denote a Markov mixed strategy over attention for the returner.

Definition 1 *If Assumption 1 holds and (ρ^*, α^*) are part of a MPE of the service game, then the Point Outcome Probabilities (POPs) Π are well defined conditional probabilities given by:*

$$\begin{aligned}\pi(\text{in}|x, m, d) &\equiv \int \ell(x, m, d, s) d\rho^*(s|x, m) \\ \pi(\text{win}|x, m, d) &\equiv \int \int \omega(x, m, d, s, a) d\rho^*(s|x, m) d\alpha^*(a|x, m)\end{aligned}$$

Notice that the mixing probabilities (ρ^*, α^*) will generally depend on the state of the game (x, m) , so the POPs will depend on (x, m) even if the underlying conditional probabilities ℓ and ω do not. Given any MPE strategies (ρ^*, α^*) , the probabilities π define a single agent “game against nature,” a dynamic optimization problem in which the server chooses a serve direction at each node in Figure 1 in order to maximize his chances of winning the service game. Figure 2 illustrates the extensive form of the *point game*; namely the subset of the larger directed graph starting at every odd point state x . In the point game, the server chooses a serve direction for the first serve, d_1 , and in the event of a fault, the direction of a second serve, d_2 . The point game ends with the server winning or losing a point at each pink node.

Figure 2: Detail on the point subgame of tennis



Following Norman (1985) we now describe the server's DP problem given π . Let $W_S(x, m)$ and $W_R(x, m)$ denote the *conditional win probabilities* in state (x, m) for the server and returner, respectively, assuming the server behaves optimally. Since tennis is a constant sum game, we have $W_S(x, m) + W_R(x, m) = 1$ for all states (x, m) ; and thus, it is sufficient for us to focus on calculating the win probability for the server W_S . Let $W_S(x, m, d)$ be the conditional win probability for the server assuming he serves to direction d on the current serve and behaves optimally on all following server. Finally let $x^+(x)$ and $x^-(x)$ denote the successor state in the event that the server wins the point or losses the point on the current serve, respectively.

The optimal serve strategy can be recursively calculated by DP using the following Bellman equations given by

$$W_S(x, m) = \max_{d \in \{l, b, r\}} W_S(x, m, d) \quad (1)$$

where $\{l, b, r\}$ denote serving to the returner's left, body or right side, respectively, and

$$\begin{aligned} W_S(x, m, d) &= \pi(\text{in}|x, m, d) [\pi(\text{win}|x, m, d)W_S(x^+(x), m') + [1 - \pi(\text{win}|x, m, d)]W_S(x^-(x), m')] \\ &+ [1 - \pi(\text{in}|x, m, d)]W_S(x + 1, m'), \end{aligned} \quad (2)$$

when x is a first serve state (i.e. x is one of the odd numbered circular nodes in Figure 1), and

$$W_S(x, m, d) = \pi(\text{in}|x, m, d) [\pi(\text{win}|x, m, d)W_S(x^+(x), m) + [1 - \pi(\text{win}|x, m, d)]W_S(x^-(x), m)] \\ + [1 - \pi(\text{in}|x, m, d)]W_S(x^-(x), m), \quad (3)$$

when x is a second serve state (i.e. x is one of the even numbered square nodes in Figure 1). The optimal serve strategy, denoted by $D_S^*(x, m)$, is the set of serve directions that maximize the win probability

$$D_S^*(x, m) = \underset{d \in \{l, b, r\}}{\text{argmax}} W_S(x, m, d). \quad (4)$$

A necessary condition for a mixed serve strategy is that $D_S^*(x, m)$ contains more than one serve direction. In particular, the server will only mix across all three locations if $W_S(x, m, l) = W_S(x, m, b) = W_S(x, m, r)$. The Bellman equation can be written more compactly as

$$W_S = \Gamma(W_S) \quad (5)$$

where Γ is the Bellman operator implicitly defined in equations (1), (2) and (3) for non-terminal first and second serve states, respectively.

While most service games are reasonably short in practice (fewer than 10 points), there is no fixed upper bound on the duration of the *deuce end game*, the sub-game starting at $x = 31$.¹² For example, if the server wins the point game at deuce ($x = 31$), then the score state transitions to ad-in ($x = 33$), and if the server loses the point at ad-in, the score state transitions back to deuce. Formally this dynamic programming problem is an infinite horizon single agent *directional dynamic game* (DDG) as defined by Iskhakov, Rust, and Schjerning (2016).¹³ There are two key elements of their approach to solving such games: use *state recursion* rather than recursion over “time” (aka serves), and solve separately for values over collections of states for which directionality is not present, in our case the deuce endgame, and then solve the overall game by backward induction across directionally connected collections of states.

¹² The longest deuce endgame that we are aware of was between Anthony Fawcett and Keith Glass in 1975. The score reverted back to deuce 37 times before Glass won the game, although Fawcett won the match.

¹³ Norman (1985) recognized the directionality of tennis and grasped the essence of state recursion when he described how the optimal tennis serve strategy and corresponding win probabilities could be calculated by DP: “One way to compute the optimal policy would be to compute the optimal decisions and state values for the six-non-absorbing states in the loop, using the method of approximation in policy space, and then to compute the optimal decisions and state values for the remaining 28 states using the method of computation for directed states” (p. 75–76). Actually the correct number of remaining states is 30, not 28, since there are a total of 36 non-terminal nodes in state transition graph in Figure 1.

To apply this approach to our game, note that score states $x \in \{1, \dots, 30\}$ are transient directed states (i.e. once a transition occurs from a transient state x the game never returns to it). The deuce end game is the only collection of states for which directionality is not present. Let W_S^e denote the value functions restricted to the 6 possible deuce endgame states $x \in \{31, \dots, 36\}$. Since the Bellman equation (5) still holds for this subset of the state space, we can solve for the fixed point $W_S^e = \Gamma(W_S^e)$.¹⁴ Even though there is no discount factor in the Bellman equations for tennis, the probability of winning serves plays an equivalent role and ensures that Γ is a contraction mapping with a unique fixed point W_S^e . Since the state transition function is monotone for all $x \in \{1, \dots, 30\}$, i.e. $x^+(x) > x$ and $x^-(x) > x$, once the win probabilities for the deuce end game states have been calculated, the win probabilities can be calculated for these score states by backward induction over states.

2.3 Calculating Win Probabilities for Stationary Serve Strategies

The theoretical model above assumes the unobserved elements of choice (speed, spin, returner location, etc.) constitute a MPE. While this is sufficient for our empirical analysis, it is not necessary. Instead we often make the following assumption directly on these probabilities.

Assumption 2 (Stationarity II) *The actual POPs (those implied even if players are not using MPE strategies) are given by families of conditional probabilities $\{\pi(in|x, m, d), \pi(win|x, m, d)\}$ that may depend on the server and returner but do not vary over time (independent of (x, m)) or across service games.*

While we have used the same notation here, as we did in Definition 1, we stress that we are not imposing equilibrium behavior in this Assumption. Formally, this new stationarity assumption is weaker than Definition 1; namely, Assumption 1 and MPE strategies. But Assumption 2 does impose implicit restrictions on the nature of tennis that rule out effects discussed in the introduction such as whether players can recognize and adapt to changes in strategies of their opponent. Assumption 2 rules out learning effects by players, or the possibility of multiple equilibria if selection of different equilibria across service games creates non-stationarity in Π . While Assumption 2 does not impose equilibrium behavior, it does implicitly assume that the players are

¹⁴Since there are 9 possible muscle memory states in $\{l, r, b\}^2$, the solution W_S^e will be $9 \times 6 = 54$ probabilities.

unaware if they are failing to play mutual best responses, since otherwise they would have an incentive to alter their strategies to gain an advantage, touching off a learning and adaptation process that would likely violate stationarity.

When stationarity holds and we have enough data, we can consistently estimate Π to recreate the server's environment and use DP to numerically calculate best response serve strategies. We then compare the DP strategies to the ones servers actually use (which can also be consistently estimated when we have sufficient observations on serve directions). In order to compare optimal strategies to observed strategies, we must calculate win probabilities given some arbitrary (potentially suboptimal) Markovian serve strategy $P(d|x,m)$, which is the probability that the server chooses direction d given the current state (x,m) . Let $W_P(x,m)$ denote the probability that a server employing strategy P will win the game conditional on the current state being (x,m) . First, similar to the Bellman equation (1), we can define $W_P(x,m)$ in terms of a new set of win probabilities we refer to as *conditional win probabilities* $W_P(x,m,d)$, which is the conditional probability of winning the game if the server is in state (x,m) and chooses serve direction d . We have, analogous to the Bellman equation (1)

$$W_P(x,m) = \sum_{d \in \{l,b,r\}} W_P(x,m,d)P(d|x,m). \quad (6)$$

The conditional win probabilities $W_P(x,m,d)$ are given by the same Bellman equations (2) and (3) above when we substitute W_P in place of W_S . These equations, plus (6) make it clear that W_P is actually an implicit function of the POPs, Π , and the serve strategy, P .¹⁵ In fact, we can write an expression for W_P as the solution to a system of linear equations, as is well known in the dynamic programming literature on *policy evaluation*. Since there are 298 distinct states (x,m) ,¹⁶

$$W_P = w_P(P,\Pi) + M_P(P,\Pi)W_P, \quad (7)$$

where $w_P(P,\Pi)$ is a 298×1 vector providing the probability of directly winning the service game in each state (in most states this is zero), and $M_P(P,\Pi)$ is a 298×298 Markov sub-transition

¹⁵ In our econometric analysis we will estimate models under two different higher level hypotheses about Π : *rational expectations*: the server's beliefs about the environment correspond to our consistent estimates of the POPs, and *subjective beliefs*: Π embodies the server's subjective beliefs about his/her own abilities and the ability/strategy of the receiver that may or may not coincide with objective reality.

¹⁶ There is only one possible muscle memory state at the start of the service game $x = 1$, three possible muscle memory states for $x = 2, 3$, and 9, and 9 possible muscle memory states for the remaining 32 score states. Thus, $1 * 1 + 3 * 3 + 32 * 9 = 298$ states (x,m) .

matrix (i.e. not all of its rows sum to 1)¹⁷ representing the probability of transiting from any given state to any new state under the transition law for tennis induced by the serve strategy P and the POPs Π . Since $M_P(P, \Pi)$ is a Markov sub-transition matrix, $\|M_P(P, \Pi)\| < 1$, where $\|\cdot\|$ denotes the linear operator or matrix sup-norm, which implies that the linear system (7) has a unique solution W_P . We can see from (7) that W_P is an implicit function of both (P, Π) and we use this result later in the paper to rapidly calculate win probabilities, and via the Implicit Function Theorem, the gradients of the win and conditional win probabilities with respect to model parameters. This enables us to compute standard errors for win probabilities and conduct efficient Wald tests of the hypothesis of equal win probabilities implied by the existence of a unique mixed strategy equilibrium.

Given sufficient number of observations of service games between a given server and returner, $W_P(x, m, d)$ can be consistently estimated as the fraction of service games the server won, when the server chose direction d in state (x, m) . And as long as the POPs and server strategies are stationary and Markovian, the W_P so estimated will obey identity (7). However, since there are 298 possible states (x, m) , we would need to estimate $298 * 3 = 894$ separate conditional probabilities $W_P(x, m, d)$. In order to estimate 894 probabilities with sufficient precision to have adequate power to test the hypothesis of equal win probabilities for all serve directions at all states (x, m) , we conjecture that we would need roughly 10000 service games. Unfortunately in our data set we typically have only 100 to 200 service games per server-returner pair.

Our analysis also requires an estimate of the actual Markovian serve strategy P and the observed POPs Π . An unrestricted or non-parametric estimate of P involves 298 parameters, and for Π a total of $894 * 2 = 1788$ parameters. While we can exploit identity (7) to estimate all three vectors W_P , P and Π using far fewer than $894 + 298 + 1788 = 2980$ parameters, direct non-parametric estimation of these conditional probabilities is not feasible given the size of our data set. To overcome this data limitation, we introduce reduced form “parametric models” for serve probabilities and the POPs in the next section. These are flexibly parameterized logit models for P and Π that depend on far fewer parameters than would be required to estimate these probabilities non-parametrically. In the next section we will introduce flexible parametric models of P and

¹⁷ The rows of M_P do not all sum to 1, due to the probability of directly winning the game, as captured by $w_P(P, \Pi)$.

Π that depend on only 44 parameters. Further, while the non-parametric estimates of W_P use only the binary win/loss outcome for each game, our reduced-form model is estimated using the much greater amount of information generated from a tennis game, namely the observations of all serve directions and serve outcomes. This effectively brings far more information to bear, allowing us to obtain precise estimates of the parameter of flexible specification for P and Π ; using these, we can then derive conditional win probabilities via identity (7) to construct much more powerful *omnibus tests* of the hypothesis of equal win probabilities over all states (x, m) simultaneously.

2.4 The Monotonicity Condition and Myopic Optimality

We conclude this section by discussing the decomposition result of Walker, Wooders, and Amir (2011) (WWA) and its implications for whether serial independence in serve directions is or is not a key implication of a mixed strategy Nash equilibrium. WWA showed that tennis is in the class of *binary Markov games* which are two player constant-sum games with only two possible outcomes, both for the overall game and all of its component subgames. The state of a binary Markov game advances via the earning of points at each subgame, or *point subgame*. WWA denote the payoff functions in the point subgames, *point game payoff functions*, which are analogous to our POPs, but in their model the point game payoff functions are independent of the current score and prior choices (i.e. there are no muscle memory effects). They appeal to the von Neumann (1928) Minimax Theorem to establish the existence of a mixed strategy equilibrium to each point subgame. They define a *minimax-stationary strategy* for the overall game as one where the mixed strategy/minimax equilibria of the point subgame is played in each point game. Thus, the minimax-stationary strategy is myopic, focusing only on winning each point. WWA ask: under what conditions does the minimax-stationary strategy coincide with the MPE of the overall game?

The key condition is the *monotonicity condition* (MC). Absent muscle memory effects, the value functions $W_S(x)$ and $W_R(x)$ only depend on x . Normalizing the payoff from winning the game to 1 and the payoff from losing to 0, WWA have $W_S(x) + W_R(x) = 1$ for all states x ; and thus, it suffices to describe the equilibrium in terms of player S . Let x be an odd (first serve) non-terminal score state, and abuse notation and let $x^+(x)$ be the successor state after a win in the

current point game and $x^-(x)$ be the successor state after a loss in the current point game, then WWA’s MC for all non-terminal odd states x is:

$$W_S(x^+(x)) > W_S(x^-(x)) \quad (8)$$

Using the MC and the assumption that $0 < W_S(x) < 1$ for each non-terminal state x , WWA prove that the minimax-stationary strategy is a MPE for the overall dynamic game, and the value functions W_S and W_R satisfy the conditions for the von Neumann (1928) Minimax Theorem, viewing the game in normal form. They also establish the converse of this, namely that if each non-terminal point game has a unique Nash equilibrium then the only MPE of the overall game is the one given by the unique minimax-stationary strategies.

We refer to WWA’s theorem as a *decomposition result*, since when it holds the overall game decomposes into independent static subgames that be solved “myopically” by assuming the players only consider the short term objective of maximizing the probability of winning each point, without concern for how their actions impact the subsequent state of the game. WWA’s decomposition result can be extended to our context with muscle memory, provided a *generalized monotonicity condition* (GMC) holds for all non-terminal odd score states x :

$$W_S(x^+(x), (m_2, d)) > W_S(x^-(x), (m_2, d')) \quad \forall m \text{ and } d, d' \in \{l, r, b\}. \quad (9)$$

In words, this states that the server would rather win a point than lose a point, regardless of the future impact on the muscle memory state of the current serve direction. Intuitively, GMC will hold provided muscle memory effects are small. Unfortunately is not easy to establish verifiable sufficient conditions for GMC to hold *a priori*. Generally one must first solve for a MPE using more general methods that do not require the GMC, such as the RLS algorithm of Iskhakov et al. (2016), and then check to see if GMC holds.¹⁸ In section 4 we show, using empirically estimated POPs and solving for optimal serve strategies using state recursion (i.e. without assuming the GMC holds), that there are server/returner pairs for which GMC fails. In these cases, the myopic solution strategy of solving each point subgame independently without regard for the future state of the game is suboptimal, though we show that typically the cost of suboptimality in terms of reduced win probability is small.

¹⁸ Recall that values W_S are unique in our hypothesized model by Theorem 1. Thus, if GMC fails in one MPE, it will fail in all MPEs.

Consider the implications for serial independence of serve directions. If there are no muscle memory effects and MC holds, then WWA's decomposition result implies that serve directions are conditionally independent random variables. That is, the server uses the same mixed strategy to select her current serve direction on all first serves and the same (generically different) mixed strategy on all second serves, independent of the history of play. WW use this result to conclude that "In addition to equality of players' winning probabilities, equilibrium play also requires that each player's choices be independent draws from a random process" (p. 1522). Of course, this independence result is a consequence of their *assumption* that the mapping from strategies to outcomes are identical across all point games, independent of previous choices and outcomes. When there is history dependence due to effects such as muscle memory m , the mixed strategy equilibria of each point game can depend on both x and m . We show in Appendix B that serial correlation in serve locations is (robustly) consistent with MPE. Thus, we conclude that serial independence of serve directions is generally *not* a testable implication (i.e. necessary condition) of a mixed strategy Nash equilibrium.

In summary, our empirical analysis in the next two sections focuses on testing three key general implications of game/decision theory on the behavior of professional tennis servers:

1. **Nash equilibrium:** there should not be any other serve strategy that increases the server's probability of winning
2. **Mixed strategy equilibrium:** the probability of winning in state (x, m) should be equal for all serve directions chosen with positive probability in state (x, m) .
3. **Option value of the second serve:** the serve strategy for the first and second serves should differ.

We also test the following behavioral implications of GMC (9) and muscle memory:

4. **Optimality of myopic serve strategies:** When GMC holds, it is optimal for the server to adopt a myopic strategy that focuses only on the goal of maximizing the probability of winning each point
5. **Serial independence:** If GMC holds and there are no muscle memory effects in the POPs, the direction of a first serve should not depend on the direction of any previous first serve.

3 Reduced-form analysis of serve strategies

In this section we start with a “model-free” descriptive analysis of our data, and then introduce a flexible “reduced form” model of tennis that we use to test several of the key implications of game theory summarized in section 2, particularly the implication that conditional win probabilities are the same for all serve directions.¹⁹ Most of our analysis focuses on a set of elite professional tennis players, who have all been ranked number one in the world and won multiple Grand Slams. These players are Roger Federer, Rafael Nadal, Novak Djokovic, Andy Murray, Pete Sampras, and Andre Agassi.²⁰ We focus on these players for two reasons: First, we have the most observations for them, and second, if we can show that they serve suboptimally, that means even the best of the best are susceptible to strategic errors.

3.1 Analysis of play of specific server-returner pairs

We have sufficient observations to analyze serve decisions of specific server-returner pairs. Table 1 summarizes some of the key statistics for our elite server-returner pairs, and it reveals a great deal of player-specific heterogeneity that would be masked in pooled statistics. The Table presents the total number of service games and the number of serves we observe for each pair. A typical service game ends after 7 to 9 serves. The Serve column breaks down the total number of serves we observe into first and second serves. We can see that the “crude fault rate” (fraction of total serves that are 2nd serves) differs across servers, ranging from a low of 21% for Nadal serving to Federer to a high of 30% for Sampras serving to Agassi.

The three columns labelled L, B and R provide the fraction of first and second serves to the returner’s left, body, and right for each server. We see that in general servers use mixed strategies, but the mixing probabilities for second serves differ significantly from the first serve. The last column of the table includes the P-value of a likelihood ratio test of the null hypothesis that the mixing probabilities for the first and second serves are equal. We see that for all servers, we can decisively reject this hypothesis. In general we see that the fraction of body serves is significantly

¹⁹ Our analysis is not “assumption-free” however, as we maintain Assumption 2, stationarity, for the validity of our statistical tests and estimates. See section 2.

²⁰ Unfortunately, we do not have enough observations of top female players to estimate our models on them.

Table 1: Win probabilities and mixed serve strategies for selected elite server-returner pairs

Server → Returner	Games, serves Serves/game	1st serves 2nd serves	Serve directions			Win prob (std)
			L	B	R	P-value: $P_1 = P_2$
Roger Federer → Rafael Nadal	523, 4732 8.36	3208 1164	.4402 .2174	.1007 .2698	.4592 .5129	.7686 (.0184) 5.1×10^{-60}
Rafael Nadal → Roger Federer	519, 4081 7.86	3227 854	.6616 .5937	.2048 .3208	.1336 .0855	.8092 (.0172) 6.3×10^{-12}
Roger Federer → Novak Djokovic	411, 3501 8.52	2524 977	.4521 .4084	.0939 .3408	.4540 .2508	.8200 (.0190) 6.7×10^{-68}
Novak Djokovic → Roger Federer	407, 3653 8.98	2696 957	.4640 .4389	.1565 .3365	.3795 .2247	.8010 (.0198) 1.0×10^{-33}
Rafael Nadal → Novak Djokovic	346, 2937 8.49	2230 707	.3964 .4073	.2825 .5403	.3211 .0523	.7197 (.0241) 2.4×10^{-64}
Novak Djokovic → Rafael Nadal	356, 2877 8.08	2149 728	.4067 .1484	.1619 .2940	.4314 .5577	.7528 (.0222) 1.2×10^{-40}
Novak Djokovic → Andy Murray	230, 1958 8.51	1447 511	.4651 .2192	.1244 .4618	.4105 .3190	.7696 (.0278) 3.0×10^{-53}
Andy Murray → Novak Djokovic	230, 2141 9.31	1522 619	.3863 .4233	.0841 .4782	.5296 .0985	.7435 (.0288) 5.8×10^{-122}
Pete Sampras → Andre Agassi	140, 1275 9.11	884 391	.4434 .4680	.0724 .1765	.4842 .3555	.9000 (.0254) 5.3×10^{-8}
Andre Agassi → Pete Sampras	135, 1125 8.33	825 300	.5127 .5766	.1115 .2700	.3758 .1533	.8666 (.0293) 7.2×10^{-16}

higher, often more than double, for the second serve relative to the first serve.

We also see that servers adjust their serve strategy for different returners. For example from table 1 we can see that Nadal uses a very different serve strategy when serving to Federer compared to when he is serving to Djokovic. The final column of Table 1 shows the empirical game win probability for the server and its estimated standard error (i.e. the fraction of games the server won). We see quite a bit of variation in service game win probabilities across different server-returner pairs, ranging from a low of 72% for Nadal serving to Djokovic, to a high of 90% for Sampras serving to Agassi. Even controlling for the same server, we see a fairly big variation in win probabilities depending on the returner: for example, Nadal has an 81% service game win probability when serving to Federer, as Federer is a weaker returner than Djokovic. Given the relatively small standard deviations in estimated win probabilities, we can strongly reject the null

hypothesis that variation in estimated win probabilities is due to sampling error.

The evidence presented so far seems consistent with the predictions of game theory; servers use mixed strategies, these strategies differ across first and second serves (reflecting the effect of the option value of the 2nd serve discussed in section 2); and servers appear to adjust their serve strategy to exploit the relative weaknesses of their opponents. The variation in win probabilities across server-returned pairs reflects differences in relative physical abilities of different players.

3.2 A flexible, agnostic reduced-form probability model of tennis

In order to test the key necessary condition for a mixed strategy equilibrium – equality of win probabilities for all serve directions – a deeper econometric analysis is required. As we suggested in section 2, we do this by estimating a flexibly parameterized reduced-form specification for serve strategies $P(d|x, m)$ and the POPs ($\pi(in|x, m, d), \pi(win|x, m, d)$). Let $f(x, m, d)$ be a $1 \times K_P$ vector of indicators for various subsets of the state/action space. We will describe specific choices for f below. In general, f will partition the state space into different subsets where serve direction probabilities are similar. Let θ_P be a conformable $K_P \times 1$ vector of coefficients to be estimated. Then we propose the following flexible logit model to approximate serve probabilities

$$P(d|x, m, \theta_P) = \frac{\exp\{f(x, m, d)' \theta_P\}}{\sum_{\delta \in \{l, b, r\}} \exp\{f(x, m, \delta)' \theta_P\}} \quad (10)$$

Similarly let $g_{in}(x, m, d)$ and $g_{win}(x, m, d)$ be $1 \times K_{in}$ and $1 \times K_{win}$ vectors of indicators used to define the following binary logit models for $\pi(in|x, m, \theta_{in})$ and $\pi(win|x, m, \theta_{win})$ that depend on parameter vectors ($\theta_{in}, \theta_{win}$):

$$\pi(in|x, m, d, \theta_{in}) = \frac{\exp\{g_{in}(x, m, d)' \theta_{in}\}}{1 + \exp\{g_{in}(x, m, d)' \theta_{in}\}} \quad (11)$$

$$\pi(win|x, m, d, \theta_{win}) = \frac{\exp\{g_{win}(x, m, d)' \theta_{win}\}}{1 + \exp\{g_{win}(x, m, d)' \theta_{win}\}} \quad (12)$$

We estimate the parameter vector $\theta = (\theta_P, \theta_{in}, \theta_{win})$ by maximum likelihood using the log-likelihood function $L(\theta)$ given by

$$L(\theta) = \sum_{g=1}^G \sum_{s=1}^{S_g} [\log(P(d_{s,g}|x_{s,g}, m_{s,g}, \theta_P)) + \log(f(o_{s,g}|x_{s,g}, m_{s,g}, d_{s,g}, \theta_{in}, \theta_{win}))], \quad (13)$$

where G is the total number of service games observed for a particular server-returned pair, S_g is the number of serves in game g , and $(d_{s,g}, x_{s,g}, m_{s,g})$ is the observed direction of serve, game state

and muscle memory state at serve s in game g . The variable $o_{s,g}$ is the outcome of serve s of game g and takes three possible values: $o_{s,g} = 1$ if the serve is in (not faulted) and the server wins the subsequent rally, $o_{s,g} = 2$ if the serve is in and the server loses the subsequent rally, or $o_{s,g} = 3$ if the serve faulted. In all non-terminal first serve states (odd values of x) the game state transits to a second serve in the event that $o = 3$, but in any second serve state the server loses the point when $o = 3$ (i.e. the server “double faults”). The conditional probability $f(o|x, m, d, \theta_{in}, \theta_{win})$ is defined in terms of the POPs as follows

$$f(o|x, m, d, \theta_{in}, \theta_{win}) = \begin{cases} \pi(in|x, m, d, \theta_{in})\pi(win|x, m, d, \theta_{win}) & \text{if } o = 1 \\ \pi(in|x, m, d, \theta_{in})[1 - \pi(win|x, m, d, \theta_{win})] & \text{if } o = 2 \\ 1 - \pi(in|x, m, d, \theta_{in}) & \text{if } o = 3 \end{cases} \quad (14)$$

We evaluated different specifications for these models that partition the state/action space in different ways. By using increasingly fine partitions, the models above encompass the unrestricted or “non-parametric” specifications for (P, Π) . However, as we noted in section 2, a completely unrestricted specification has 2512 parameters, and we can see from Table 1 that we do not have sufficient observations to reliably estimate an unrestricted model for most server-returned pairs in our data set. Thus, we face a classic tradeoff between a desire to have the most flexible possible model with many parameters, and the desire to have sufficiently many observations per parameter estimated to guard against the possibility of “overfitting” where a few outlier observations could distort key parameter estimates.

We manage this tradeoff using model selection techniques, particularly the Akaike Information Criterion (AIC) which penalizes model complexity. Specifically, we have $AIC = 2[K - L(\hat{\theta})]$ where K is the total number of parameters estimated in a given model, $L(\hat{\theta})$ is the maximized value of the log-likelihood function and $\hat{\theta}$ is the maximum likelihood estimate of the parameters of the particular model. We evaluated several different models (i.e. choices for f , g_{in} and g_{win} with different numbers of parameters and different partitions of the state space) and chose as our preferred specification the model with the smallest AIC.²¹

²¹ We also evaluate models in terms of the Bayesian Information Criterion $BIC = K * \log(n) - 2L(\hat{\theta})$ which has a stronger penalty for model complexity, but we found that the higher complexity penalty caused BIC to select models with fewer parameters. In cases where one model specification was nested within another encompassing specification, BIC would choose the more parsimonious restricted specification even though likelihood ratio tests lead us to reject the parsimonious restricted specification relative to the less restricted encompassing model.

Our preferred specification still involves a large number of parameters per server-receiver pair (44 parameters to be exact with 12 parameters θ_P determining serve strategies, P , and 16 parameters each for $(\theta_{in}, \theta_{win})$ that determine the POPs). Unfortunately we do not have the space to present all these parameter estimates and the associated standard errors for each of the 12 server-receiver pairs we analyzed, though we are happy to provide them to interested readers on request. As we will describe further in the next sections, our preferred specification balances the tradeoff described above: it provides an accurate probability model of the entire service game for individual server-receiver pairs while avoiding the dangers of overfitting. In the remainder of this section we will use this model to test several of our key assumptions, including the key hypothesis of Nash equilibrium play in tennis.

3.3 Testing the stationarity assumption

We now test necessary implications of stationarity (Assumption 2). In particular, stationarity implies that the stochastic process of serves and serve outcomes in any given tennis game between a given server and returner on a given type of court (in our case, hard courts), is Markovian and the realizations of these Markov processes are *IID* across successive service games. That is, while the presence of muscle memory and the scoring rules of tennis imply that the sequence of serve directions and serve outcomes in a given service game will generally be serially correlated, there will be no dependence across successive service games. This is because we assume that muscle memory is “reset” across successive service games (to $m = \emptyset$ at the start of each game), so there are no effects linking serve choices and serve outcomes across successive service games. Further, stationarity implies that the stochastic process $\{d_t, o_t, m_t, x_t\}$ of serve directions d_t , serve outcomes o_t , muscle memory state m_t , and score state x_t for each service game is a stationary Markov process with transition density $g(d_{t+1}, o_{t+1}, m_{t+1}, x_{t+1} | d_t, o_t, m_t, x_t)$ given by

$$g(d_{t+1}, o_{t+1}, m_{t+1}, x_{t+1} | d_t, o_t, m_t, x_t) = P(d_{t+1} | x_{t+1}, m_{t+1}) I\{x_{t+1} = T(x_t, o_{t+1})\} I\{m_{t+1} = M(m_t, d_t) f(o_{t+1} | x_t, m_t, d_t)\} \quad (15)$$

where M is a deterministic updating rule for muscle memory, such that $m_{t+1} = (d_t, d_{t-1}) = M(d_{t-1}, d_{t-2}, d_t)$ via the “right shift” operator, and $T(x, o)$ is the deterministic transition rule encoded in the tennis score transition directed acyclic graph given in Figure 1 of section 2. Note

that this transition density also satisfies a conditional independence restriction, i.e. g can be written as $g(d_{t+1}, o_{t+1}, m_{t+1}, x_{t+1} | d_t, m_t, x_t)$ since given the last score x_t , the previous outcome o_t has no effect on serve directions, muscle memory, or the new serve outcome o_{t+1} .

It is easy to think of reasons why these independence, Markovian, stationarity, and conditional independence restrictions implicit in our model formulation may not hold. For example, if a server injures his shoulder, this can persistently affect the POPs (and thus the outcome transition density $f(o'|x, m, d)$, which depends on the POPs) in an adverse way. Also there might be psychological effects within a tournament, such as confidence or a “hot hand,” that could lead to serial correlation across successive service games served by the same player. Finally, if a player is learning and adapting, his strategy may slowly evolve as he learns more about his opponent’s weaknesses and trains to exploit them.

In light of this, why do we make the stationarity assumption? Due to limited numbers of observations, of course! We need to pool over successive service games to have a sufficient number of observations to estimate a sufficiently unrestricted and flexibly parameterized reduced form model of serves $P(d|x, m)$ and the POPs $\{\pi(in|x, m, d), \pi(win|x, m, d)\}$. From the previous section, our preferred reduced form model has a total of 44 parameters under the muscle memory specification and 32 parameters under the no muscle memory specification. Given that a typical service game lasts for about 8 to 9 serves, we need at least 100 service games of data to estimate these 44 or 32 parameters with any semblance of accuracy. We are particularly concerned with the issue of *overfitting*, along with the possibility that the model’s predictions of serves to particular directions or particular outcomes will have incredibly high or low probabilities due simply to the lack of sufficient observations to estimate each cell, or parameter.

However, since the stationarity assumption is testable, we present results from a particularly simple way of testing for stationarity in Table 2. For a subset of the server-returned pairs where we had the most full game observations (over 300), we estimate separate reduced form models. Since service games are ordered chronologically over the players’ careers, we estimate separate reduced form models using the first 100 and the last 100 service games, respectively. Then we estimate a “pooled” model using 200 games and calculated likelihood ratio test statistic. The “unrestricted” model is one that sums the separate maximized log-likelihood values for the first and last 100 observations, each estimated separately. The unrestricted model has a total of $2*44=88$ (with

Table 2: Tests for stationarity of $P(d|x, m)$ and $(\pi(in|x, m, d, \theta_{in}), \pi(win|x, m, d, \theta_{win}))$

Server → Returner	No muscle memory			Muscle Memory		
	Restricted	Unrestricted	LR test	Restricted	Unrestricted	LR test
	LL, AIC	LL, AIC	df, P value	LL, AIC	LL, AIC	df, P value
Roger Federer → Rafael Nadal	-3132.5 6329.1	-3110.4 6348.8	32 .074	-3104.7 6297.5	-3077.6 6331.2	44 .138
Rafael Nadal → Roger Federer	-3047.8 6159.7	-3012.36 6152.7	32 8.9×10^{-5}	-3043.61 6175.2	-2995.9 6167.8	44 1.2×10^{-5}
Roger Federer → Novak Djokovic	-3352.99 6770.0	-3320.0 6768.1	32 3.9×10^{-4}	-3325.7 6739.4	-3285.7 6747.3	44 7.2×10^{-4}
Novak Djokovic → Roger Federer	-3627.1 7318.2	-3592.3 7312.62	32 1.3×10^{-4}	-3600.6 7289.3	-3559.2 7293.5	44 3.6×10^{-4}
Rafael Nadal → Novak Djokovic	-3577.0 7218.0	-3569.0 7266.0	32 .992	-3566.4 7220.8	-3554.8 7285.7	44 .996
Novak Djokovic → Rafael Nadal	-3152.3 6368.6	-3136.29 6400.6	32 .466	-3131.3 6350.5	-3111.7 6399.4	44 .683

muscle memory) or $2*32=64$ (without muscle memory) parameters that are estimated separately without placing any equality restrictions across the two sample subsets. The restricted model imposes an equality restriction that the same parameters (and thus the same serve probabilities and POPs) hold for the first and last 100 observations, respectively. Thus, the LR test has 32 total degrees of freedom for the specification without muscle memory, and 44 total degrees of freedom for the specification with muscle memory.

From Table 2, we see some mixed results. We are not able to reject our test of the stationarity assumption at the 5% critical level for Federer serving to Nadal, Nadal serving to Djokovic, or Djokovic serving to Nadal. But for the other three pairs the test does reject the restrictions implied by stationarity. However, we also calculate the AIC criterion for these models and present these values in the table as well. We see that the AIC chooses the restricted specification with muscle memory for all but one of the server-returner pairs: Nadal serving to Federer. Ultimately, we are more concerned about overfitting and the danger of spurious variability in estimated POPs (which could lead to spurious rejection of the hypothesis of equal win probabilities) than we about the possibility that non-stationarity in server behavior or the POPs across successive games could bias or invalidate our results. Consequently, we assume stationarity is a reasonable assumption,

which enables us to pool sufficient numbers of service games to get the most reliable possible estimates of serve probabilities and the POPs.

3.4 Testing for equality of win rates by serve direction

As we noted in section 2, once we have the serve probabilities P and POPs Π , we can calculate the implied win probabilities and conditional win probabilities using equations (2), (3), and (6). In this section we use this approach to construct an omnibus test of the hypothesis of equal win probabilities for all serve directions by testing the equality restrictions

$$W_P(x, m, l) = W_P(x, m, b) = W_P(x, m, r) \quad \forall (x, m). \quad (16)$$

Since W_P is an implicit function of (P, Π) , which are in turn nonlinear functions of the parameters $\hat{\theta}_{rf} = (\hat{\theta}_P, \hat{\theta}_{in}, \hat{\theta}_{wi})$, we use the delta method to calculate the omnibus Wald test of the equality restrictions (16). The test results are presented in Table 4 below. Before discussing these results, we present Table 3 which compares the implied win and conditional win probabilities to the non-parametric estimates of these probabilities at the first serve of each game (the serve for which we have the most data to produce reliable non-parametric estimates of these quantities). The Table also presents in the final column the results of a Hausman-Wu-Durbin test of our preferred reduced form specification.²²

From Table 3 we see that the calculated win probabilities $W_P(1, 1)$ are close to the non-parametric estimates of these quantities, and almost always within a standard deviation of each other. The P-values of the Hausman-Wu-Durbin specification tests in the final column of the table show that for all servers except Federer serving to Nadal we are unable to reject the reduced-form specification and its implied win probability. In the case of Federer serving to Nadal, the RF estimate of the win probability, $W_P(1, 1) = .796$, is slightly more than one standard deviation away

²² The Hausman-Wu-Durbin specification test compares two estimators of a given quantity or parameter: an inefficient but \sqrt{N} -consistent estimator that is consistent both under the null and alternative hypotheses, and an efficient estimator that is also \sqrt{N} -consistent for the true parameter under the null hypothesis but may be inconsistent under the alternative hypothesis, where N denotes the sample size. In our case the relevant null hypothesis is that our reduced form specification for (P, Π) is correct, and the non-parametric estimates of the win probabilities in Table 3 are inefficient but consistent even if the null hypothesis is false (i.e. our reduced form model is misspecified). Under the null hypothesis the Hausman test statistic equal to the square of the two estimates of the win probability divided by the differences in the asymptotic variances converges to a Chi-squared random variable with 1 degree of freedom.

from the non-parametric estimate of the win probability, .829. The middle columns compare the non-parametric estimates of the conditional win probabilities with the corresponding estimates implied by the reduced-form model, $W_P(1, 1, d)$ for $d \in \{l, b, r\}$. Again we see that the two estimates are generally close to each other, though due to low numbers of observations for certain serve directions, there are cases where we find relatively big differences between the two estimates. For example in the last row, Pete Sampras serving to Andre Agassi, due to the low probability that Sampras serves to the body (approximately 7%, see Table 1), combined with the relatively low number of games in which we observe him serving (140), the non-parametric estimate of the conditional win probability of serving to the body equals 1. Of course the non-parametric estimate is probably not a reasonable estimate in this case: instead it is likely to be a statistical fluke where Sampras happened to win every one of the 8 games where he served to Agassi's body on the very first serve of the game.

This discussion emphasizes that the non-parametric estimates of win probabilities depend only on the number of observations of *entire games*, whereas the reduced-form estimates of win probabilities are more efficient estimators that depend on the much greater information we have on *all serves and serve outcomes in every game*. An alternative strategy is to focus on using the less efficient (and fairly noisy) non-parametric estimates of conditional win probabilities to test the hypothesis of equal win probabilities. As a result, we can expect these testing strategies to have low statistical power and thus, unlikely to reject the null hypothesis when it is false.

However, the other problematic aspect is the restriction to first serves of each service game. We do this since subsequent serves would require us to condition on lagged serve directions, adding even more probabilities to estimate. From Table 1 we see that we typically observe at most a few hundred service games between a given server-returner pair. Thus, we have relatively few observations available to estimate conditional win probabilities at the first serve of a game, and fewer observations at higher score states further down in the game tree illustrated in Figure 1. In addition, we expect low the correlation between the direction chosen at the very first serve of the game and the ultimate game outcome, due to the fact that a typical game will have on average 6 to 8 intervening serves before the final game outcome is determined. So there could be higher correlations between serve directions and win/loss outcomes at states closer to the endgame states in this Figure.

Table 3: Estimated 1st serve win and conditional probabilities. selected elite servers

Server → Returner	Est.	Win prob 1st serve	Conditional win probability, 1st serve			Spec test P-value
			L	B	R	
Roger Federer → Rafael Nadal	NP RF	.796 (.026) .829 (.023)	.816 (.025) .828 (.024)	.650 (.030) .819 (.027)	.803 (.025) .833 (.022)	.004
Rafael Nadal → Roger Federer	NP RF	.786 (.026) .807 (.023)	.748 (.028) .808 (.023)	.896 (.020) .807 (.025)	.762 (.028) .803 (.025)	
Roger Federer → Novak Djokovic	NP RF	.810 (.024) .818 (.020)	.844 (.022) .826 (.020)	.867 (.021) .812 (.023)	.767 (.025) .813 (.021)	.504
Novak Djokovic → Roger Federer	NP RF	.782 (.025) .781 (.022)	.769 (.026) .792 (.022)	.710 (.028) .769 (.026)	.815 (.024) .774 (.024)	.910
Rafael Nadal → Novak Djokovic	NP RF	.712 (.035) .712 (.034)	.685 (.036) .712 (.034)	.726 (.035) .701 (.035)	.750 (.034) .718 (.034)	.992
Novak Djokovic → Rafael Nadal	NP RF	.829 (.029) .848 (.023)	.868 (.026) .854 (.023)	.735 (.034) .830 (.027)	.833 (.029) .849 (.023)	.278
Novak Djokovic → Andy Murray	NP RF	.794 (.034) .791 (.029)	.759 (.036) .796 (.031)	.750 (.036) .758 (.034)	.841 (.031) .799 (.029)	.871
Andy Murray → Novak Djokovic	NP RF	.721 (.038) .717 (.036)	.816 (.033) .735 (.036)	.500 (.042) .712 (.039)	.701 (.038) .703 (.037)	.675
Pete Sampras → Andre Agassi	NP RF	.885 (.028) .866 (.024)	.894 (.027) .866 (.025)	1.00 (.000) .839 (.029)	.859 (.030) .872 (.024)	.150
Andre Agassi → Pete Sampras	NP RF	.874 (.029) .859 (.024)	.907 (.026) .861 (.026)	.867 (.030) .853 (.026)	.852 (.032) .859 (.024)	.362

Thus, we expect much greater power from an omnibus test of equal win probabilities (16) that tests the restrictions for all states (x, m) simultaneously. We see this is indeed the case in Table 4, where the last column shows that the omnibus test of equal win probabilities is able to decisively reject this hypothesis for all of the server-returner pairs we tested.²³

For comparison, the middle column of Table 4 reports Wald tests of equal win probabilities

²³ The omnibus Wald tests amount to a test of 648 equality restrictions of the form given in (16). Since the conditional win probabilities are implicit function of (P, Π) and the latter are functions of the 44-dimensional parameter vector $\hat{\theta} = (\hat{\theta}_P, \hat{\theta}_\Pi, \hat{\theta}_{win})$, when we use the delta method to construct the omnibus Wald test statistic, a quadratic form with the 648×1 vector of differences in conditional win probabilities between directions l and b and b and r over all states and the implied 648×648 covariance matrix for these differences, the latter covariance matrix is expressed as a sandwich formula in terms of the 44×44 variance covariance matrix for the reduced-form parameter vector $\hat{\theta}$. Thus, the rank of this matrix, which also equals the degrees of freedom of the Chi-squared distribution of the test statistic under the null hypothesis, is at most 44. However, the rank is generally lower since in the initial states of tennis (e.g. states $x \in \{1, 3, 9\}$) there is no muscle memory, so in these states $W_P(x, m, d)$ is only defined for $m = 1$ and the other values are redundant.

Table 4: Tests of equal conditional win probabilities for all serve directions, selected elite servers

Server → Returner	Estimator	P-value, df Wald test of equal win probs, 1st serves	P-value, df Omnibus test of equal win probs, all serves
Roger Federer → Rafael Nadal	NP RF	.048, 2 .614, 2	0, 36
Rafael Nadal → Roger Federer	NP RF	.148, 2 .891, 2	4.9×10^{-24} , 37
Roger Federer → Novak Djokovic	NP RF	.402, 2 .237, 2	0, 37
Novak Djokovic → Roger Federer	NP RF	.382, 2 .067, 2	0, 35
Rafael Nadal → Novak Djokovic	NP RF	.780, 2 .267, 2	0, 36
Novak Djokovic → Rafael Nadal	NP RF	.351, 2 .235, 2	0, 36
Novak Djokovic → Andy Murray	NP RF	.643, 2 .058, 2	0, 39
Andy Murray → Novak Djokovic	NP RF	.002, 2 .057, 2	0, 39
Pete Sampras → Andre Agassi	NP RF	.471, 2 .037, 2	0, 38
Andre Agassi → Pete Sampras	NP RF	.889, 2 .850, 2	0, 37

restricted to the first serve of each game, i.e. we test the two restrictions $W_P(1, 1, l) = W_P(1, 1, b)$ and $W_P(1, 1, b) = W_P(1, 1, r)$, so the relevant Wald statistic is asymptotically Chi-squared with two degrees of freedom under the null hypothesis. We present the test statistics for both the non-parametric estimates of the conditional win probabilities (the rows labelled NP) as well as the conditional win probabilities implied by the reduced form estimates (the rows labelled RF). We see from the generally high P-values that these more limited Wald tests have much lower power in detecting deviations from the null hypothesis. Indeed, in only three cases we can reject the null hypothesis of equal win probabilities for all serve directions at the very first serve at the 5% level. Overall, we conclude that our approach to testing for equal win probabilities, combined with the much greater number of observations of service games compared to WW's original analysis explains why we are able to decisively reject the key implication of a mixed strategy

Nash equilibrium.

3.5 Testing for effects of “muscle memory”

We conclude this section by presenting evidence of serial dependence in serve directions and to a lesser extent, the POPs. We have already shown in section 3.1 that there are significant differences between the mixture probabilities servers use in first and second serves, so it should not be surprising that we also find significant serial dependence between first and second serves. However, as we noted in Section 2, this serial dependence is not necessarily inconsistent with equilibrium play: the server considers the option value of the second serve when choosing the speed and direction of the first serve.

The more important question is whether there is serial correlation *across successive first serves*. Note that for our preferred specification of the reduced form model there are only two ways for there to be serial dependence in successive first serves: 1) via the presence of the muscle memory state variable m which is effectively a lagged dependent variable capturing previous serve directions; and 2) via our ad/deuce and 1st/2nd serve partition of the state space which allows the probability distributions governing serve directions to differ depending on whether they are to the ad or deuce courts, or are first or second serves. Note the subtle distinction between “serial dependence” and “serial correlation.” Though allowing for the probability distributions over serve directions to differ over serves to the ad and deuce courts and first and second serves is indeed a type of serial dependence, there will still be zero serial correlation in serves in the absence of muscle memory effects, since the no muscle memory version of our reduced form specification implies that the direction of any serve is conditionally independent of the direction of any previous serve, and this implies zero correlation in the directions of successive serves.

To test for serial dependence in serve directions we use likelihood-ratio tests of a restricted version of our reduced-form model of tennis that excludes the muscle memory variable m . As we showed in Section 2.4 the directions of serves become serially independent under this specification. Table 5 presents the results of LR tests of the hypothesis of “no muscle memory effects” in the last column of the table show that except for the case of Nadal serving to Federer, we can reject the hypothesis of no muscle memory in serves at the 5% significance level. How-

Table 5: Tests for muscle memory effects in (P, Π) , selected elite servers

Server → Returner	Model	No muscle memory		Muscle memory		LR test
		AIC	LL	AIC	LL	P-value
Roger Federer → Rafael Nadal	Serves POPs	3764.7 3928.3*	-1874.4 -1940.1	3713.5* 3932.6	-1844.8 -1934.3	4.3×10^{-12} .170
Rafael Nadal → Roger Federer	Serves POPs	3397.8* 3814.5*	-1690.9 -1883.3	3400.3 3824.9	-1688.2 -1880.9	.249 .779
Roger Federer → Novak Djokovic	Serves POPs	4603.7 4617.4*	-2293.9 -2284.8	4554.1* 4625.4	-2265.1 -2280.7	9.3×10^{-12} .414
Novak Djokovic → Roger Federer	Serves POPs	4925.5 4871.3*	-2454.8 -2411.7	4871.6* 4871.9	-2423.8 -2403.9	1.1×10^{-12} .048
Rafael Nadal → Novak Djokovic	Serves POPs	2891.0 2879.6*	-1437.5 -1415.8	2889.3* 2892.4	-1432.6 -1414.2	.044 .921
Novak Djokovic → Rafael Nadal	Serves POPs	2744.8 2656.9*	-1364.2 -1304.5	2718.9* 2668.1	-1347.4 -1302.1	9.0×10^{-7} .779
Novak Djokovic → Andy Murray	Serves POPs	2458.0 2425.5*	-1221.0 -1188.7	2427.1* 2430.3	-1201.6 -1183.2	7.7×10^{-8} .202
Andy Murray → Novak Djokovic	Serves POPs	2525.9 2623.9	-1254.9 -1287.9	2524.0* 2524.1*	-1250.0 -1280.1	.044 .049
Pete Sampras → Andre Agassi	Serves POPs	2208.8 2296.2*	-1096.4 -1124.1	2194.9* 2299.7	-1085.4 -1117.9	2×10^{-4} .134
Andre Agassi → Pete Sampras	Serves POPs	1906.5 2113.2*	-945.3 -1032.6	1887.6* 2126.2	-931.8 -1031.1	2×10^{-5} .934

ever, when it comes to the POPs we have far weaker evidence of serial correlation. For most of the server-returner pairs in Table 5 we are unable to reject the hypothesis of no muscle memory effects.

Why would that be the case? We think it may have to do with the returner’s behavior. Specifically, if muscle memory effects are real, and the returner shifts his position accordingly, then the returner would effectively cancel out any effect that muscle memory would impart on the POPs. As a result, we would observe serial correlation in the server’s directional choices but not in the POPs. This would be consistent with a Nash equilibrium, as we demonstrate in appendix B. If play is not consistent with Nash equilibrium, the serial dependence in serve directions that we find could be another manifestation of disequilibrium play.

4 Dynamic structural analysis of serve strategies

In the previous section we estimated a reduced form model of tennis serve directions and POPs and showed that this flexible agnostic model of tennis decisively rejects the key implication of a mixed strategy Nash equilibrium: namely that the probability of winning the game is the same regardless of serve direction. In this section we attempt to get deeper insight into the behavior of elite pro servers by estimating three different structural models of serve behavior that impose the restriction that serve directions are chosen to maximize the server's probability of winning the game. The three structural models are 1) a *fully dynamic model* that assumes the server chooses a strategy that maximizes the probability of winning the entire *service game*; 2) a *myopic model* that assumes the server chooses a strategy that maximizes the probability of winning each *point*; and 3) a *fully myopic model* that assumes the server chooses a strategy that maximizes the probability of winning each *serve*. The fully myopic model ignores the option value of the second serve. For each of these models we estimate the server's *subjective POPs*, i.e. we find POPs that rationalize observed serve behavior as a best response to the server's potentially subjective beliefs about their own performance and the performance of the returner.

We use discrete choice models to construct the implied mixture probabilities over serve directions $P(d|x, m)$. We assume that at the moment each serve is made, the server's choice of direction reflects *trembles*. These are *IID* shocks that affect their perception of the probability of winning when serving to different directions d . We assume that these trembles or preference shocks are observed only by the server but not by the opponent or the econometrician. For example, the server may feel more comfortable hitting to a certain direction at some point in the match due to a psychological factor. Let $\varepsilon(d)$ be the tremble associated with serving to direction d . We assume the trembles are independently distributed across all three possible serve directions, $\{l, b, r\}$ and *IID* across successive serves, and have a Type 1 extreme value distribution with location parameter normalized so that $E\{\max_d \varepsilon(d)\} = 0$ and scale parameter $\lambda \geq 0$.

Let $\sigma_{FD}(x, m, \varepsilon)$ be the serve strategy under the fully dynamic structural model as function of the observed state (x, m) and the unobserved trembles $\varepsilon = (\varepsilon(l), \varepsilon(b), \varepsilon(r))$. The fully dynamic model presumes that for each (x, m, ε) the server chooses the serve direction that maximizes the

probability of winning the game, given by

$$\sigma_{FD}(x, m, \varepsilon) = \underset{d \in \{l, b, r\}}{\operatorname{argmax}} [\lambda \varepsilon(d) + V_\lambda(x, m, d)] \quad (17)$$

where $V_\lambda(x, m, d)$ is a conditional value function, the analog of the conditional win probability $W_S(x, m, d)$ defined in equations (2) and (3) of Section 2, where the analog of the function $W_S(x, m)$ given by the Bellman equation (1) is replaced by $V_\lambda(x, m)$ given by

$$V_\lambda(x, m) = \lambda \log \left(\sum_{d \in \{l, b, r\}} \exp \{V_\lambda(x, m, d) / \lambda\} \right). \quad (18)$$

The serve direction MP implied by the fully dynamic model is denoted by $P_{FD}(d|x, m)$ and is given by

$$P_{FD}(d|x, m) = \Pr \{d = \sigma_{FD}(x, m, \varepsilon) | x, m\} = \frac{\exp \{V_\lambda(x, m, d) / \lambda\}}{\sum_{d' \in \{l, b, r\}} \exp \{V_\lambda(x, m, d') / \lambda\}}. \quad (19)$$

The choice probability $P_{FD}(d|x, m)$ gives the probability of choosing to serve to direction d in observed state (x, m) accounting for the randomness of the unobserved trembles, ε . Though $\sigma_{FD}(x, m, \varepsilon)$ is a pure strategy from the standpoint of the server, it appears to be a mixed strategy from the standpoint of someone who does not observe ε . This device enables us to rationalize or fit observed mixed serve strategies without imposing equal win probabilities, i.e. imposing equality of $V_\lambda(x, m, d)$ over serve directions d . Since the trembles are *IID* across serves, it would appear that this model should also imply conditional independence in serve directions across successive first and second serves. However, that will actually only be true if there is no muscle memory, i.e. the variable m does not enter $V_\lambda(x, m, d)$ (recall that m is a vector that stores the directions of the two most recent first serves). With muscle memory present, we can still have serial correlation in serves even though the trembles are *IID*.

Theorem 3 of Iskhakov, Jørgensen, Rust, and Schjerning (2017) establishes that the following limit holds as $\lambda \downarrow 0$:

$$W_S(x, m, d) = \lim_{\lambda \downarrow 0} V_\lambda(x, m, d) \quad (20)$$

uniformly for all (x, m, d) . This result implies that the only way for $P_{FD}(d|x, m)$ to converge to a mixed strategy as $\lambda \downarrow 0$ is if the limiting conditional win probabilities $W_S(x, m, d)$ obey the equal win probability constraints, $W_S(x, m, l) = W_S(x, m, b) = W_S(x, m, r)$ for all (x, m) .

The myopic and fully myopic models have the same general structure as the fully dynamic model, so the serve strategies, value functions, and choice (mixing) probabilities are given by the same equations, (17), (18), and (19). The difference is the equations defining V_λ . In the fully myopic model we have

$$V_\lambda(x, m, d) = \pi(\text{in}|d, x, m)\pi(\text{win}|d, x, m), \quad (21)$$

i.e. $V_\lambda(x, m, d)$ is the probability of winning the serve. Thus, the fully myopic server chooses a serve strategy to maximize the probability of winning each serve, without any concern about the effect of winning or losing on the future state of the game. The myopic server's objective is to win each *point*, but the server does recognize the option value provided by the second serve in the event of a faulted first serve. Thus, the myopic server conducts a two period backward induction calculation. If x is a second serve state (i.e. x is an even number between 2 and 36 in our numbering of tennis states in Figure 4), then $V_\lambda(x, m, d)$ coincides with the fully myopic formula given in equation (21) above. However in any non-terminal first serve state (any odd value of x from 1 to 35), V_λ is given by

$$V_\lambda(x, m, d) = \pi(\text{in}|d, x, m)\pi(\text{win}|d, x, m) + [1 - \pi(\text{in}|d, x, m)]V_\lambda(x + 1, m') \quad (22)$$

where $m' = (d^{-1}, d)$ and $V_\lambda(x, m)$ is given by equation (18). As we noted in Section 2, the myopic serve strategy coincides with the fully dynamic serve strategy in the limit as $\lambda \downarrow 0$ when the GMC (9) holds.

Note that all three structural models have mixed serve probabilities that are implicit functions of the POPs. Thus, the mixed serve directions for the three structural models are entirely determined by the POPs and the single parameter λ controlling the magnitude of the trembles. In comparison, the reduced-form model of serve directions is estimated separately from the POPs with flexible parameterization of serve directions. The structural models can be viewed as restricted special cases of the most flexible specification of the reduced form serve model. This enables us to conduct likelihood ratio specification tests for the three structural models relative to the unrestricted reduced form specification.²⁴

²⁴ Strictly speaking for a likelihood ratio specification test to be valid, we would need to estimate a fully unrestricted version of the reduced form model with a total of 624 parameters so that it has the flexibility to replicate

4.1 Structural estimation results

We estimated the three structural models by maximum likelihood using the full panel likelihood function (13) with data for hard courts for the ten elite server-returner pairs listed in Table 6. We used the same specification for the POPs as in our reduced form results presented in Section 3 (for the specification with muscle memory), so our structural models involve a total of 33 parameters: the 32×1 vector of POP parameters $(\theta_{in}, \theta_{win})$, plus the extreme value scaling parameter λ . Note that unlike the reduced form specification, the structural serve probabilities are functions of the POP parameters; and thus, the likelihood function is no longer block-diagonal between the POP parameters $(\theta_{in}, \theta_{win})$ and λ , whereas we do have block diagonality between the reduced form serve parameters θ_P and POP parameters $(\theta_{in}, \theta_{win})$.

Our structural estimates of the POPs can be regarded as estimates of the server’s *subjective beliefs* that may or may not correspond to rational beliefs about the true POPs which we obtain from our reduced-form estimates of the POPs. That is, the only way for the structural models to simultaneously fit both the observed serve direction data and the serve outcome data is by distorting the POPs to help rationalize the server’s choice of serve directions. This is in contrast to the reduced form model where there are no cross equation constraints linking the parameters of the POPs and the serve probabilities. The reduced form model maximizes separate likelihoods for the POP parameters $(\theta_{in}, \theta_{win})$ and the serve parameters θ_P , whereas the structural models fit the parameters $(\lambda, \theta_{in}, \theta_{win})$ to maximize a joint likelihood for serves and POPs, forcing a tradeoff between fitting serve directions and game outcomes.

Table 6 summarizes the structural estimation results for the same 5 elite server-returner pairs that we analyzed in Section 3.²⁵ For comparison, we show the optimized log-likelihood function for the reduced form model and the number of serve observations used to estimate the parameters,

any conditional probability $P(d|x, m)$. As we discussed above, given the limited number of observations for specific server-returner pairs, our specification for $P(d|x, m)$ depends on only 12 parameters, though it produces estimates that fit the data well. Though our specification does not strictly nest the structural models, the reduced form model has sufficient flexibility to closely approximate the structural serve probabilities. We use this as a justification for “quasi likelihood ratio” tests of the structural models relative to the reduced form model. We can also do tests using the non-nested specification test of Vuong (1989), however we prefer to rely on the AIC model selection criterion to select our preferred structural specification, similar to the way we used it to select our preferred specification for the reduced form model.

²⁵ Due to limited space we do not provide the 32 parameter estimates of $(\theta_{in}, \theta_{win})$ and their standard errors for all 10 servers for all 3 structural models. We are happy to provide these results to interested readers on request.

Table 6: Summary of structural estimation results for selected elite pro server-returner pairs

Player pair	Reduced form	Fully myopic	Myopic	Fully dynamic
Server→ Returner	LL, N BIC	LL, $\hat{\lambda}$ AIC, LR P-value	LL, $\hat{\lambda}$ AIC, LR P-value	LL, $\hat{\lambda}$ AIC, LR P-value
Roger Federer → Rafael Nadal	-3779.1, 2011 7646.1	-3788.2, 7.9×10^{-4} 7642.7, .074	-3783.8, 5.8×10^{-3} 7633.7, .571	-3817.3, 1.1×10^{-4} 7700.7, 6.9×10^{-12}
Rafael Nadal → Roger Federer	-3569.1, 1882 7226.2	-3571.3, 6.1×10^{-3} 7208.6, .957	-3570.6, 2.7×10^{-3} 7207.3, .990	-3632.4, 2.2×10^{-4} 7330.7, 8.8×10^{-22}
Roger Federer → Novak Djokovic	-4545.8, 2333 9179.5	-4551.2, .010 9168.4, .457	-4552.2, 4.5×10^{-3} 9170.4, .300	-4576.0, 9.1×10^{-4} 9128, 7.5×10^{-9}
Novak Djokovic → Roger Federer	-4827.7, 2372 9743.5	-4840.0, .011 9746.0, .010	-4842.0, 1.9×10^{-3} 9750.0, 2.6×10^{-3}	-4844.8, 2.4×10^{-4} 9755.7, 3.3×10^{-4}
Rafael Nadal → Novak Djokovic	-2846.8, 1405 5781.7	-2853.8, 1.1×10^{-4} 5773.7, .232	-2853.2, 8.2×10^{-5} 5772.4, .310	-2864.5, 1.2×10^{-5} 5795.0, 2.2×10^{-4}
Novak Djokovic → Rafael Nadal	-2649.5, 1344 5387	-2659.9, .070 5385.9, .035	-2656.1, .097 5378.2, .285	-2656.7, 1.7×10^{-5} 5375.3, .505
Novak Djokovic → Andy Murray	-2384.7, 1201 4857.5	-2396.2, 9.9×10^{-3} 4858.3, .018	-2396.9, .044 4859.8, .011	-2413.0, 5.5×10^{-4} 4892.0, 4.0×10^{-8}
Andy Murray → Novak Djokovic	-2649.5, 1328 5387.0	-2536.4, .014 5138.9, .309	-2539.8, 6.9×10^{-3} 5145.7, .051	-2556.3, 9.7×10^{-5} 5350.0, 2.2×10^{-7}
Pete Sampras → Andre Agassi	-2203.3, 1181 4494.6	-2219.6, .031 $4505.3, 5.9 \times 10^{-4}$	-2217.7, .037 $4501.4, 2.5 \times 10^{-3}$	-2240.2, 1.9×10^{-6} $4546.4, 2.3 \times 10^{-11}$
Andre Agassi → Pete Sampras	-1962.9, 1050 4013.8	-1973.0, 1.8×10^{-3} 4011.9, .043	-1970.8, 1.9×10^{-5} 4007.6, .145	-2004.7, 5.7×10^{-5} $4075.5, 2.8 \times 10^{-13}$

along with the point estimates of λ for each of the structural models. The second row of numbers for each server-returner pair reports the AIC value along with the P-value of a “likelihood ratio test” of each structural model relative to the reduced form model. As per our discussion above, these models are not strictly nested within each other, though the reduced form model is the more flexible specification with a total of 44 parameters.

In view of this, we followed the approach in Section 3 and selected our preferred model as the one with the smallest value of AIC, labelled in bold font. Notice that the best-fitting model selected by AIC is also the model that has the highest P value for a quasi-likelihood ratio test of each structural model relative to the reduced form model. We see that the best model selected by AIC is generally is also the model for which there is the least evidence for rejecting it in favor of the reduced form model via the likelihood ratio test. In two cases, Djokovic serving to

Table 7: Omnibus tests of equal win probabilities for selected elite pro server-returner pairs

Player pair	Reduced form	Fully myopic	Myopic	Fully dynamic
Server → Returner	Wald stat, df P-value	Wald stat, df P-value	Wald stat, df P-value	Wald stat, df P-value
Roger Federer → Rafael Nadal	2.6×10^6 , 36 0	6135, 28 0	12258, 28 0	0.8, 29 1.000
Rafael Nadal → Roger Federer	199, 37 4.9×10^{-24}	4717, 29 0	20, 27 .825	11, 29 .999
Roger Federer → Novak Djokovic	105946, 37 0	3412, 29 0	931, 28 6.0×10^{-178}	117, 29 1.9×10^{-12}
Novak Djokovic → Roger Federer	1.1×10^6 , 35 0	37094, 29 0	12, 28 .995	37, 29 .142
Rafael Nadal → Novak Djokovic	11511, 36 0	68074, 29 0	.0001, 25 1.0000	1.0×10^{-5} , 25 1.000
Novak Djokovic → Rafael Nadal	146629, 36 0	219705, 29 0	247715, 29 0	.008, 27 1.000
Novak Djokovic → Andy Murray	4152, 39 0	7021, 29 0	109771, 29 0	8.5, 29 1.000
Andy Murray → Novak Djokovic	94519, 39 0	5784, 29 0	30, 29 .380	1.1, 28 1.000
Pete Sampras → Andre Agassi	2268, 38 0	853141, 29 0	138153, 29 0	33, 29 .265
Andre Agassi → Pete Sampras	44190, 37 0	177305, 29 0	1.9×10^{-5} , 24 1.000	5, 28 1.000

Federer and Sampras serving to Agassi, the AIC criterion selects the reduced form model and the likelihood ratio test strongly rejects all three structural models.

For the other 8 servers, we see that the AIC selects the full DP model in only one case, Djokovic serving to Nadal. AIC selects the myopic two-period DP model as the best model for four other servers, and it selects the fully myopic model for three of the servers. Note that the scale parameters $\hat{\lambda}$ for all specifications are uniformly small, which means that the data finds limited role for “trembles” to explain the observed mixed serve strategies of these players. Instead, as we noted above, the maximum likelihood estimates of the POPs, $(\hat{\theta}_{in}, \hat{\theta}_{win})$ are distorted in a manner that results in conditional win probabilities much closer to equality than the ones implied by the reduced form estimates of the POPs.

Note that the λ estimates decline for the structural models that require increasingly “far

sighted” calculations by the server: the λ values for the fully dynamic model are typically very close to zero; those for the myopic model are small but somewhat larger; and the λ estimates for the fully myopic serve model are typically the largest. When λ is sufficiently small, conditional value functions $V_\lambda(x, m, d)$ are extremely close to the conditional win probabilities as per the limiting result in equation (20). But when λ is larger the trembles play a more important role in the mixed serve strategies, allowing more freedom for the conditional value functions (and the conditional win probabilities) to differ across serve directions.

We can see this in Table 7 which reports the results of the omnibus Wald tests of equal win probabilities implied by each of the four specifications. For convenience, we repeat the results of the omnibus tests for the reduced form model in the first column. Due to the larger estimated values of λ for the fully myopic model, we strongly reject the hypothesis of equal win probabilities for this specification, just as for the reduced-form model. However for the fully dynamic specification, which has the smallest estimated values of λ , we are unable to reject the hypothesis of equal win probabilities for any of the 10 elite server-returner pairs except for Federer serving to Djokovic. Note that Table 6 shows that the fully dynamic specification with the highest estimated $\hat{\lambda}$ is for Federer serving to Djokovic.

Table 8 provides the estimated win probabilities and the P-values of Hausman-Wu-Durbin specification tests of the different model specifications. Recall this test is based on a comparison of the implied win probabilities calculated via equation (7) to the non-parametric estimate of the win probability, where the latter is simply the fraction of the games between a given server-returner pair that the server won. The first column of Table 8 presents the non-parametric estimate of the win probability and its standard error, and the remaining columns present the estimated win probabilities implied by equation (7) with standard errors calculated via the delta method.²⁶

We see that the specification tests strongly reject the fully dynamic specification, with the exception of Djokovic serving to Nadal. Recall from Table 6 that the AIC criterion selects the fully dynamic model as the preferred specification for Djokovic serving to Nadal, so it is reassuring to know that it is not rejected by the specification tests. But for the the other servers, we note that the

²⁶ Note that the model estimates are relatively efficient estimates of the win probability (as reflected by their smaller standard errors) but they are consistent only if the model specification is correct. The less efficient non-parametric estimator of the win probability is consistent regardless of whether any of the model specifications are correct or not.

Table 8: Win probabilities and Hausman tests for selected elite pro server-returner pairs

Player pair	Nonparametric win probability	Reduced form	Fully myopic	Myopic	Fully dynamic
Server→ Returner	$\hat{W}(1, 1)$	$W(1, 1)$ P-value	$W(1, 1)$ P-value	$W(1, 1)$ P-value	$W(1, 1)$ P-value
Roger Federer → Rafael Nadal	.796 (.026)	.829 (.023) .004	.825 (.021) .050	.830 (.021) .025	.749 (.021) 1.3×10^{-3}
Rafael Nadal → Roger Federer	.786 (.026)	.807 (.023) .107	.806 (.022) .147	.807 (.022) .132	.641 (.022) 1.0×10^{-22}
Roger Federer → Novak Djokovic	.810 (.024)	.818 (.020) .504	.818 (.020) .506	.818 (.020) .509	.759 (.020) 1.0×10^{-4}
Novak Djokovic → Roger Federer	.781 (.025)	.781 (.023) .910	.779 (.022) .838	.778 (.022) .756	.746 (.021) .011
Rafael Nadal → Novak Djokovic	.712 (.035)	.712 (.034) .992	.703 (.033) .485	.706 (.033) .675	.650 (.032) 4.1×10^{-5}
Novak Djokovic → Rafael Nadal	.829 (.029)	.848 (.023) .278	.846 (.023) .318	.850 (.023) .231	.797 (.024) .052
Novak Djokovic → Andy Murray	.794 (.034)	.792 (.029) .871	.791 (.030) .840	.791 (.030) .844	.750 (.029) .012
Andy Murray → Novak Djokovic	.721 (.038)	.717 (.036) .675	.717 (.034) .792	.718 (.035) .825	.584 (.032) 1.6×10^{-11}
Pete Sampras → Andre Agassi	.885 (.028)	.866 (.024) .150	.863 (.024) .130	.866 (.024) .187	.757 (.028) 0
Andre Agassi → Pete Sampras	.874 (.029)	.859 (.024) .362	.854 (.026) .183	.853 (.026) .150	.715 (.028) 1.4×10^{-58}

fully dynamic specification typically significantly underestimates the true win probability. This is caused by the need to distort the POPs to rationalize serve behavior as a best response to the estimated POPs in the fully dynamic specifications. As we will show in the next subsection, the serve strategy for the fully dynamic specification is close to the “true” serve strategy captured by the reduced form model, but the estimated POPs from the fully dynamic model imply far less favorable performance for the server than the POPs estimated from the reduced form model. That is, the fully dynamic POPs generally imply a higher probability of faults and a lower probability of winning the rally given a serve is in compared to the reduced form POPs. As a result of this, the calculated win probabilities from the fully dynamic model are significantly lower than the true win probabilities (i.e. the non-parametric estimate given in the first column). In contrast, the specification tests are generally unable to reject the fully myopic and myopic serve models. This

is consistent with the results we reported in Table 6 where we showed that these models were the ones most frequently selected as having the lowest AIC values.

We do not have the space to show the 32 estimated $(\theta_{in}, \theta_{win})$ parameters for all of the 8 model specifications (RF, fully myopic, myopic, fully dynamic with and without muscle memory) for the ten server-returner pairs we have analyzed in this section, nor do we have space to show the nearly 2000 values of the POP arrays $(\pi(in|x, m, d), \pi(win|x, m, d))$ these parameters imply. In order to get a better sense of how the structural models distort or *attenuate* the POPs to rationalize the observed mixed serves strategies, we present estimated POPs for the first serve of a tennis game, $(\hat{\pi}(in|1, 1, d), \hat{\pi}(win|1, 1, d))$ in Table 9.

The left hand side columns present the point estimates $(\hat{\pi}(in|x, m, d), \hat{\pi}(win|x, m, d))$ for the three serve directions $d \in \{l, b, r\}$, and the right hand columns present the estimated standard errors of these probabilities computed from the POP parameters $(\hat{\theta}_{in}, \hat{\theta}_{win})$ using the delta method. Though at first glance the estimated POPs seem relatively similar across the four different models, there are a number of big differences, especially when comparing the POPs from the fully dynamic model with the unrestricted reduced form estimates. For example, the $\hat{\pi}(win|1, 1, l)$ estimates for Nadal serving to Federer differ by more than two standard errors: the value for the fully dynamic model is .568 whereas the reduced form estimate is .653. Similarly, for Federer serving to Djokovic, the fully dynamic model estimate of $\hat{\pi}(win|1, 1, l) = .732$, which is more than two standard deviations below the reduced form estimate, $\hat{\pi}(win|1, 1, l) = .803$.

However, we also see how the structural estimates of the POPs are *attenuated* in order to satisfy the equal win probability constraints. We observe this in a much lower difference between the highest and lowest estimates of the POPs across different serve directions for the fully dynamic model compared to the reduced form estimates. For example, for Nadal serving to Djokovic, the maximum difference in the fully dynamic estimates of $\hat{\pi}(win|1, 1, d)$ across the 3 serve directions is $.642 - .616 = .026$ which is approximately equal to the estimated standard deviation of these probabilities. However, the maximum difference in the reduced form estimates is $.716 - .573 = .143$, which is more than three times the estimated standard deviation.

Note that when λ is sufficiently small the structural models predict that the effect of trembles are negligible, and servers will choose to serve to the direction with the highest win probability. In this situation, in order to fit the observed mixed serve strategies, the model is forced to equate

Table 9: Estimated point outcome probabilities (POP) for selected elite server-returner pairs

Model	$\hat{\pi}(in 1, 1, d)$			$\hat{\pi}(win 1, 1, d)$			$se(\hat{\pi}(in 1, 1, d))$			$se(\hat{\pi}(win 1, 1, d))$		
Direction	L	B	R	L	B	R	L	B	R	L	B	R
	Roger Federer → Rafael Nadal											
fully dynamic	.517	.763	.700	.769	.669	.689	.027	.051	.022	.026	.029	.020
myopic	.516	.730	.718	.809	.706	.724	.027	.054	.022	.025	.032	.019
fully myopic	.567	.722	.694	.846	.661	.694	.022	.047	.020	.026	.048	.023
reduced form	.518	.726	.720	.818	.665	.720	.029	.059	.026	.034	.070	.032
	Rafael Nadal → Roger Federer											
fully dynamic	.673	.760	.587	.568	.557	.579	.022	.031	.044	.020	.018	.026
myopic	.673	.772	.579	.644	.629	.651	.023	.031	.046	.022	.019	.026
fully myopic	.677	.751	.628	.655	.582	.692	.020	.028	.036	.025	.029	.042
reduced form	.673	.772	.572	.653	.622	.614	.027	.032	.052	.028	.042	.065
	Roger Federer → Novak Djokovic											
fully dynamic	.580	.699	.645	.732	.687	.704	.025	.040	.023	.021	.024	.019
myopic	.581	.707	.657	.762	.703	.730	.026	.040	.024	.022	.024	.019
fully myopic	.592	.685	.653	.787	.658	.712	.021	.030	.021	.025	.041	.023
reduced form	.592	.704	.649	.803	.681	.689	.028	.048	.026	.029	.054	.033
	Novak Djokovic → Roger Federer											
fully dynamic	.602	.665	.688	.664	.654	.652	.027	.022	.026	.021	.018	.018
myopic	.587	.678	.700	.685	.669	.668	.027	.027	.025	.022	.018	.018
fully myopic	.599	.670	.692	.730	.635	.626	.020	.021	.021	.025	.026	.024
reduced form	.592	.720	.678	.750	.615	.608	.026	.042	.029	.031	.046	.034
	Rafael Nadal → Novak Djokovic											
fully dynamic	.641	.724	.587	.631	.616	.642	.034	.036	.034	.026	.023	.028
myopic	.651	.719	.584	.646	.634	.660	.035	.036	.035	.026	.024	.029
fully myopic	.658	.714	.582	.632	.582	.714	.030	.032	.029	.033	.033	.035
reduced form	.654	.712	.593	.640	.573	.716	.040	.041	.036	.046	.048	.040
	Novak Djokovic → Rafael Nadal											
fully dynamic	.535	.773	.719	.772	.668	.677	.033	.047	.029	.030	.022	.022
myopic	.548	.749	.730	.757	.574	.710	.037	.049	.032	.030	.042	.024
fully myopic	.591	.720	.713	.809	.550	.659	.028	.046	.028	.032	.051	.032
reduced form	.548	.753	.728	.780	.591	.680	.037	.048	.035	.043	.069	.042

conditional win probabilities. We see this most clearly in the inability of the omnibus Wald test to reject the hypothesis of equal conditional win probabilities for the fully dynamic specification in Table 7. For the myopic and fully myopic models we showed that the estimated λ values were larger, so trembles play a greater role in explaining serves. This greater freedom allows these models to rationalize the observed mixed strategies without having to equate conditional win probabilities, and this is reflected by the greater frequency of rejection of the equal win probability hypothesis for these specifications in Table 7, especially for the fully myopic structural model. The reduced-form model places no constraint on the estimation of the POPs since it estimates separate parameters and likelihoods for serves and POPs. We have shown that the flexibility of this specification results in nearly unbiased estimates of the POPs and implied win probabilities.

We also observe significant *dynamic attenuation* in the POPs. That is, as we noted in the previous section, the reduced form estimation results reveal much stronger evidence of serial correlation in serve behavior compared to the POPs. In the fully dynamic model, the degree of serial correlation in both serves and the POPs is attenuated (i.e. closer to zero; and thus, more likely to be statistically insignificant). In fact, for most servers the fully dynamic model does not exhibit any statistically detectable serial correlation in the structural estimates of the POPs, though it does predict serial correlation in serves. What explains this paradox? The explanation is that when λ is close to zero, *serve strategies are very sensitive to small changes in the POPs* since trembles play a negligible role and the server chooses to serve to the direction with the highest win probability. Thus, it is possible to produce significant muscle memory effects in serve strategies (i.e. dependence of the current serve direction to the direction of the previous serve to the same court) via *very tiny, oscillations in the POPs which are hard to detect statistically*.

Now we return to the key question of this paper: do these distorted/attenuated estimates of the POPs enable the structural models to rationalize observed serve behavior as mixed strategies consistent with Nash equilibrium? We have shown that at best, the structural models are able to rationalize observed serve behavior as a best response, but only relative to the server's *subjective perception* of their environment and returner, as captured by the structural estimates of the POPs. These subjective beliefs are distorted estimates of the true POPs which are consistently estimated by the unrestricted reduced-form model. A Nash equilibrium entails a key assumption of *rationality* i.e. the players' subjective beliefs about each other coincide with the truth. In the

next section, we will directly calculate best response strategies to our estimates of the true POPs using dynamic programming and compare how well these strategies perform relative to the mixed serve strategies the players actually use.

4.2 Calculating best response serve strategies

In this section we provide a more powerful direct test of Nash equilibrium play in tennis: we construct alternative *deviation* serve strategies that significantly increase a server's chance of winning the game compared to the mixed strategy they are currently using. If the hypothesis of Nash equilibrium is correct, it should be impossible to construct any such deviation strategies. We construct deviation strategies via dynamic programming, which are generally pure strategies. The DP serve strategies exploit the unequal win probabilities captured by our reduced form estimates of the POPs. At each stage of the game, the DP strategy chooses the serve direction that has the maximum conditional win probability, see equation (4) of Section 2, where the optimal conditional win probability $W_S(x, m, d)$ is calculated via the Bellman equations given in equations (1), (2) and (3) of Section 2.

For comparison purposes, we also calculate suboptimal serve strategies based on the myopic and fully myopic specifications described at the start of this section. However, in all three cases, we do the calculations using the reduced form estimates of the POPs, not the structural estimates of the POPs which we showed in the previous section were distorted estimates of the true POPs. We also do the calculations with $\lambda = 0$, i.e. we do not allow for any "trembles" in our calculated serve strategies. Note that when the GMC (9) holds, the optimal myopic serve strategy (calculated via a 2 period DP within for each point separately, without considering the option value of the future state of the game that the full DP calculation accounts for) coincides with the full DP solution. Therefore, our calculations will be able to reveal when the GMC holds and does not hold.

Our general approach to testing for Nash equilibrium is to test the necessary condition that there is no alternative serve strategy than can increase the server's chance of winning the game *given* the strategy used by the returner. Under our stationarity assumption, the strategy used by the returner is encoded in our estimates of the POPs. The POPs also capture unobserved aspects

of the server’s strategy such as serve speed and spin, the server’s chance of faulting first and second serves, and the server’s playing ability during rallies, which affects $\pi(\text{win}|x, m, d)$. While we are interested in statistically significant rejections of Nash equilibrium, our primary goal is to assess the *magnitude* of the difference between the actual serve strategies used by elite players and the optimal strategies we estimate. A natural measure of this difference is the increase in the win probability from switching from the observed serve strategy to the optimal DP serve strategy.

In order to test for Nash equilibrium we appeal to the *one shot deviation principle* which states that there is *no deviation at any stage of a dynamic game that can increase the server’s chance of winning, given the strategy of the returner*. Indeed, we find that there are profitable one shot deviations at many stages of the game. While each such deviation yields a modest improvement in the win probability, the cumulative effect of all profitable deviations is often a large improvements in the overall game win probability. Of course, if a server were to switch to the optimal serve strategy we estimate, the returner may detect the change and adjust their own strategies, which would then result in changes in the POPs, likely mitigating the gain we estimate.

It is important to acknowledge one key shortcoming of our approach to testing the hypothesis of Nash equilibrium: we only have *estimates* of the POPs rather than the *true* POPs. We acknowledge that estimation error in the POPs could result in spurious, upward biased, estimates of the win probability when we use a noisy estimate of the POPs to calculate a best response strategy via DP instead of using the true POPs. Appealing to the usual common knowledge assumption underlying Nash equilibrium, the server and returner play with knowledge of the true POPs and true serve strategy. As econometricians we only have noisy estimates of the POPs and the server’s mixed serve strategy.

In order to account for this, we use the principle of maximum likelihood estimation to calculate an approximate probability distribution for the true POPs based on the data we observe. The true POPs are asymptotically distributed about the maximum likelihood estimate of the POPs according to the normal asymptotic distribution of the reduced form POP parameters $(\hat{\theta}_{in}, \hat{\theta}_{win})$. Via stochastic simulation, we can draw from the distribution of the true POPs by drawing values of $(\tilde{\theta}_{in}, \tilde{\theta}_{win})$ from an asymptotic normal distribution centered at the MLE $(\hat{\theta}_{in}, \hat{\theta}_{win})$ with a covariance matrix equal to the asymptotic covariance of the MLE. This gives is a probability distribution over the true POPs that the server might actually be facing (and knows, via the com-

mon knowledge assumption underlying Nash equilibrium), and we evaluate our calculated best response serve strategy using a robust control approach, where we calculate the win probability for our DP serve strategy for a random sample of POPs that the DP strategy was not “expecting”.

Recall that σ_S was used in Section 2 to denote the optimal serve strategy, which is an implicit function of the POPs, Π , which we now make explicit by writing $\sigma_S(\Pi)$. Assume a Nash equilibrium and let Π^* and P^* denote the true equilibrium POPs mixed serve strategy, respectively. By assumption, the players have common knowledge of these POPs. While we do not directly observe Π^* and P^* , we can consistently estimate them with sufficient data. In particular, the hypothesis of Nash equilibrium implies that for any alternative serve strategy σ we have

$$W_S(P^*, \Pi^*) \geq W_S(\sigma, \Pi^*). \quad (23)$$

Let $\sigma_S(\Pi^*)$ be the optimal dynamic serve strategy (generally a pure strategy) calculated by dynamic programming for the true Nash equilibrium POPs, Π^* . Then by definition of optimality we have

$$W_S(\sigma_S(\Pi^*), \Pi^*) \geq W_S(P^*, \Pi^*) \geq W_S(\sigma, \Pi^*) \quad (24)$$

for all stationary Markovian serve strategies σ . Together, inequalities (23) and (24) imply the key equality

$$W_S(P^*, \Pi^*) = W_S(\sigma_S(\Pi^*), \Pi^*), \quad (25)$$

that serves as the basis for our direct test of a mixed strategy Nash equilibrium in tennis: the optimal DP serve strategy should not result in a higher win probability compared to the mixed serve strategy P^* that the server actually used.

To illustrate the problem with comparing the win rate given the optimal best response to the *estimated POPs* $\hat{\Pi}$ to the win rate given the estimated strategy \hat{P} and the estimated POPs, note that by definition of optimality we have

$$W_S(\sigma_S(\hat{\Pi}), \hat{\Pi}) \geq W_S(\hat{P}, \hat{\Pi}) \quad (26)$$

That is, the optimal win probability using the estimated POPs will always be at least as high as the win probability implied by the estimated mixed serve strategy \hat{P} and the estimated POPs, $\hat{\Pi}$.

To develop a meaningful test of the key equality (25) we rely on the Continuous Mapping Theorem and the fact that $W_S(\sigma, \Pi)$ is a continuous function of the serve strategy σ and POPs Π .

Thus, we have

$$\sqrt{N} [W_S(\hat{P}, \hat{\Pi}) - W_S(P^*, \Pi^*)] \implies N(0, \Omega(P^*, \Pi^*)), \quad (27)$$

and

$$\sqrt{N} [W_S(\sigma_S(\hat{\Pi}), \hat{\Pi}) - W_S(\sigma_S(\Pi^*), \Pi^*)] \implies N(0, \Omega(\sigma_S(\Pi^*), \Pi^*)), \quad (28)$$

where $\Omega(P^*, \Pi^*)$ and $\Omega(\sigma_S(\Pi^*), \Pi^*)$ are the asymptotic variances of the win probability, which can be calculated from the asymptotic variance covariance matrix of the reduced-form estimates of the serve parameters, $\hat{\theta}_P$ and POP parameters $(\hat{\theta}_{in}, \hat{\theta}_{win})$ using the delta method. Let $\tilde{\Pi}_1$ and $\tilde{\Pi}_2$ be two *independent* random POPs that are drawn from the same normal asymptotic distribution as the reduced form estimate of the POPs, $\hat{\Pi}$. Then if the hypothesis of Nash equilibrium holds, using the key equality (25) and the limiting results in (27) and (28) above, we have

$$\sqrt{N} \frac{W_S(\hat{P}, \tilde{\Pi}_1) - W_S(\sigma_S(\hat{\Pi}), \tilde{\Pi}_2)}{\sqrt{\Omega(\hat{P}, \hat{\Pi}) + \Omega(\sigma_S(\hat{\Pi}), \hat{\Pi})}} \implies N(0, 1), \quad (29)$$

where we have used the independence between the simulated POPs $\tilde{\Pi}_1$ and $\tilde{\Pi}_2$ to derive the limiting $N(0, 1)$ asymptotic distribution for our Nash equilibrium test statistic in (29).

Notice that the use of randomly drawn POPs circumvents the tautological inequality (26). Specifically, while $\sigma_S(\hat{\Pi})$ is a best response to the reduced form POPs $\hat{\Pi}$, it is not necessarily a best response to the *randomly drawn POPs* $\tilde{\Pi}$; and thus, it is possible that $W_S(\hat{P}, \tilde{\Pi}) > W_S(\sigma_S(\hat{\Pi}), \tilde{\Pi})$. That is, it is possible that the estimated reduced form serve strategy \hat{P} will have a higher win probability than the DP serve strategy $\sigma_S(\hat{\Pi})$ for a randomly drawn POP $\tilde{\Pi}$. However if the Nash equilibrium hypothesis is true, and the two independently drawn POPs $\tilde{\Pi}_1$ and $\tilde{\Pi}_2$ are drawn from the normal asymptotic distribution for $\hat{\Pi}$, then both win probabilities will be close to each other and close to the common Nash equilibrium win probability given in equation (25); and thus, the Nash test statistic (29) will have an asymptotic $N(0, 1)$ under the Nash null hypothesis.

We do not need to rely on only a single pair of randomly drawn POPs, $(\tilde{\Pi}_1, \tilde{\Pi}_2)$ we can use T *iid* randomly drawn pairs of POPs $\{(\tilde{\Pi}_{1,t}, \tilde{\Pi}_{2,t})\}$ (each drawn from the asymptotic normal distribution of the reduced form estimate of the POPs, $\hat{\Pi}$) to obtain the following test statistic that has an asymptotic χ^2 distribution with T degrees of freedom

$$N \sum_{t=1}^T \frac{[W_S(\hat{P}, \tilde{\Pi}_{1,t}) - W_S(\sigma_S(\hat{\Pi}), \tilde{\Pi}_{2,t})]^2}{\Omega(\hat{P}, \hat{\Pi}) + \Omega(\sigma_S(\hat{\Pi}), \hat{\Pi})} \implies \chi^2(T). \quad (30)$$

The advantage of using many randomly drawn POPs $\{(\tilde{\Pi}_{1,t}, \tilde{\Pi}_{2,t})\}$ is that this provides a strong test of *robustness* of the estimated mixed serve strategy and the calculated DP serve strategy over a wide range of environments that these strategies were not “expecting.” With a large number of random draws T we can tabulate the full distributions of win probabilities for these two serve strategies, allowing us to judge how they perform in relatively extreme situations that differ significantly from the estimated POPs $\hat{\Pi}$ that they were expecting.

Table 10 presents our simulation results and Nash equilibrium test statistics. We drew $T = 500$ simulated POPs and calculated the game win probability for each of the four serve strategies (which were fixed at the empirical estimates and were the same for all simulation draws). We see that the mean win probabilities for the DP strategy (last column) or semi-DP serve strategy (second to last column) are typically close to each other but in all cases are significantly higher than the mean win probability of the reduced form estimate of the server’s mixed serve strategy (first column). Note that the numbers in parentheses are the estimated standard errors of the simulated sample of 500 win probabilities, not the standard errors of the mean of these simulations. The latter standard error is of course equal to the reported values divided by $1/\sqrt{500} = .0447$. Thus, we find *very statistically significant* increases in the win probabilities for the DP serve strategies relative to the reduced-form estimate of the servers’ mixed serve strategies.

Below the estimated win probabilities we report the P values of our omnibus test of Nash equilibrium (30) which has a χ^2 distribution with $T = 500$ degrees of freedom under the null hypothesis. In all cases the test statistics are huge (typically in the order of billions) so we round the reported P values to 0, indicating decisive rejections of the hypothesis that servers’ mixed strategies are consistent with Nash equilibrium. Thus, we can construct alternative *robust* serve strategies with significantly higher win probabilities, contradicting the key restriction of Nash equilibrium (23) that there is no deviation strategy that results in a higher probability of winning.

A surprising finding from our simulations is that large improvements in win probabilities can be obtained even if a server were to switch to a demonstrably suboptimal serve strategy, namely, the fully myopic serve strategy where the server’s objective is to maximize the probability of winning each serve, ignoring the option value of a second serve in the event of a faulted first serve. There is additional improvement in adopting a myopic serve strategy, i.e. one where we solve the two period DP problem to maximize the probability of winning each point but ignoring

Table 10: Improvements in simulated win probabilities for selected elite pro server-returner pairs

Player pair	Reduced form	Fully myopic	Myopic	Fully dynamic
Server → Returner	$W(1,1)$	$W(1,1)$ P-value	$W(1,1)$ P-value	$W(1,1)$ P-value
Roger Federer → Rafael Nadal	.821 (.035)	.850 (.037) 0	.888 (.029) 0	.890 (.028) 0
Rafael Nadal → Roger Federer	.798 (.045)	.830 (.052) 0	.870 (.049) 0	.867 (.048) 0
Roger Federer → Novak Djokovic	.810 (.028)	.816 (.049) 0	.865 (.036) 0	.869 (.033) 0
Novak Djokovic → Roger Federer	.776 (.027)	.843 (.038) 0	.856 (.032) 0	.861 (.034) 0
Rafael Nadal → Novak Djokovic	.704 (.042)	.830 (.059) 0	.888 (.071) 0	.886 (.073) 0
Novak Djokovic → Rafael Nadal	.838 (.029)	.929 (.023) 0	.908 (.043) 0	.931 (.022) 0
Novak Djokovic → Andy Murray	.780 (.040)	.890 (.034) 0	.893 (.031) 0	.893 (.032) 0
Andy Murray → Novak Djokovic	.709 (.045)	.833 (.056) 0	.858 (.057) 0	.858 (.062) 0
Pete Sampras → Andre Agassi	.857 (.031)	.921 (.052) 0	.940 (.028) 0	.939 (.027) 0
Andre Agassi → Pete Sampras	.843 (.035)	.895 (.059) 0	.921 (.041) 0	.920 (.042) 0

the option value of winning or losing the point on the subsequent play of the game. However in all cases except Djokovic serving to Nadal, we see that there is negligible improvement in win probabilities from adopting a fully dynamic serve strategy. For all of the other servers, we would be unable to reject the hypothesis that the mean win probabilities for the myopic and full DP serve strategies are the same using a two sample t-test.

Recall that for Djokovic serving to Nadal, we find that the fully dynamic DP specification was the best fitting specification according to the AIC model selection criterion. Though this suggests that Djokovic is behaving “as if” he had solved the full DP problem to determine his serve strategy against Nadal, recall that this was with respect to distorted subjective beliefs about the POPs. When we solve the DP problem using a random sample of POPs drawn from the asymptotic distribution centered at the reduced form estimate of the POPs, $\hat{\Pi}$, we find that Djokovic can

improve his mean probability of winning by nearly 10 percentage points: from .838 for the reduced form estimate \hat{P} of his serve strategy to .931 for the DP best response strategy $\sigma_S(\hat{\Pi})$.

However the case of Djokovic serving to Nadal is a puzzle: why does the demonstrably suboptimal fully myopic serve strategy result in a mean simulated win probability, .929, that is nearly equal to the mean simulated win probability of the fully dynamic serve strategy, .931? We believe the explanation is that the calculated serve strategies are pure strategies, and so there are sets of POPs that imply the same strategy. The reduced form estimates of the POPs for Djokovic serving to Nadal are such that the full DP serve strategy happens to coincide with the fully myopic strategy over a majority of the (x, m) states of the game. Altogether, the full DP and fully myopic serve strategies have the same win probabilities under $\hat{\Pi}$, at least up to 5 significant digits, .93683.

We conclude that for most cases optimizing in the point game is nearly equivalent to optimizing over the game, and in all cases we analyze the consequences of ignoring dynamics between points of the game is small. In other words, tennis players can safely focus on winning each point and be assured that the resulting strategy will either be exactly optimal or nearly so. However our results also demonstrate that for most of the servers, there is a significant reward to being able to solve at least a two period DP that takes into consideration the option value of the second point. We have shown that serve behavior does differ significantly across first and second serves in a way that suggests players are in fact taking the option value of the second serve into consideration in how they determine their strategy for their first serve.

We conclude this section by providing some graphs that provide more insight into the nature of the improvements in win probabilities provided by our DP “best response serve strategies.” The overall conclusion is that the DP (and myopic DP) do not just result in higher *mean* win probabilities, they result in distributions of simulated win probabilities that *first order stochastically dominate* the distribution of win probabilities implied by the estimated reduced form mixed serve strategies for the case of Nadal serving to Djokovic. Figure 3 graphs four CDFs: the black CDF is the distribution of win probabilities implied by the reduced form estimate of the serve strategy, the blue CDF is the one resulting from the fully myopic serve strategy, the red line corresponds to the myopic DP serve strategy, and the green line is the CDF of win probabilities resulting from the full DP serve strategy.

We also plot a CDF marked with stars (*). This is the CDF of win probabilities resulting

from using the structural estimate of the fully dynamic server strategy (allowing for trembles, i.e. with the estimated $\hat{\lambda}$). We see that this CDF nearly overlaps the CDF for the reduced form serve strategy. This shows that the maximum likelihood estimates of the fully dynamic structural model results in a serve strategy that closely approximates the reduced form serve strategy. Paradoxically, though by construction the objective of the fully dynamic structural model is to maximize the probability of winning the game, the maximum likelihood estimation is forced to distort the POPs in order to rationalize the observed serve behavior. The distorted POPs are sufficiently far away from the true POPs that even though DP is used to compute the serve strategy, it does not constitute a best response to the true POPs, but only to the distorted subjective POPs. The structurally estimated DP serve strategies are suboptimal when we evaluate win probabilities with POPs that are random perturbations about the true values.

Figures 3 and 4 show that our conclusions about the robustness and stochastic dominance of the full DP and myopic DP serve strategies hold for the other elite pro servers, and provide further insight into the performance of the different serve strategies we analyzed. In Figure 3 we see that the fully myopic serve strategy stochastically dominates the mixed serve strategy Nadal actually uses, and by a significant margin. The red and green CDFs are the distributions of win probabilities for the myopic DP and full DP serve strategies and these in turn stochastically dominate the blue CDF for the fully myopic serve strategy. The red and green CDFs are nearly the same, which is an indication that the GMC nearly holds in this case.

Figure 4 plots the CDFs of simulated win probabilities for the case of Federer serving to Djokovic. We see a gap between the CDFs of win probabilities for the myopic DP and full DP serve strategies, which is an indication that the GMC does not hold in all states of the game given the estimated POPs for Federer serving to Djokovic. We also see a smaller gap between the CDF of win probabilities for Federer's mixed serve strategy relative to the DP serve strategies, which indicates that Federer's serve strategy is less suboptimal when serving to Djokovic compared to when Nadal is serving to Djokovic.

We conclude this section by providing intuitive descriptions of the optimal serve strategies that we have calculated by dynamic programming. The full serve strategy is embodied in the serve probabilities $P(d|x, m)$ which we have indicated is a pure strategy and we have also denoted this as $\sigma_S(\hat{\Pi})$ above to emphasize that it was calculated using our reduced form estimate of the

Figure 3: Distribution of win probabilities, Rafael Nadal serving to Novak Djokovic

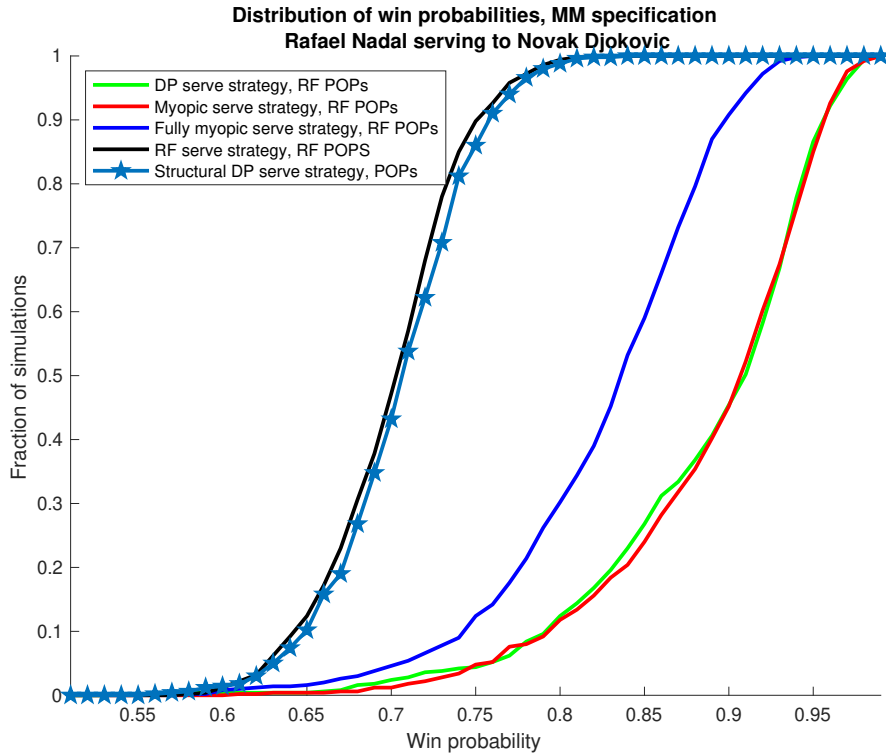
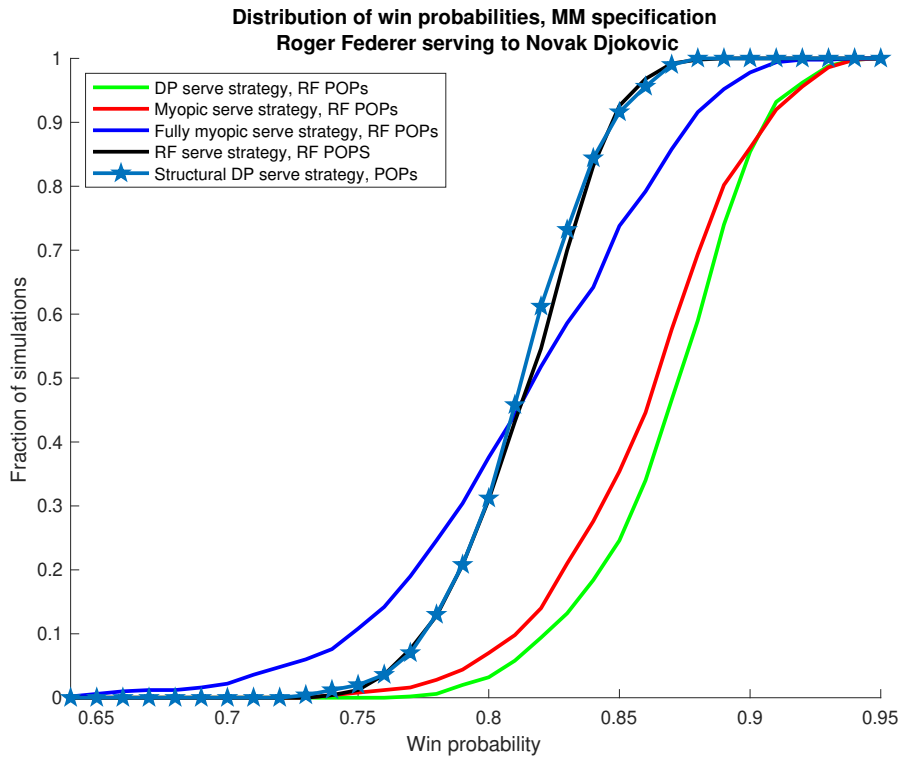


Figure 4: Distribution of win probabilities, Roger Federer serving to Novak Djokovic



POPs, $\hat{\Pi}$. These strategies simply require the server to choose the direction with the highest conditional win probability in every stage of the game.

For example in the case of Djokovic serving to Nadal the full DP serve strategy generally entails serving to Nadal's right (i.e. backhand since Nadal is a lefty) on first serves, whereas on second serves, the optimal direction depends on whether Djokovic is serving to the deuce or ad court. To the deuce court, he should serve to Nadal's backhand, whereas to the ad court, he should serve to Nadal's forehand. In other words, Djokovic should hit his second serve wide.

Under the optimal strategy, the loss from one shot deviations from the recommended serve direction are typically small. For example, on first serves, the conditional probability of winning if Djokovic serves to Nadal's right (the optimal choice) is 0.937, but the worst choice, to serve to Nadal's body, still results in a conditional win probability of 0.921. The reason the loss is not larger is that the full DP strategy entails an automatic recovery from "mistakes" at subsequent stages of the game and provided Djokovic follows the strategy most of the time, the expected losses from an occasional deviation from the optimal strategy are not too large.

Conversely, when we consider the conditional win probabilities implied by the suboptimal mixed reduced form strategy that Djokovic uses, we see bigger gains at each stage. For example at the first serve of the service game, the direction with the highest win probability is to Nadal's left, 0.854. The direction with the lowest win probability is to Nadal's body, 0.831. Thus, there is a bigger penalty in terms of forgone win probability from serving to Nadal's body, yet Djokovic serves to Nadal's body with probability 0.155. Though the gain in win probability, *on the first serve* to serving to Nadal's left is only 0.023, this presumes that Djokovic reverts to his suboptimal mixed serve strategy for the remainder of the game. Dynamic programming exploits the "profitable deviations" in serve directions at every stage of the game and these gains cumulate to a much larger overall increase in win probability at the start of the game. As we showed in Table 10 the total gain in win probability from switching from his current mixed serve strategy to the optimal DP serve strategy is nearly 10 percentage points, from .831 to .938.

5 Conclusions

There is substantial evidence against Nash equilibrium and minimax play in laboratory experiments: see, e.g. Brown and Rosenthal (1990) and Camerer (2003). However a standard critique is that laboratory subjects are not sufficiently trained and incentivized to behave sufficiently closely to the predictions of game theory. The influential study by Walker and Wooders (2001) concludes that “the theory has performed far better in explaining the play of top professional tennis players in our data set.” (p. 1535). Similar results have been found in to other sports such as soccer (see, e.g. Chiaporri, Levitt, and Groseclose (2002)) who study the direction of penalty kicks. The general conclusion is encapsulated in the title of the study by Palacios-Huerta (2003), “Professionals Play Minimax.” In contrast, we show that the serve strategies of elite tennis pros are inconsistent with the minimax prediction. Though they use mixed strategies, the probability of winning is not the same for all serve locations — the key restriction of the Nash equilibrium/minimax solution.

There has also been considerable work on testing for serial independence in serve directions as an additional implication of mixed strategy equilibrium. We argue that serial dependence, which has been found in many previous studies including Walker and Wooders (2001), is not necessarily inconsistent with equilibrium play when we account for *muscle memory effects* that reflect natural improvement from repeating recently performed actions. Our empirical analysis confirms that muscle memory effects are important and can induce both positive and negative serial correlation in serve directions. We contribute to the theoretical literature by introducing a new dynamic model of tennis that accounts for muscle memory effects that demonstrates that serial correlation *per se* is not evidence against equilibrium play.

Our empirical analysis exploits a new source of data, the Match Charting Project, that allows us to analyze a large number of professional tennis matches at the level of individual server-returned pairs. Unlike previous analyses, we have also used a feature of the MCP data, which records body serves, in addition to the L vs. R serves that have been the focus of the previous literature. Tennis players and coaches consider body serves to be an important component of an optimal server strategy, a view supported by our analysis, since they are used frequently in the data and in the calculated optimal serve strategies.

The main innovation in our empirical analysis is to provide new, more powerful tests of

Nash equilibrium. We have introduced an omnibus Wald test for equal win probabilities for all serve direction that decisively rejects the hypothesis of equal win probabilities for the 10 elite professional server-returned pairs we analyzed. We also introduced an alternative direct test of the key implication of Nash equilibrium: namely, that there is no deviation strategy that can strictly improve the payoff of the players. Using numerical dynamic programming and our econometric estimates of the point outcome probabilities (POPs) that capture the probabilistic outcomes of serves to each possible direction, we have been able to decisively reject the hypothesis that the observed mixed strategies of these elite pro servers are best responses to their opponents. We have used dynamic programming to construct numerical best responses and we show that these best responses significantly increase the probability of winning a service game. We have used stochastic simulations to show that our calculated deviation serve strategies are *robust*. That is, they result in significantly higher win probabilities even if the true environment, as captured by the POPs, deviates from the estimated POPs which we used to calculate them — the environment that these strategies were “expecting.”

Our conclusion that many elite tennis pros fail to discover and play serve strategies that are best responses to their opponents may seem surprising and is clearly contrary to the consensus in the literature noted above. Of course, the reader may be suspicious that there is something wrong with our analysis, since it seems it should not be possible that pros have not exploited every possible angle of the game given that the stakes are so high. We believe that we have convincing evidence of suboptimal serve strategies, but the ultimate test would be to run field experiments to verify whether our DP serve strategies really do deliver the increased win probabilities that we predict. However these gains depend on how quickly the returner can recognize and adapt to a change in the server’s strategy. We can artificially randomize serve directions to hinder the returner’s ability to detect a change in strategy, but doing this reduces the potential gain in win probability. The issues raised by the possibility of learning and adaptation to changes in strategy are beyond the scope of this analysis. Similar to WW, our conclusion is based on a key *stationarity assumption* that all learning and strategy experimentation has already taken place, that strategies do not change across games, and if they are in equilibrium, it is unique. The purpose of our analysis is not to give service advice to tennis pros, but to provide a new approach to testing the hypothesis of Nash equilibrium play in a dynamic game.

Though we are convinced that many of the elite pro tennis players are not playing best responses, we are not entirely sure why they have failed to discover and implement best responses. Clearly the rewards to doing so are very high. The usual presumption in economics and much of the previous literature on tennis is that when there are high rewards we can expect to see behavior that is consistent with Nash equilibrium. Implicit in the Nash equilibrium concept is an undefined process of learning and experimentation and the presumption that highly motivated agents will not stop learning and experimenting until they come to a rest point where further adjustments in strategies have no incremental payoff. Any such rest point is a Nash equilibrium.

An alternative to the Nash equilibrium hypothesis is the principle of *satisficing* of Simon (1956): “Both from these scanty data and from an examination of the postulates of the economic models it appears probable that, however adaptive the behavior of organisms in learning and choice situations, this adaptiveness falls far short of the ideal of ‘maximizing’ postulated in economic theory. Evidently, organisms adapt well enough to ‘satisfice’; they do not, in general, ‘optimize.’ ”

We estimated dynamic structural models of serve choices to gain insight into why these elite pro servers fail to use Nash equilibrium strategies. These models rationalize serve strategies as best responses, but instead of to the true POPs, they are best responses to distorted, subjective POPs that deviate significantly from the true POPs. We have no explanation for this failure of rational expectations on the part of elite pro servers, but it constitutes our best explanation for the behavior we observe. We also introduced a generalized version of the monotonicity condition of Walker et al. (2011) and confirmed empirically that it typically holds in tennis, even in the presence of “muscle memory” that can result in serial correlation in serves. It follows that in contrast to board games such as chess or Go, optimal play in tennis does not take heroic mental calculating ability. When the generalized monotonicity condition holds, the tennis server can maximize the chance of winning the overall game by solving a much simpler two period DP problem that maximizes the probability of winning each point rather than having to solve a full (infinite horizon) DP problem that maximizes the probability of winning the overall game. We conclude the disequilibrium play in tennis we find is unlikely to be due to the inability of top tennis players to do the relevant mental calculations; for instance, it does not take complex calculations for Djokovic to hit more first serves to Nadal’s backhand or hit more second serves wide.

Thus, our findings lead us to conclude that even the elite pro tennis players may have inadequate statistical knowledge or an inadequate mental model of the POPs, the point outcome probabilities that implicitly embody their own strengths and weaknesses as tennis players as well as their opponents. The rising industry of *sports analytics* may lead to new awareness and changes in behavior that motivates tennis players to change their strategies to gain advantages over their opponents. The steady state outcome of such learning and experimentation could well be something that more closely approximates Nash equilibrium play.

We note that we are not the first study to have provided evidence that suggests highly compensated and motivated sports professionals may not be behaving optimally. There is the famous book *Moneyball* by Lewis (2003) that showed how sports analytics could improve the performance of entire baseball teams. In football, Romer (2006) provided convincing evidence of suboptimal decisions regarding when teams should go for first downs or kick a field goal, using dynamic programming. We feel that tennis may be another sport where econometrics, dynamic programming, and analytics could affect thinking, change behavior, and help guide players to play in a way that more closely corresponds to the predictions of Nash equilibrium.

A Minimax: Proof of Theorem 1

Formally, the service game described in Section 2.1 is a *recursive constant-sum game* as introduced in Everett (1957). We will apply Theorem 6 in Everett (1957) to secure uniqueness of the Nash equilibrium value $W_S(x, m)$ in every *recursive constant sum sub-game* $\Gamma(x, m)$.

STEP 1: CONTINUOUS STATE TRANSITIONS. Let $\lambda(x', m', x, m, d, s, a)$ be the probability that the state becomes (x', m') when the server chooses location d and serve and spin s , and the returner chooses attention vector a in state (x, m) . This transition chance is implicitly defined by the score state transitions in Figure 1, the muscle memory updating process, and the probabilities $\ell(x, m, d, s)$ and $\omega(x, m, d, s, a)$ introduced in Section 2.1. The precise details of λ are unimportant for the current proof,²⁷ only that it inherits continuity in (s, a) from ℓ and ω .

STEP 2: A UNIQUE EQUILIBRIUM VALUE $W_S(x, m)$. For any function $v : \{1, 2, \dots, 38\} \times$

²⁷ Precisely defining λ and establishing continuity in (s, a) is trivial. We suppress the tedious details to save space.

$\{l, r, b\}^2 \mapsto [0, 1]$ with $v(38, m) = -1$ and $v(37, m) = 1$ (i.e. server payoffs for the terminal loss and win states), define the static zero-sum game with the same strategy sets as in our original game and server payoff:

$$u(d, s, a|x, m, v) \equiv \sum_{(x', m')} \lambda(x', m', x, m, d, d, a) v(x', m') \quad (31)$$

Let \mathcal{B} be the set of probability distributions over $\{l, r, b\} \times \mathcal{S}$ α be receiver's probability distribution over his attention vector, and let \mathcal{A} be the set of probability distributions over returner attention vectors. Since λ is continuous in (s, a) , u is continuous in (s, a) for any fixed v ; thus by the Minimax theorem in Ville (1938):²⁸

$$\min_{\alpha \in \mathcal{A}} \max_{\beta \in \mathcal{B}} \int u(d, s, a|x, m, v) d\beta(d, s) d\alpha(a) = \max_{\beta \in \mathcal{B}} \min_{\alpha \in \mathcal{A}} \int u(d, s, a|x, m, v) d\beta(d, s) d\alpha(a)$$

Altogether, the recursive sub-game $\Gamma(x, m)$ meets the premise of Theorem 6 in Everett (1957); and thus, there exists at most one value $v^* : \{1, 2, \dots, 38\} \times \{l, r, b\}^2 \mapsto [0, 1]$ in each $\Gamma(x, m)$. In other words, the equilibrium service game win probability $W_S(x, m) = v^*$ is uniquely defined.

STEP 3: EXISTENCE OF MPE MIXED STRATEGIES. Recall that $\mathcal{S} \subset \mathbb{R}^2$ is non-empty, closed and bounded (and so, compact), serve directions are discrete and finite, and attention a is in the unit simplex in \mathbb{R}^2 (non-empty and compact). That is, the strategy triple (d, s, a) is restricted to a non-empty compact set. And since the payoff function (31) is continuous in (d, s, a) for any function v , the function $u^*(d, s, a|x, m, v^*(x, m))$ is continuous in (d, s, a) . Thus, there exists a mixed strategy Nash equilibrium (β^*, α^*) of the static game with payoffs u^* by Glicksberg (1952). Since these strategies constitute a Nash equilibrium of the *static* game, they only depend on (x, m) , and by construction they constitute a Nash Equilibrium in sub-game $\Gamma(x, m)$. Altogether, the mixed strategies $(\beta^*(x, m), \alpha^*(x, m))$ constitute a MPE of the service game of tennis. ■

B Muscle Memory and Serial Correlation in Serve Locations

We now explore serial correlation in serve location choices in MPE. We do this in a simple version of the model, but the core insights remain valid in the general model.

²⁸ For an English translation of Ville's Minimax Theorem see Raghavan (2009) page 749.

B.1 Three Sources of Serial Correlation in Serve Locations

Assume that attention to serve location d , only directly impacts win rates if the server chooses location d , i.e. $\omega(m, d, s, a) = \omega(m, d, a^d)$, and that $\omega(m, d, a^d)$ is differentiable, strictly decreasing, and weakly convex in a^d for all m . Let speed and spin choices in some MPE be given by $s^*(x, m)$ and define the induced conditional probabilities:

$$\ell^d(x, m) \equiv \ell(m, d, s^*(x, m)) \quad \text{and} \quad \omega^d(x, m, a^d) \equiv \omega(m, d, s^*(x, m), a^d),$$

Recall that $W_S(x, m)$ is the server's chance of winning the service game in state (x, m) and let $w^d(x, m, a^d)$ be the conditional chance that the server wins the service game when choosing location d given that the returner choosing attention a^d at location d . Then for first serves (x odd):

$$\begin{aligned} W^d(x, m, a^d) &= \ell^d(x, m) \left(\omega^d(x, m, a^d) W_S(x^+(x), (m_2, d)) + (1 - \omega^d(x, m, a^d)) W_S(x^-(x), (m_2, d)) \right) \\ &\quad + (1 - \ell^d(x, m)) W_S(x+1, (m_2, d)) \end{aligned}$$

To further simplify, assume an MPE in which the server strictly mixes over l, r with respective chances $\sigma_S(x, m), 1 - \sigma_S(x, m)$ on first serves.²⁹ Since the returner has no direct effect on the muscle memory state, the returner will best respond by setting $a^b = 0$; and thus, we have $a^r = 1 - a^l$. Altogether, the chance that the server wins the service game in state (x, m) for first serves given returner attention a^l is:

$$W(x, m, a^l) \equiv \sigma_S(x, m) W^l(x, m, a^l) + (1 - \sigma_S(x, m)) W^r(x, m, a^r)$$

Further assuming $a^l \in (0, 1)$,³⁰ it must be the case that the returner cannot lower this probability by adjusting a^l up or down; and thus, the returner's MPE attention $a^l(x, m)$ must obey $W_{a^l}(x, m, a^l(x, m)) = 0$, i.e.:

$$\frac{\sigma_S(x, m)}{1 - \sigma_S(x, m)} = \frac{W_{a^r}^r(x, m, 1 - a^l(x, m))}{W_{a^l}^l(x, m, a^l(x, m))} = \frac{\ell^r(x, m) \omega_{a^r}^r(x, m, 1 - a^l(x, m)) \Delta^r(x, m_2)}{\ell^l(x, m) \omega_{a^l}^l(x, m, a^l(x, m)) \Delta^l(x, m_2)} \quad (32)$$

²⁹ A sufficient condition for an MPE with no body first serves is that serving left or right on first serves gives the server a better chance of winning the current point *and* enhances his chances of winning future points: $\ell^b(x, m) \omega^b(x, m, a^b) < \ell^d(x, m) \omega^d(x, m, a^d)$ for all $m, a, d \in \{l, r\}$, and x odd *and* $\ell^d(x, (d'', b)) \omega^d(x, (d'', b), a^d) \leq \ell^d(x, (d'', d')) \omega^d(x, (d'', d'), a^d)$ for all x, d'', a^d and $d, d' \in \{l, r\}$.

³⁰ An assumption that *implies* $a^l \in (0, 1)$ is $\omega_{a^l}^l(x, m, a^l) / \omega_{a^r}^r(x, m, 1 - a^l)$ converging to ∞ as $a^l \rightarrow 0$ and converging to 0 as $a^l \rightarrow 1$.

where $\Delta^d(x, m_2) = W_S(x^+(x), (m_2, d)) - W_S(x^-(x), (m_2, d))$ is the increase in the service game win chance from winning vs. losing the current point on the first serve.

Equation (32) affords a way to formalize equilibrium serial correlation in first serve strategies. Specifically, a sufficient condition for negatively serial correlation is that the server is less likely to serve left following a left serve, i.e. when $\sigma(x, (d, l)) < \sigma(x, (d, r))$, for all odd x and first serve locations d chosen two first serves prior, which by (32) is equivalent to:

$$\frac{W_{ar}^r(x, (d, l), 1 - a^l(x, (d, l)))}{W_{al}^l(x, (d, l), a^l(x, (d, l)))} < \frac{W_{ar}^r(x, (d, r), 1 - a^l(x, (d, r)))}{W_{al}^l(x, (d, r), a^l(x, (d, r)))} \quad (33)$$

Since $W_{ad}^d = \ell^d \omega_{ad}^d \Delta$, inequality (33) compares the product of three separate ratios, and thus, there are three logically separate ways to generate negative serial correlation with muscle memory. One is when muscle memory affects the server's chance of landing a serve in, as follows:

$$\frac{\ell^r(x, (d, l))}{\ell^l(x, (d, l))} < \frac{\ell^r(x, (d, r))}{\ell^l(x, (d, r))} \quad (34)$$

This comparison of likelihood ratios states that the server's relative chance of landing a right serve in is higher following a right serve. While this makes intuitive sense, it is an empirical question whether such short term muscle memory effects exist for elite pro serves.

Muscle memory can also generate (33) if the following inequality holds:

$$\frac{\omega_{ar}^r(x, (d, l), 1 - a^l(x, (d, l)))}{\omega_{al}^l(x, (d, l), a^l(x, (d, l)))} < \frac{\omega_{ar}^r(x, (d, r), 1 - a^l(x, (d, r)))}{\omega_{al}^l(x, (d, r), a^l(x, (d, r)))} \quad (35)$$

For an interpretation of this condition, notice that $\omega_{ar}^r/\omega_{al}^l$ is the the marginal rate of technical substitution (MRTS) between attention at location r and attention at location l . Inequality (35) states that this MRTS is larger following a serve to the right than it is following a serve to the left. This could be the result of a *direct effect* of muscle memory, the MRTS larger following a serve to the right holding the returner's attention constant, or an *indirect effect*, the returner's attention changes following a serve to the right inducing a larger MRTS.

Notice that inequalities (34) and (35) are about the impact of *past* serve locations on the *current* point game ratios ℓ^r/ℓ^d and $\omega_{ar}^r/\omega_{al}^l$. These effects may be present even if the server and returner behave *myopically*, maximizing their chances of winning the current point and ignoring the future. When the players are forward looking, there is third potential source of negative serial correlation; namely:

$$\frac{\Delta^r(x, l)}{\Delta^l(x, l)} < \frac{\Delta^r(x, r)}{\Delta^l(x, r)}$$

which after substituting in for all Δ^d becomes:

$$\frac{W_S(x^+(x), (l, r)) - W_S(x^-(x), (l, r))}{W_S(x^+(x), (l, l)) - W_S(x^-(x), (l, l))} < \frac{W_S(x^+(x), (r, r)) - W_S(x^-(x), (r, r))}{W_S(x^+(x), (r, l)) - W_S(x^-(x), (r, l))} \quad (36)$$

For an interpretation, recall that muscle memory $m = (m_1, m_2)$, where m_1 is the location of the previous first serve and m_2 is the location of the second to last first serve. Now, arbitrarily order serve locations $l > r$ (or vice versa), then inequality (36) states that the increase in the probability of winning the service game from winning vs. losing the *current* point $W_S(x^+(x), m) - W_S(x^-(x), m)$ is strictly log-supermodular in (m_1, m_2) . This necessarily requires muscle memory to depend on the two previous serve locations.

B.2 Example: Serial Correlation in A Linear Model

We now simplify the model further in order to sign the equilibrium serial correlation in serve locations and show that this serial correlation can be strictly negative (or strictly positive), even if the POPs in Definition 1 are independent of muscle memory.

The *two location linear model* removes spin, speed and body serves as choice variables and assumes that ℓ and ω only depend on the prior first serve location, m_1 , and that the conditional win chance ω is linear in attention, i.e. $\omega^d(m_1) = \bar{\omega} - \eta^d(m_1)a^d$ with $\bar{\omega} \in (0, 1)$ and $\eta^d(m_1) \in (0, \bar{\omega})$. Thus, this model is fully determined by the nine scalars: $\bar{\omega}$ and $\ell^d(m_1), \eta^d(m_1)$ for $(d, m_1) \in \{l, r\}^2$. The two location linear model is *log-supermodular* when $\ell^r(l)\ell^l(r)\eta^r(l)\eta^l(r) < \ell^r(r)\ell^l(l)\eta^r(r)\eta^l(l)$ and *log-submodular* when the opposite inequality obtains. The *symmetric two location linear model* further restricts: $\ell^d(m_1) = \bar{\ell}$, $\eta^r(l) = \eta^l(r) = \eta$, and $\eta^l(l) = \eta^r(r) = \hat{\eta}$.

Theorem 2 *Serve locations are negatively (positively) serially correlated in the log-supermodular (log-submodular) two location linear model in any MPE in which the server strictly mixes over first serve locations. The POPs are independent of muscle memory in the symmetric two location linear model.*

STEP 1: SERIAL CORRELATION IN SERVE LOCATIONS. Direct substitution establishes that inequality (33) is equivalent to $\ell^r(l)\ell^l(r)\eta^r(l)\eta^l(r) < \ell^r(r)\ell^l(l)\eta^r(r)\eta^l(l)$ (i.e. log-supermodularity) in the two location linear model; and thus, negative serial correlation obtains. Similarly, under log-submodularity, inequality (33) flips, implying positive serial correlation.

STEP 2: SERIALY INDEPENDENT POPs. The chance of serving in ℓ is a constant by assumption. Routine algebra establishes that the following strategies constitute a MPE:

$$\frac{a^l(r)}{1 - a^l(r)} = \frac{1 - a^l(l)}{a^l(l)} = \frac{\sigma_S(r)}{1 - \sigma_S(r)} = \frac{1 - \sigma_S(l)}{\sigma_S(l)} = \frac{\hat{\eta}}{\eta} \quad (37)$$

Given these strategies, the win chance is $\omega^d(m_1) = \bar{\omega} - \frac{\eta\hat{\eta}}{\eta+\hat{\eta}}$, independent of muscle memory. This implies the continuation value function W_S is independent of muscle memory; and thus, the generalized monotonicity condition (9) holds. Altogether, the service game can be decomposed into a sequence of identical static games. It is straightforward to verify that strategies (37) constitute the unique equilibrium in these static games. ■

The symmetric model is log-supermodular when $\eta < \hat{\eta}$ and log-submodular when $\eta > \hat{\eta}$; and thus, serve locations are generically either negatively serially correlated or positively serially correlated in the symmetric two location linear model, despite the fact that the POPs are independent of muscle memory.

References

- Brown, J. N., & Rosenthal, R. (1990). Testing the minimax hypothesis: A re-examination of o'neill's game experiment. *Econometrica*, 58(5), 1065-1081.
- Camerer, C. F. (2003). *Behavioral game theory: Experiments on strategic interaction*. Princeton University Press.
- Chiaporri, P., Levitt, S., & Groseclose, T. (2002). Testing mixed-strategy equilibria when players are heterogeneous: The case of penalty kicks in soccer. *American Economic Review*, 92(4), 1139–1151.
- Everett, H. (1957). Recursive games. In M. Dresher, A. Tucker, & P. Wolfe (Eds.), *Contributions to the theory of games, in ann. of math. stud., vol. iii* (pp. 47–78). Princeton, NJ: Princeton University Press.
- Gauriot, R., Page, L., & Wooders, J. (2018). *Nash at wimbledon: Evidence from half a million serves* (Tech. Rep.).
- George, S. L. (1973). Optimal strategy in tennis: a simple probabilistic model. *The Journal of the Royal Statistical Society. Series C (Applied Statistics)*, 22(1), 97–104.
- Glicksberg, I. (1952). A further generalization of the kakutani fixed point theorem, with application to nash equilibrium points. *Proceedings of the American Mathematical Society*, 3, 170–174.
- Hsu, S., Huang, C., & Tang, C. (2007). Minimax play at wimbledon: comment. *American Economic Review*, 97(3), 517–523.
- Iskhakov, F., Jørgensen, T. H., Rust, J., & Schjerning, B. (2017). The endogenous grid method for discretecontinuous dynamic choice models with (or without) taste shocks. *Quantitative Economics*, 8, 317–365.
- Iskhakov, F., Rust, J., & Schjerning, B. (2016). Recursive lexicographical search: Finding all markov perfect equilibria of finite state directional dynamic games. *Review of Economic Studies*, 83(2), 658–702.
- Klaassen, F., & Magnus, J. (2001). Are points in tennis independent and identically distributed? evidence from a dynamic binary panel data model. *Journal of the American Statistical Association*, 96(454), 500–509.

- Klaassen, F., & Magnus, J. (2009). The efficiency of top agents: An analysis through service strategy in tennis. *Journal of Econometrics*, 148(1), 72–85.
- Lewis, M. (2003). *Moneyball: The art of winning an unfair game*. W.W. Norton & Company.
- Norman, J. M. (1985). Dynamic programming in tennis-when to use a fast serve. *The Journal of the Operational Research Society*, 36(1), 75–77.
- Palacios-Huerta, I. (2003). Professionals play minimax. *Review of Economic Studies*, 70, 395–415.
- Raghavan, T. (2009, 01). Zero-sum two person games. *Handbook of Game Theory*, 2.
- Rive, J., & Williams, S. C. (2011). *Tennis skills & drills*. Human Kinetics, Inc.
- Romer, D. (2006). Do firms maximize? evidence from professional football. *Journal of Political Economy*, 114(2), 340–365.
- Sackmann, J. (2013). *The match charting project*. Retrieved from TennisAbstract.com
- Simon, H. A. (1956). Rational choice and the structure of the environment. *Psychological Review*, 63(2), 129–138.
- Ville, J. (1938). Note sur la théorie générale des jeux où intervient l'habileté des joueurs. In *Traité du calcul des probabilités et de ses applications*.
- von Neumann, J. (1928). Zur theorie der gesellschaftsspiele. *Mathematische Annalen*, 100, 295–320.
- Vuong, Q. (1989). Likelihood ratio tests for model selection and non-nested hypotheses. *Econometrica*, 57-2, 307–333.
- Walker, M., & Wooders, J. (2001). Minimax play at wimbledon. *American Economic Review*, 91(5), 1521–1538.
- Walker, M., Wooders, J., & Amir, R. (2011). Equilibrium play in matches: Binary markov games. *Games and Economic Behavior*, 71, 487–502.
- Wiles, J. (2006). *Mixed strategy equilibrium in tennis serves* (Tech. Rep.).