Is Game Theory Real?

Player 2; Player 3
Computer Science Department at University of British Columbia

December 22, 2011

Abstract

In this survey paper we will study the application of game theory in professional sports. Are the players really rational and do they play using their theoretical best strategy? Although Nash equilibrium theorem is a well known and widely accepted theory to predict players' strategies, there is still a lack of empirical results consistent with the Nash equilibrium prediction models. Laboratory experiments differ in their conclusions about whether players adopt their equilibrium strategy and they rather reject this hypothesis. Due to the limitation of the laboratory experiments' environment, new studies have been conducted using data from professional games and they support the hypothesis that agents play according to the equilibrium.

1 Introduction

In Game Theory lots of attention is devoted to the notion of Nash equilibrium - the situation when neither of the players can get a better payoff by choosing a different strategy. Indeed, this kind of equilibrium seems to be a valid candidate for the strategies that real people will end up using in any type of game. However, is this how people *really* play the game? We may think that if the game is really complex people just do not have the computational power to find the best strategy (take the game of chess as an example), and thus rely on some suboptimal heuristic solutions. But what about games with the complexity that people can handle? Can the classic Game Theory be applied to them?

Some researchers do try to address this issue, and in this survey we are going to describe the results they got. But first of all, we want to mention the idea that goes throughout our entire paper - the skill level is very important when we talk about the strategies players adopt. Indeed, it is quite reasonable to expect that a novice player may not recognize the optimal strategy, but when we talk about people who play the game all their life it is also reasonable to expect that the best of them should be able to play the best strategy available (otherwise, why someone else can't use their imperfection if the game is simple enough?). But where do we find such people who devote their life and career to playing games? The most obvious answer is - in professional sports.

However, it turns out that out it is not easy to model sports game as a finite game. Indeed, most sport games are highly dynamic in nature (like soccer) - everyone moves around in almost chaotic way, and the action space for each player is vast (and continuous). But researchers still found a nice way to tie it to the classic theory - they model a very specific part of a game, which is relatively easy to analyze. For soccer, they modeled penalty kick. For tennis, they modeled the serve. Note, that even for these very specific parts of the game still a lot of simplifications are needed to model it as a finite game (we will describe them in section 3). In this way, the researchers are able to compare their results to the results of other studies about what strategies people adopt to play finite games.

Our survey is organized as follows:

- First, we present the fundamental study of O'Neill [1]. He was doing experiments on what strategies people play in a laboratory setting when they know about each other's imperfection. Although this paper is not directly related to sport games, it provides a foundation and a good baseline for pretty much all the papers for sports. The difference is that in sports the players are highly skilled professionals, while O'Neill was basically studying novice players.
- Second, we describe the papers related to modeling a game of soccer, more specifically modeling a penalty kick procedure as a finite 2-player game. The actions of both players (a kicker and a goalkeeper) are easily observable, providing rich data for analysis.
- Next, we describe the papers related to modeling a game of tennis, more specifically modeling a serve procedure as a finite 2-player game. Somewhat counter-intuitively, it turns out that there is no clear way to observe the actions of one player and thus all the conclusions about the strategies have to be done in an indirect way, through theory predictions.
- Finally, we discuss the papers's implications and try to identify promising future directions. We put especial emphasis on the idea that Game Theory can potentially benefit real people playing games in the real world.

2 Early Studies

Mixed-strategy Nash equilibrium [2] is unanimity in the game theory to explain and predict players best strategy. Despite its importance, it is not clear whether real-game agents play according to the Nash equilibrium.

Many different experiments were realized in laboratories in order to test Nash equilibrium theorem for two-person zero-sum games (Von Neumann's Minimax theorem) and their results have been mixed and were often rather negative.

In the late 80's, O'Neill [1] presented the result of an ingenious experiment designed to assess players strategies and whether they play according to minimax strategies. This was a famous study due to its novel approach where the game was simple enough so that players might reasonably be expected to adopt a (nontrivial) strategic behavior similar to the minimax strategy. O'Neill pointed out flaws in previous experiments that invalidated them as tests of mixed strategy equilibrium, like dependence on quantitative assumptions about agents utility for money, games without real opponent and so on.

Consequently, he improved experimental design and the understanding of empirical validity of mixed strategy play.

In O'Neill's experiment, the subjects played a matching-card game with four cards only. Each player would choose their cards simultaneously and revealed them to each other. The payoff matrix $(4x4)$ only had two types of outcome (a win or a loss) and the winner was dictated by the combination of cards chosen. He adopted this game because it had the same minimax strategy no matter what (reasonable) utility functions do agents have. Thus, the subjects were freed from having to consider relative magnitudes of money payoff.

He also took care to eliminate the case when a game is completely symmetrical in strategies, since games too trivial to solve would not be testing the full logic of the minimax solution. The game was then played 5250 times between two random subjects.

His results suggested that people could deviate from minimax play since they assumed that their opponents had limited information processing ability, however, the deviation was small and not evident, so their own payoffs wouldn't be diminished.

Besides his novel approach, O'Neills' experiment was also very important because it motivated new experimental tests of the minimax theorem in games with unique equilibrium.

James Brown and Robert Rosenthal [3] were the first to contradict O'Neill's result. They published a paper demonstrating that O'Neill's support for the minimax hypothesis was much weaker than indicated by him and many players displayed statistically significant dependence on the most recent choices of their opponents. Additionally, they claim that O'Neill's analysis was focused on the aggregate level (a pool of all players strategies) instead of individual level. The minimax hypothesis was also rejected by others experiments like Rapoport and Boebel [4].

Despite all criticism about O'Neill's results [1], that was the first known study to question the laboratory testing environments used to test minimax theorem.

However, one important limitation that was not addressed by O'Neill was the fact that in laboratory experiments, despite their controlled structure, individuals are exposed to games and environments different from real-games. Additionally, most of the situations are new to the participants, thus, they may not be able to learn the best strategy during the experiment period.

To overcome both issues, recent studies have been conducted by observing professional players strategies (natural setting).

3 Professional Player Strategies

In professional sports, players are experts, i.e.:, they are the most proficient agents of the situations of their games. This can improve the quality of experiments to test the

Minimax theorem since they can bypass the complexities of learning and agents are supposed to play according to their best strategies.

3.1 Penalty Kicks

A penalty kick can be considered a two players zero sum game, which involves two players - a kicker and a goalkeeper. Both kicker and goalkeeper should choose their action without any knowledge on the other player's choice and they should move simultaneously[5]. Besides, a penalty kick has only two possible outcomes: score or no score. Thus, the payoff for the kicker is the probability of scoring the goal, and the payoff for the goalkeeper is the probability of *not* scoring the goal.

In addition, penalty kick is a good testing environment since players have few strategies available (where to kick or defend the ball) and their actions are observable. The rules are also very clear and the outcome is declared immediately after players choose their strategies.

Palacios-Huerta [5] presented an experiment where he exploits a data set of 1417 penalty kicks and 22 kickers and 20 goalkeepers total in professional soccer games. The data set included very detailed information on all relevant aspects of the play, especially actions and outcomes. The author claims that this was the first study that supports the notion of Nash equilibrium in mixed strategies using real data, however, Chiappori et al [6] published a paper on the same subject one year before with 459 penalty kicks.

The model presented in the study of Palacios-Huerta has a fair complexity - the actions of each player correspond to the region of the goal that they aimed (right, left and center from the point of view of the goalkeeper). If players choose different regions of the goal, then the score rate was essentially 100%. On the other hand, equal choices does not imply a no score outcome - the score rate was still over 60% (indeed, it is well-known that in professional soccer games most penalty kicks end up as a goal).

The model is also sensitive to a kicker's stronger side, i.e., once goalkeeper and kicker choose the same side, the probability of scoring differs whether it is the kicker's stronger side or not, furthermore, it is higher in the first case (however, the author didn't specify the difference). A "stronger side" in this case means that the kicker has better skills in shooting that direction.

Hence, the best strategy for the kicker is to randomize between both sides but shooting more often to his stronger side than to the opposite side. Due to that, the goalkeeper's best strategy is also to randomize but weighting up the stronger side of the kicker he is facing.

The mixed strategy Nash equilibrium prediction was very close to the empirical values (table 3.1).

Two others studies about players strategies in penalty kicks were also published by Chiappori et al [6] and Azar and Bar-Eli [7].

Both studies used a similar model as presented before, where the goal is divided into three distinct regions (left, right and center) but now related to the kicker position. Additionally, kicking to the left is considered the stronger side for all players - data obtained from left-footed players are inverted to respect that pattern.

	$g_L(\%)$	$1 - g_L(\%)$	$k_L(\%)$	$1 - k_L(\%)$
Nash predicted	41.99	58.01	38.54	61.46
Actual frequencies	42.31	57.69	39.98	60.02

Table 1: Nash equilibrium predicted frequencies and actual frequencies (g_L and k_L denote the non-strong sides of the goalkeeper and kicker respectively)

One interesting aspect introduced by Chiappori et al [6] is the probability of a kick to go out or hit the post. The model is then constructed in the following way:

P_S: Is the probability of the goal to be scored if the kicker and the goalie choose the same side $S \in \{R, L\}$.

π_S: Is the probability of the goal to be scored if the kicker choose a side $S \in R, L$ while the goalie chooses the opposite side or center (corresponds to the probability of a kick to go out or hit the post).

μ: Is the probability of the goal to be scored once the kicker decides to kick in the center and the goalkeeper jumps to one side.

And the payoff matrix of the kicker is given by:

		G L	**G** C	**G** R
K	L	P_L	π_L	π_L
K	C	μ	0	μ
K	R	π_R	π_R	P_R

Table 2: Payoff matrix of the kicker given players action

With the following properties:

$$\forall S \in R, L. \pi_S > P_S \tag{1}$$
$$\pi_L \geq \pi_R \tag{2}$$
$$P_L \geq P_R \tag{3}$$

On one hand this model is richer than the first one presented, on the other hand the number of samples analyzed is much smaller (459 penalty kicks, 162 kickers and 88 goalkeepers). Authors again found that their predictions hold using the data they collected on penalty kicks.

Finally, the third experiment was presented by Azar and Bar-Eli [7]. In addition to the previous studies, the authors explored various prediction alternatives (like probability matching) besides Nash equilibrium model and they showed that Nash equilibrium yields the best predictions - which supports the notion that Nash equilibrium is indeed a crucial idea even in real-world games.

3.2 Tennis Services

Above we described how researchers were able to model one single aspect (penalty kick) of a very complex game (soccer) as a finite 2-player game. The same approach was used to model a single aspect of a tennis match - the *serve*.

As you might all know, in tennis game there are two players on the opposite side of the court, and each point starts with one of them (a ”server”) hitting the ball to the other one (a ”receiver”). The simplest model presented by Walker and Wooders [8] allows the server to serve to the left (action L) and right (action R), while the receiver actions are described as *anticipating* the serve to be on the right (action R) or left (action L) - see Figure 1.

The payoff for a server is then the probability of the server’s winning the point after the serve, and the payoff for a receiver is the probability of the server losing a point. Obviously, when both server and receiver play actions RR or LL then the chances of winning the point for the server are lower than if they play RL or LR, which implies that such a game should have a unique mixed strategy equilibrium.

Figure 1:
A Typical Point Game

Server \ Receiver	L	R
L	π_{LL}	π_{LR}
R	π_{RL}	π_{RR}

Figure 1: Outcomes (cell entries) are probability Server wins the point. (From [8]).

However, this modeling turns out to be more complicated than the modeling of penalty kicks. Although tennis as a game in general has much more statistics available (e.g. serve percentage, unforced errors etc.), and some researchers are successfully extracting valuable insights from it [9], it is not easy to understand what actions are *executed* by both players even for a very restricted setting of serving procedure. Indeed, while we can assume that the server puts the ball to the side he/she intended to do it (since the professional players are quite skilled), it is very hard to say where the receiver *expected* the ball to land since the posture of a player remains the same.

Being unable to observe the actions of the receiver, Walker and Wooders [8] decided to check whether or not the professional players adopt the equilibrium strategy in an indirect way. They *assumed* that there is a finite 2-player game with the two actions described above and that the players do play the equilibrium strategy of this game. Then the authors made a prediction based on that assumption - the (empirical) probability of winning the point by serving left or right should be the same for the server (given how often he/she plays R or L).

And indeed, by analyzing a few Wimbledon matches (which yields almost 10000 samples of the game described above played), they found that the prediction actually holds true. However, note that it does not prove that the players do use equilibrium strategy, it just supports that claim. Indeed, Walker and Wooders [8] also found that if we look at point-by-point play, we will find out that the server switch his/her serve direction significantly more often than it would have been if the player was truly randomizing at each serve independently (i.e. the authors did not find *serial independence* in the data). In other words, the professional players tend to alternate the serve (from left to right and vice versa) - slightly, but noticeably. This shows that either the players do not play the optimal strategies, or the model of the tennis serve that Walker and Wooders [8] presented is too simple for describing players' behaviour.

Wiles [10] adopted the second explanation - and tried to augment the model in a way that would explain this lack of serial independence in the data. He assumed that for the receiver it is easier to respond if the current serve is going the same direction as the previous one. This assumption seems to be reasonable enough, since it makes sense that players may be primed by recent actions of their opponents. Thus, the payoffs in each point play may be slightly different from the payoffs in the previous play. Wiles modeled this short-term effect by introducing what he calls a ”timing variable”. As an example, if the previous serve was to the left, the winning probability of the server serving to the left again is decreased by a specific amount determined by this timing variable.

While the specific details of this model are not very important, it should be mentioned that Wiles [10] was able to fit this model to the data, and according to his estimates this timing effect accounts to approximately half of the variance presented in the original Walker and Wooders's model [8]. In short, although this model was able to fit the data better, the players still do switch the direction of the serve more than the model predicts. Again, it could mean two things - the first explanation is that the players do not play optimally, and the second one is that the model is too simple.

4 Discussion and Conclusion

As shown above, the empirical evidence on professional penalty kicks and tennis serve are somewhat consistent with the hypothesis that in real-life situations two skilled players should play the optimal Nash Equilibrium strategy.

These findings differ from most laboratory experiments, which typically reject the assumption that subjects plays according to the theoretical implications of equilibrium play. Despite their importance, many arguments explain the inherit limitations to such studies:

Testing Environment: In experimental settings individuals are often exposed to situations that they have not faced previously and differs from real life situation.

These studies also face the lack of motivation and incentives of subjects.

Payoff Intuition: Some experiments establish numerical payoffs to their games, which are unfamiliar to players and make the game more complicated to analyze and infer better strategies.

Role of Learning: Despite the games simplicity, most of them are new to the subjects and it may not be possible for them to become proficient in the time frame of an experiment [5].

All these limitations can be overcome by observing real life games, thus, their results provide substantial support for the hypothesis that agents do play according to equilibrium when they are skilled enough players. However, while the studies of soccer penalty kicks clearly support the hypothesis, the studies of tennis serve suggest that even the most skilled players did not learn how to randomize their actions well enough independently of what they did before. Alternatively, it is possible that the players actually do play optimal strategies, but the model used by researchers is not complex enough to grasp and describe the behaviour of players.

If the former alternative is true, then the studies on tennis serve described above actually present a foundation for improving professional players' skills - because it literally means that a serve is predictable to some extent, and the players can capitalize on that. And it will be interesting to see if anything like teaching to randomize correctly (in some or another way) is currently present in tennis academies' practices. It is quite possible that teaching new tennis players to use this predictability of a serve can actually make a difference for their success in professional sports.

However, if the latter alternative (that the model is incorrect) is true, then the researchers still need to catch up with the knowledge of skilled players, before Game Theory could give something valuable to the real-world games.

One interesting point yet to be analyzed is the influence that learning has over players strategies. Many studies support that learning games can generate convergence into Nash Equilibria, and various models has been examined for the laboratory experiments [11]. But it would be interesting to learn how people acquire these skills in real-world settings - for instance, how soccer (or tennis) players *learn* the minimax strategy. The studies on strategies adopted by players of different levels could give us invaluable data that could help us to build better models of people's learning. In this way the abstract mathematical theory can potentially help people learn the best strategies faster and better.

References

[1] B. O'Neill, "Nonmetric test of the minimax theory of two-person zerosum games." *Proc Natl Acad Sci U S A*, vol. 84, no. 7, pp. 2106–2109, April 1987.

[2] J. Nash, "Equilibrium points in n-person games," *Proceedings of the National Academy of Sciences of the United States of America*, vol. 36, pp. 48–49, 1950.

[3] J. N. Brown and R. W. Rosenthal, "Testing the minimax hypothesis: A re-examination of o'neill's game experiment," *Econometrica*, vol. 58, no. 5, pp. 1065–81, September 1990. [Online]. Available: http://ideas.repec.org/a/ecm/emetrp/v58y1990i5p1065-81.html

[4] A. Rapoport and R. Boebel, *Mixed strategies in strictly competitive games: a further test of the minimax hypothesis*. Institute of Information Processing and Decision Making, University of Haifa, 1991. [Online]. Available: http://books.google.com.br/books?id=nOMwHQAACAAJ

[5] I. Palacios-Huerta, "Professionals play minimax," *Review of Economic Studies*, vol. 70, no. 2, pp. 395–415, 04 2003. [Online]. Available: http://ideas.repec.org/a/bla/restud/v70y2003i2p395-415.html

[6] P.-A. Chiappori, S. Levitt, and T. Groseclose, "Testing mixed-strategy equilibria when players are heterogeneous: The case of penalty kicks in soccer," *American Economic Review*, vol. 92, no. 4, pp. 1138–1151, September 2002. [Online]. Available: http://ideas.repec.org/a/aea/aecrev/v92y2002i4p1138-1151.html

[7] O. H. Azar and M. Bar-Eli, "Do soccer players play the mixed-strategy nash equilibrium?" *Applied Economics*, vol. 43, no. 25, pp. 3591–3601, 2011. [Online]. Available: http://www.tandfonline.com/doi/abs/10.1080/00036841003670747

[8] M. Walker and J. Wooders, "Minimax play at wimbledon," *American Economic Review*, vol. 91, no. 5, pp. 1521–1538, December 2001. [Online]. Available: http://ideas.repec.org/a/aea/aecrev/v91y2001i5p1521-1538.html

[9] J. R. Magnus and F. J. G. M. Klaassen, "On the advantage of serving first in a tennis set: Four years at wimbledon," *Journal of the Royal Statistical Society: Series D (The Statistician)*, vol. 48, no. 2, pp. 247–256, 1999. [Online]. Available: http://dx.doi.org/10.1111/1467-9884.00186

[10] J. Wiles, *Mixed strategy equilibrium in tennis serves*, 2006. [Online]. Available: http://books.google.ca/books?id=VT6RPgAACAAJ

[11] I. Erev and A. E. Roth, "Predicting how people play games: Reinforcement learning in experimental games with unique, mixed strategy equilibria," *American Economic Review*, vol. 88, no. 4, pp. 848–81, September 1998. [Online]. Available: http://ideas.repec.org/a/aea/aecrev/v88y1998i4p848-81.html