The Tennis Coach Problem:
A Game-Theoretic and Experimental Study

Ayala Arad*

Abstract

The paper introduces a new allocation game, related to Blotto games: each tennis coach assigns his four different skilled players to four positions, and then each team plays all other teams in the tournament. The winning team is the one with the highest total score.

The set of equilibria is characterized and experimental behavior in variants of the game is analyzed in light of an adapted level-*k* model which is based on an appealing specification of the starting point (Level-0). The results exhibit a systematic pattern- a majority of the subjects used a small number of strategies. However, although level-*k* thinking is naturally specified in this context, only a limited use of (low) level-*k* thinking was found. These findings differ from those obtained in previous studies, which found high frequencies of level-*k* reasoning among subjects in various games. Thus, the results illuminate some bounds of the level-*k* approach.

JEL classification: C72, C91

Keywords: level-*k* thinking, tennis coach problem, experimental game theory, colonel Blotto.

* The Eitan Berglas School of Economics, Tel Aviv University, Ramat Aviv, Tel Aviv 69978, Israel. *E-mail address:* aradayal@post.tau.ac.il.

I wish to thank Chaim Fershtman, Ariel Rubinstein and Ran Spiegler for their superb guidance and Eli Zvuluny for his help in programming the experiments. I also acknowledge valuable comments from Anat Bracha, Colin Camerer, Vincent Crawford, Eddie Dekel, Erik Eyster, Gabi Gayer, Gil Kalai, Matthew Rabin, Pedro Rey-Biel, Alvin Roth, and seminar participants at Tel-Aviv University and at the Max Planck Institute of Economics in Jena.

1. Introduction

This paper introduces a new allocation game called the Tennis Coach problem, which captures the essence of some interesting strategic interactions observed in competitive environments. The game is analyzed both theoretically and experimentally and serves as a platform for studying iterated reasoning and non-equilibrium models based on this concept.

1.1 The Tennis Coach problem

Consider a tournament in which each participant plays the role of a tennis coach who is planning to send his team to the tournament. Each team consists of four players with four different skill levels: A+, A, B+ and B, where A+ is the highest level and B is the lowest. The coach's task is to assign his players to positions 1, 2, 3 and 4 (one player to each position). Each team plays against each of the other teams in the tournament.

A battle between two teams includes four matches: a tennis player that was assigned by his coach to a particular position plays once against the player on the other team assigned to the same position. In any match between two tennis players of different levels, the one with the higher level wins and scores one point for his team. When two players with the same level play against each other, the outcome is a tie and each team receives half a point. Thus, a battle between two teams ends with one of the teams winning 3:1 or 2.5:1.5, or in a tie of 2:2. The team's score at the end of the tournament is the total number of points it received in all the battles. The only goal of the coaches is to win the tournament, i.e. to achieve the highest score among all the teams.

The strategic interaction between the coaches will be referred to as "the Tennis Coach problem", or "the coach problem" for short.

1.2 Experimental motivation

The first strategy that comes to mind is the allocation of the tennis players according to their correct ranking (i.e. A+ in Position 1, A in Position 2, B+ in Position 3 and B in Position 4). Such an assignment immediately suggests itself because of its special characteristics (levels and positions are perfectly correlated) and since it is observed in numerous real-life situations. Therefore, this strategy is a natural starting point for iterative reasoning in the coach problem. A coach *i* who believes that this instinctive strategy will be chosen frequently will best-respond to it; a coach *j* who believes that many coaches use *i*'s reasoning will best respond to the strategy chosen by *i* and so on.

A classic game-theoretic analysis of the coach problem ignores the existence of the salient strategy (A+, A, B+, B) and the induced framing effect. Thus, the game's structure and its psychological properties call for addressing solution concepts other than equilibrium, which are based on iterative reasoning. In particular, the coach problem will be used for exploring the concept of level-*k* thinking which has recently become increasingly popular.[1]

Level-*k* non-equilibrium models assume that the population of players consists of several types, each of which follows a different decision rule. *L0* is a non-strategic type who chooses his action naively by following a particular rule of behavior that depends on the context and is determined by the modeler. *L1* best responds to the belief that all other players are *L0*, *L2* best responds to the belief that all other players are *L1*, and so on. Thus, a type *Lk*, for *k*>0, is behaving rationally in the sense that he best responds to his belief regarding other players' actions. However, the belief held by *Lk* is not the "correct" belief as required by Nash equilibrium. Level-*k* models were first introduced by Stahl and Wilson (1994, 1995) and Nagel (1995). Since then, they have been developed extensively and used to explain experimental results in a variety of settings. For example, Crawford and Iriberri (2007b) apply the model to explain behavior in auctions.[2]

Papers that use level-*k* models to explain experimental results usually estimate the frequency of each type in a particular context. The appeal of this approach is due to a finding stated clearly in Crawford and Iriberri (2007b, page 1725): "The estimated distribution tends to be stable across games, with most of the weight on *L1* and *L*2. Thus the anchoring *L0* type exists mainly in the minds of higher types."

Applying the level-*k* approach to explain experimental results requires a reasonable specification of *L0* and of the belief held by type *Lk* in that particular context. Often (though not always) *L0* is taken to be a uniform randomization over the strategy space. In the Tennis Coach problem, the specification of *L0* is intuitively appealing due to the existence of a salient strategy (A+, A, B+, B), which is the natural starting point for iterated reasoning.[3]

---

[1] The term "iterated reasoning" is usually associated with "iterated dominance", although the term is more general and describes a process in which a player applies arguments recursively. In this paper, I do not discuss iterated elimination of dominated strategies since there are no dominated strategies in the coach problem. Thus, throughout the paper I refer to level-*k* thinking as "iterated reasoning".

[2] Some other examples are: Ho, Camerer and Weigelt (1998), Costa-Gomes, Crawford, and Broseta (2001), Crawford (2003) and Costa-Gomes and Crawford (2006). A different model containing similar ideas is introduced in Camerer, Ho, and Chong (2004).

[3] This specification has features in common with the specification in Crawford and Iriberri (2007a, 2007b).

Indeed, this specification of the starting point will turn out to be the best for explaining the experimental data.

Decision rules based on level-*k* reasoning are expected to be reflected in subjects' choices in this game also because, given this anchor (starting point), best responding to an *Lk* type is cognitively simple (as I confirm experimentally). Furthermore, compared to many other level-*k* models, the adapted model in the Tennis Coach problem assumes weaker and more plausible assumptions on subjects' beliefs. Thus, as will be shown in Section 2.4, the typical choice of *Lk* is not only optimal given the belief that **all** (or almost all) other subjects are *Lk-1* types, but is also the unique best response to the belief that the **majority** of subjects are *Lk-1*, or to the belief that the **most frequent** type is *Lk-1* and that the rest of the choices are uniformly distributed.

Since level-*k* types are naturally specified in the coach problem and level-*k* thinking is cognitively simple here (in fact, it provides an escape from the potentially complex strategic reasoning in this game), the level-*k* approach appears to be suitable *a priori*. On the other hand, the strategy space in the game is large enough and the structure of the game rich enough to leave room for other kinds of decision rules which are not based on iterated reasoning (examples will be discussed at a later stage). Therefore, the Tennis Coach problem is an ideal platform for testing the extent to which level-*k* models are capable of explaining behavior in novel settings.

As expected, experimental behavior in the one-shot game was not consistent with any equilibrium predictions. The adapted model of level-*k* reasoning explained only some of the behavior in the game. Patterns based on iterated reasoning were indeed found, but most choices seemed to be driven by other kinds of deliberations. The subjects' ex-post explanations of their decisions supported this finding. The distribution of strategies reflects a low level of reasoning – even the first step of iterated reasoning was not very common and the second and higher steps were almost totally absent. These frequencies are much lower than those reported in the literature for the parallel steps in other games. Thus, the results illuminate some bounds of the level-*k* approach. For other examples of games in which this approach is not successful see Rey-Biel (2008) and the references there.

1.3 Theoretical motivation

The coach problem will be analyzed as a tournament of N players (coaches). Aside from its literal interpretation, one can think of the tournament model as describing the following situation: players are occasionally involved in a two-person interaction (the pairs of players are randomly matched) in which they receive some payoff. However, a player does not wish to maximize the sum of payoffs obtained in the various interactions but rather to have the highest total payoff among the players in the population. The assumption that players care only about their relative ranking in the population, though an extreme one, is consistent with many real-life situations. The set of equilibria in the tournament will turn out to be equivalent to that of the two-person game in which players maximize score. The tournament and this two-person game are not identical (since the best response functions differ for some beliefs) but the strategic reasoning is similar and the equilibria are identical.

As such a two-person game, the coach problem can be seen as an intuitively appealing version of the popular Colonel Blotto game, introduced by Borel (1921). In the Colonel Blotto game, two players simultaneously allocate a fixed number of troops to N battlefields. A player wins a battle if the number of troops he assigns to a particular battlefield is higher than that assigned by his opponent, and each player aims to maximize the number of battlefields won.[4] The game has been widely interpreted as a competition between two players, in which each distributes his limited resources across N tasks and succeeds in a task if he assigns more resources to it than his opponent. A well-known application of the game involves the interaction between vote-maximizing parties in an election campaign, in which the promises made by the parties are modeled as the various ways to divide a homogeneous good and are assumed to determine the outcome of the election. The basic idea is that an individual votes for party X if it has promised him more than party Y. This scenario could also be interpreted as vote-buying.[5]

Whereas in the Blotto game all partitions (and in some versions only discrete partitions) of the total resources are possible, in the coach problem a player is restricted to a finite number of allocations. This does not make the coach problem a special case of the Blotto game, but rather a different and somewhat simpler version, yet one which captures much of its strategic spirit. Moreover, in many cases, the coach problem reflects more realistic assumptions than the Blotto game. For example, a general might not be able to assign any number of troops to a single battlefield and may be restricted by the internal organization of his army to assigning one division to each battlefield, where the divisions

[4] A number of papers have analyzed the game on a theoretical level. Roberson (2006) provides an analysis of the continuous case and Hart (2008) characterizes the equilibrium in the discrete case. Arad and Rubinstein (2009), Avrahami and Kareev (2009), and Chowdhury et al. (2009) analyze experimental behavior in different versions of Blotto games.

[5] See some variants of the promises game in Myerson (1993), Laslier and Picard (2002) and Dekel et al. (2008).

differ in ability and strength. More generally, the coach problem is better suited to competitive scenarios in which human resources are allocated among several tasks.

The Tennis Coach problem is also able to capture the interaction in the campaign promises game, in which promises are made in the form of a list of priorities (an ordering of projects) that a candidate guarantees to adhere to after being elected. If different projects are associated with different groups (each with equal voting power) then declaring the list of priorities is equivalent to the problem of the tennis coach.

Now consider an R&D race in which each of two firms chooses the order of the routes it will follow in trying to solve a particular problem. Each firm wishes to be the first to find the solution. Assume that there are 4 possible routes but only one of them will lead to the solution. A firm's strategy is the order in which its sequential search will be conducted (i.e. which route to follow at each point in time). Interestingly, choosing the order of the search is equivalent to allocating the tennis players in the coach problem.[6] When two firms search according to their chosen ordering, the probability that a firm will be the first to find the solution is equivalent to the number of points earned by a team in the Tennis Coach problem.[7]

Characterizing the set of equilibria in the Tennis Coach problem is quite involved and relies on its special structure, in which any pure strategy has a unique “best response”[8] and the “best response” function induces a partition of the game's 24 strategies into 6 cycles of 4 strategies each (within a cycle, each strategy is the “best response” to the preceding strategy in that cycle). The characterization yields some interesting results. For example, it will be shown that the simplest mixed strategy equilibrium (simple in terms of number of strategies in the support of the equilibrium strategies) involves the use of two pure strategies, with the property that each is the “best response” to the “best response” of the other strategy.

The rest of the paper is organized as follows: Section 2 presents a game-theoretic analysis of the Tennis Coach problem and an adapted level-*k* model; Section 3 describes two experimental studies of two different versions of the coach problem; and Section 4 concludes.

---

[6] The players' skill levels are analogous to the search schedule. For example, assigning A+ to the second position in the coach problem is equivalent to following the second route first.

[7] See a related game in Fershtman and Rubinstein (1997).

[8] A strategy *S* is the “best response” to the strategy *T* if *S* achieves the highest possible score (3 points) when playing against *T*.

2. Theoretical Analysis of the Tennis Coach Problem

2.1 Formal Presentation of the Game

Players and strategies

The players in the game consist of N tennis coaches who participate in a single round-robin tournament. Coaches choose their strategies simultaneously at the beginning of the tournament. A pure strategy in this game is an assignment of the four tennis players in the team, with skill levels A+, A, B+ and B, respectively, to the four positions. Denote A+ by 1, A by 2, B+ by 3 and B by 4. Formally, denote a pure strategy by a four-tuple, which is a permutation of (1, 2, 3, 4), where the j^{th} component is the level of the player assigned to position *j*. An abbreviation will often be used to represent a strategy, where, for example, 2134 will represent the strategy (2, 1, 3, 4). Since any order of the four players is permissible, there are 24 possible strategies in the game.

Scoring

When two teams play against each other, four points are divided between them. A team receives one point when it assigns a better player to a particular position and no points if the other team assigns a better player. Each team receives half a point when the two players assigned to a position are equally ranked.

Let $score(<x_1, x_2, x_3, x_4>, <y_1, y_2, y_3, y_4>) = \left|\{i \mid x_i > y_i\}\right| + 0.5\left|\{i \mid x_i = y_i\}\right|$ be the total number of points earned by a team that uses a strategy $S = (x_1, x_2, x_3, x_4)$ against a team using the strategy $T = (y_1, y_2, y_3, y_4)$. Thus, $score(S,T) + score(T,S) = 4$ for all *S* and *T*.

Note that a team can never score less than one point in a battle against another team since the best tennis player is unbeatable and in the case that he ties, the second-best player cannot lose and at worst will tie. This implies that a team cannot earn more than 3 points in a battle and that there are five possible scores: 3, 2.5, 2, 1.5 and 1.

Payoffs

Each team will play all the other teams in the tournament. The total score of a team that chooses strategy *S* is the sum of points it scores in **all** battles. Each team wishes to score the highest number of points among all the teams in order to win the tournament but does not care about its total score per se. This is in fact characteristic of many real-life situations, in which competitors only care about winning and the total points earned or the gap between the winner and runners-up is only of secondary importance. (This was also characteristic of the experiments reported on later in the paper.) It is assumed that a prize is shared between the winning teams in the tournament. Therefore, a team prefers winning together with M other teams over winning with N>M other teams (this assumption prevents the game from having trivial equilibria in which all coaches win by choosing the same assignment). Thus, in a tournament between two coaches, the payoff structure is simple: unlike the score function which can receive five values, the payoff function can now receive only three (since each coach prefers winning the tournament over a draw and a draw over losing).

Comment: Hamilton and Romano (1998) describe a similar but not identical assignment game (that is able to capture political parties' assignments of candidates in simultaneous multiple-elections, as well as the assignment of tennis players in dual team matches). There are only two teams in their game, each consisting of n tennis players. The outcome of any match between two players is probabilistic and each team wishes to maximize its probability of winning the overall competition. They found that the (generically) unique equilibrium in this game involves uniform randomization over the strategy space. The coach problem analyzed here assumes deterministic scoring in the matches. More importantly, it is defined as a tournament rather than a game of two teams (and hence is not a special case of the former). These differences lead to a completely different analysis and the set of equilibria in the coach problem is much broader than uniform randomization, as will be shown below.

2.2 The Score Function

The possible scores in any battle between two strategies can be presented in a matrix. Presenting the score function in an illuminating way (see the appendix) requires an appropriate choice of the strategy order. This sub-section presents some properties of the score function that help direct us to it.

Partition of strategies into cycles

We say that a strategy *S* **wins** a battle against strategy *T*, if $score(S,T) > 2$. A strategy *S* **defeats** strategy *T* if $score(S,T) = 3$. Given a level $x \in \{1, 2, 3, 4\}$ and an integer $n \in Z$, denote by $x+n$ the level y satisfying $y=x+n$ (mod 4). For any strategy *S*, let $D(S) = D(x_1, x_2, x_3, x_4) = (x_1 - 1, x_2 - 1, x_3 - 1, x_4 - 1)$ be the unique strategy that defeats *S*.

The function D is reversible. Thus, for each strategy S, there is exactly one strategy $D(S)$ that defeats S and exactly one strategy $D^{-1}(S)$ that is defeated by S. If we perform D on S four times, we again obtain S.[9] This implies that the function D induces a partition of the game's strategies into six disjoint cycles of four strategies each.

Following are the basic properties of the score function:

Property 1. $score(S,S)=2$, $score(S,D(S))=1$, $score(S,D^2(S))=2$ *and* $score(S,D^3(S))=3$.

The following property, which states that any strategy that confronts a pair of non-sequential strategies in a cycle scores a total of 4 points, is of particular importance.

Property 2. *For any T and S,* $score(T,S)+score(T,D^2(S))=4$.

Cycles 1 and 2

Although the score function is invariant to any permutation of the positions, some strategies are more salient than others. In particular, the strategy 1234 immediately suggests itself because of its special characteristics (levels and positions correlate perfectly). Moreover, it is a strategy that can be observed in numerous real-life situations. Thus, the cycle that contains 1234 is of particular importance in the experimental part of the study and its role is discussed later in length. Denote 1234 by L_0, $D(L_0)=L_1$, $D(L_1)=L_2$, and $D(L_2)=L_3$. Cycle 1 is denoted as $[L_0, L_1, L_2, L_3]$.

Different notations are used for the other cycles. Thus, for any $i \in \{2,..,6\}$, denote Cycle i by $[S_0(i), S_1(i), S_2(i), S_3(i)]$. For Cycle 2, I choose $S_0(2)=4321$, which is another possible salient strategy. Thus, Cycle 2 is [4321, 3214, 2143, 1432].

Property 3. *If* $S \in Cycle\ 1$ *and* $T \in Cycle\ 2$*, then* $score(T,S)=2$

Thus, any strategy in Cycle 1 ties with each of the strategies in Cycle 2. A pair of cycles with this property will be called **twin cycles**.

[9] $D(D(S)$ is denoted by $D^2(S)$, and $D(D^2(S))$ is denoted by $D^3(S)$.

Cycles 3, 4, 5 and 6

Four other cycles will now be identified and the strategies ordered in a manner that will simplify the analysis. The first strategy in each of these cycles is chosen to be a permutation of 1234 that swaps two tennis players at adjacent levels: *x* and *x*+1. Let $S_0(3)$=1324, $S_0(4)$=4231, $S_0(5)$=1243 and $S_0(6)$=2134.

Property 4. *Cycles 3 and 4 are twin cycles, as are Cycles 5 and 6.*

Property 5. *For* $0 \le k \le 3$ *and* $3 \le i \le 6$*:* $score(L_k, S_k(i)) = 2$*,* $score(L_k, S_{k+1}(i)) = 1.5$*,* $score(L_k, S_{k+2}(i)) = 2$*, and* $score(L_k, S_{k+3}(i)) = 2.5$.

We define Cycles 3, 4, 5 and 6 as being **parallel** to Cycle 1. This term is appropriate since for *i*=3,4,5,6 L_k ties with $S_k(i)$ for any *k* and the score obtained by L_k when played against $S_m(i)$ is close to that obtained by L_k when played against L_m ($\left|score(L_k, L_m) - score(L_k, S_m(i))\right| = 0 \text{ or } 0.5$).

Due to symmetry considerations, any Cycle *i* can serve as the "starting point" for identifying parallel cycles (by identifying the order of strategies in four other cycles, which makes these cycles parallel to Cycle *i*). In this way, the score can be determined for any two strategies. The matrix presentation of the score function appears in the appendix.

2.3 Equilibrium characterization

This subsection characterizes the population equilibrium in the Tennis Coach problem. A distribution of strategies is a population equilibrium if the average score of a strategy in the support of the distribution is at least as high as any other strategy when playing against this distribution. This concept can be seen as an approximation of Nash equilibrium for a tournament with a large number of teams. It generally fits cases in which players are assumed to best respond to **a distribution of pure strategies** (even if the number of players is small) and is particularly natural when players do not know the exact number of participants.

Denote by *P*(*S*) the probability assigned by the distribution *P* to the strategy *S*. There is no equilibrium with *P*(*S*) =1 since any strategy *T* for which *score*(*T*,*S*)>2 earns a higher score than *S*. Thus, the support of an equilibrium contains at least two pure strategies.

Claim 1 *A probability distribution P is a population equilibrium if and only if the average score for all 24 strategies is 2 points.*

Proof. Obviously, if all the strategies score 2 points, then by definition *P* is a population equilibrium. The other direction: All strategies in the support of *P* yield the same average score only if the average is 2 points. The score of any strategy outside the support must be at most 2; however, if some strategy *S* receives strictly less than 2 points, property 2 implies that $D^2(S)$ receives more than 2 points. Thus, all the game's strategies score 2 points ■

Before moving on to a complete characterization of equilibrium, I present several claims concerning simple forms of equilibrium that will clarify the intuition behind the characterization.

Claim 2 *If P satisfies* $P(S)=P(D^2(S))$ *for any strategy S, then P is an equilibrium.*

Proof. By Property 2, each strategy *T* in the game receives an average score of 2 points when played against a pair of non-sequential strategies. Since for all *S*, $P(S)=P(D^2(S))$, the expected score for any *T* is 2 points. ■

Claim 3 *Any equilibrium P with a support contained in a single cycle satisfies* $P(S)\text{-}P(D^2(S))=0$ *for all S.*

Proof. If for some strategy *S*, $P(D^2(S))>P(S)$, then $D^3(S)$ earns more than 2 points. To see this, recall that $D^3(S)$ earns 2 points when played against $D(S)$ and $D^3(S)$ and more than 2 points, on average, when played against *S* and $D^2(S)$. ■

The analysis of equilibrium remains unchanged if 2 points are subtracted from any possible score in the score matrix. Such a transformation implies that in equilibrium there is no strategy with an average score different from zero. For convenience, what follows is analyzed accordingly.

Claim 4 *Any equilibrium P with a support contained in two cycles satisfies* $P(S)\text{-}P(D^2(S))=0$ *for all S.*

Proof. Assume the contrary. Consider $S \in Cycle\ i$ for which $P(S)\text{-}P(D^2(S))=A$ is maximal. Since $D(S)$ earns a positive score *A* when played against strategies in Cycle *i*, it must earn a negative payoff (-*A*) when played against strategies in Cycle *j* in order to reach the equilibrium score (0 points). This can occur only if *P(D²(T))-P(T)=2A* for the strategy $T \in Cycle\ j$, for which *score(D(S),T)*=0.5 (*score(D(S),T)*=2.5 in the original score function). However, A is the maximal difference between the probabilities of non-sequential strategies in a cycle, a contradiction. ■

Recall that a minimum of two pure strategies is used in equilibrium. Claims 3 and 4 add that these two strategies must be non-sequential in the same cycle. In other words, the simplest mixed strategy equilibrium involves the use of two strategies, with the property that each is the "best response" to the "best response" of the other strategy. (It is appropriate to say that *D(S)* is the "best response" to *S* because it is the strategy that achieves the highest score in a battle against S).

We now consider the full characterization of the game's equilibrium. Define: $\Delta = (\delta_1, \delta_2, \ldots, \delta_{12}) \equiv (p(L_2) - p(L_0), p(L_3) - p(L_1), p(S_2(2)) - p(S_0(2)), \ldots, p(S_3(6)) - p(S_1(6)))$

Proposition 1 *A probability distribution P is an equilibrium if and only if:*

$$\begin{pmatrix} \delta_7 \\ \delta_8 \\ \delta_9 \\ \delta_{10} \\ \delta_{11} \\ \delta_{12} \end{pmatrix} = \begin{pmatrix} -(\delta_1 + \delta_3 + \delta_5) \\ -(\delta_2 + \delta_4 + \delta_6) \\ 0.5 \cdot (-\delta_2 - \delta_4 - 2\delta_6 + \delta_3 - \delta_1) \\ 0.5 \cdot (\delta_1 + \delta_3 + 2\delta_5 + \delta_4 - \delta_2) \\ 0.5 \cdot (\delta_2 + \delta_4 + 2\delta_6 + \delta_3 - \delta_1) \\ 0.5 \cdot (-\delta_1 - \delta_3 - 2\delta_5 + \delta_4 - \delta_2) \end{pmatrix}$$

Outline of the proof. In equilibrium, the score earned by any strategy must be zero. Using Property 2, it is sufficient to verify that in any cycle, two arbitrary adjacent strategies, *S* and *D*(*S*), both earn 0 points (which implies that each of the other two adjacent strategies also earns 0 points). The next step is to understand that the points earned by a strategy *S* are determined only by differences between the probabilities of two non-sequential strategies that do not tie with *S*. Solving the system of 12 linear equations (see the appendix) yields the solution given in the proposition. ■

The analysis of equilibrium in this sub-section is equivalent to that of a symmetric mixed-strategy Nash equilibrium in a two-player game, in which **the payoff matrix is the**

score matrix of the Tennis Coach problem. In other words, the analysis also captures scenarios in which each of the two players aims at maximizing his objective score and not just to obtain a higher score than his opponent. In fact, P is a population equilibrium if and only if it is an equilibrium mixed strategy (possibly asymmetric) in this two-player game.

Comments

(I) All population equilibria with a support contained in three cycles belong to the class suggested in Claim 2 in which $P(D^2(S))-P(S)=0$ for any S. This is because there are 6 degrees of freedom in the system. Therefore, if we substitute zero for the 6 variables, we obtain a single solution: $\Delta = 0$. This claim does not hold for equilibria with a support contained in four cycles. By Proposition 1, the following distributions are examples of equilibria for which the condition in Claim 2 is not satisfied:

$$P = (0,0,\frac{1}{4},0,\quad 0,0,\frac{1}{4},0,\quad \frac{1}{4},0,0,0,\quad \frac{1}{4},0,0,0,\quad 0,0,0,0,\quad 0,0,0,0)$$

and $P = (\frac{1}{12},0,\frac{1}{6},0,\quad \frac{1}{12},0,\frac{1}{6},0,\quad \frac{1}{12},0,0,0,\quad \frac{1}{4},0,\frac{1}{6},0,\quad 0,0,0,0,\quad 0,0,0,0)$.

Note that these two examples induce different vectors of the type $\Delta = (a,b,a,b,-a,-b,-a,-b,0,0,0,0)$, which reflects the structure of equilibria with a support contained in the first four cycles.

(II) Consider a **two-player tournament**, in which the players' payoffs are 1 for winning the tournament, 0 for a draw and -1 for losing. It is straightforward to show that in this game, a probability distribution G is a Nash equilibrium mixed strategy if and only if $G(S)=G(D^2(S))$ for any S.

2.4 Best Response Function

This sub-section investigates the best responses to some interesting distributions of choices. In particular, I identify the best responses for distributions that I consider to be natural beliefs and which may be those actually held by coaches in the experiment. Examples of natural beliefs include: "All other coaches will choose S", "Most of the coaches will choose S" and "The most frequent choice will be S". The analysis here will be used for constructing a plausible level-k model.

As intuition suggests, the best response to the belief that "**almost all** other coaches will choose $S_k(i)$" is $D(S_k(i))$. However, given the belief that **all** other coaches will choose

$S_k(i)$, any $D(S_k(j))$ for a parallel Cycle j is also a best response (a coach who chooses $D(S_k(j))$ earns an average score of 2.5 points but wins the tournament since it is the highest score among the coaches). The next proposition refers to the natural belief that "most of the coaches will choose S". The adapted level-k model that will be constructed in Section 2.6 relies on this proposition.

Proposition 2 *If $1>P(S)>0.5$ for some S, then D(S) is the unique best response to P.*

Outline of the proof. Assume without loss of generality that $1>P(L_0)>0.5$. We need to show that no strategy earns as much as L_1. It is enough to show that for any $X \neq L_1$, if $score(X,L_0)=(3-t)$, then $score(X,Y)-score(L_1,Y)\leq t$ for any Y. In other words, X cannot compensate for its inferiority to L_1 when played against L_0 by its superiority when played against some other strategies. The proof covers all the possible strategies X and confirms that the condition on the score is satisfied (see the appendix). ■

Now consider the belief that "all choices will be in Cycle i and the most frequent choice will be S". For such a belief, the optimal choice is not necessarily $D(S)$. For example, if $P(S_0)$=0, $P(S_1)$=0.4, $P(S_2)$=0.3 and $P(S_3)$=0.3, then the optimal choice is S_3, and not S_2. The reason is that the optimal choice, when choices are in a single cycle, is determined by the differences between two non-sequential strategies. The optimal choice in this case is S_{k+1}, for k that maximize $P(S_k)-P(S_{k+2})$.

The last example also demonstrates why $D(S)$ is not necessarily the optimal strategy given the belief that "the most frequent strategy is S". However, it is easy to see, as an implication of Property 2, that $D(S)$ is the optimal strategy for the belief that the most popular choice is S and that the rest of the chosen strategies are uniformly distributed. Essentially, this claim states that $D(S)$ is the best response to a belief that attributes high probability to the strategy S and takes into account some level of uniform noise.

2.5 A Variant of the Game

In the experimental part of the paper, a second version of the game is discussed, which is denoted as Version 2. It differs from the first version only in the method of scoring. Thus, in Version 2, a team receives one point in a battle against another team **only if it wins three matches out of four**. At any other case, it does not receive any points.

In this version of the game, and given a probability distribution P, it is always optimal to choose $D(S^*)$, where S^* is the strategy for which $P(S)$ is maximal. Therefore, Proposition 2 becomes trivial in this context and can be extended to the following proposition: If none of the strategies are chosen more often than S, then $D(S)$ is a best response. If, in addition, none of the strategies are chosen as often as S, $D(S)$ is the unique optimal strategy. Equilibrium analysis also becomes simpler in this version. Thus, the probability distribution P constitutes an equilibrium if and only if, for any S and T in the support, $P(S)=P(T)$ and in any Cycle i, $P(S_0(i))= P(S_1(i))= P(S_2(i))= P(S_3(i))$.

2.6 The Adapted Level-*k* Model

In this sub-section, the equilibrium solution concept is abandoned and an alternative approach is considered in an attempt to account for the experimental behavior in the Tennis Coach problem. The game's structure and its psychological properties call for applying the concept of level-*k* thinking, which is based on iterative reasoning.

Level-*k* non-equilibrium models assume that the population consists of several different types of decision makers and that each type uses a different level of iterated reasoning. *L0* is a non-strategic type who chooses his action naively. *L1* best responds to the belief that all other players are *L0*; *L2* best responds to the belief that all other players are *L1*; and so on.[10] In each game, the specification of *L0* determines the definition of the other *Lk* types in that particular context. Type *L0* is often assumed to choose a strategy by performing a uniform randomization over the strategy space, but there are cases in which *L0* is specified differently. A relevant example is presented by Crawford and Iriberri (2007a) who construct an adapted level-*k* model to explain behavior in hide-and-seek games with non-neutral framing[11]. Their *L0* type instinctively recognizes salient actions[12] and his typical decision rule is taken to be a mixed strategy which puts greater weight on salient actions. Their specification of the naive *L0* type accurately captures a psychological effect that is also relevant in the Tennis Coach problem. Another related specification is that used by Crawford and Iriberri (2007b) in the context of auctions, in which the "truthful *L0*" bids the value that his own private signal suggests.

---

[10] In some cases (e.g. Camerer et al. (2004)), *Lk* is assumed to best respond to a combination of lower types.

[11] The games were introduced in Rubinstein, Tversky and Heller (1996).

[12] Bacharach and Stahl (1997) propose a general framework that captures this idea.

Note that any distribution of choices can be explained trivially by specifying *L0* as a decision maker who chooses according to that particular distribution. A level-*k* model attempts to explain the data primarily through the behavior of *L1*, *L2* or higher types and by considering only a small number of **natural** non-strategic types. In other words, the explanatory power of level-*k* models is based on the typical behavior of the strategic types.

Specification of *L0* in the Tennis Coach problem

The main assumption I make in this subsection is that the natural starting point for iterated reasoning in the coach problem is the salient strategy 1234 (L_0), which is associated with the non-strategic type *L0*. Since this naive strategy is a natural choice, a sophisticated coach might choose to best respond to such a strategy by choosing 4123 (L_1). Forming a belief concerning the opponent's strategy and best responding to it is the first step of iterated reasoning and thus the type who chooses this strategy is denoted as *L1*. An iteration of this process involves best responding to the belief that other coaches will choose L_1. Therefore, *L2* will typically choose the strategy 3412 (L_2) which reflects the second step of iterated reasoning. The highest level of iterated reasoning that this model takes into account is the third iteration which leads to type *L3* choosing 2341 (L_3).[13]

Note that if a coach simply wants to win the tournament and believes that **all** other coaches will choose L_0, then he actually has five possible best responses: L_1, $S_1(3)$, $S_1(4)$, $S_1(5)$ and $S_1(6)$, though the score for $S_1(i)$ against L_0 is less than that for L_1 against L_0. The justification for my definition of types is Proposition 2, which states that if “the majority of the coaches choose *T*” (rather than all the coaches), then the only optimal strategy is $D(T)$. This kind of belief reflects a rough estimation of the opponents' choices and is likely to be more common than the belief that all other coaches choose a specific strategy. Therefore, the assumption made here concerning coaches' beliefs is more plausible than those made in other level-*k* models.[14] In fact, the typical choices of types defined in the model can be sustained also under a different reasonable assumption, according to which type *Lk* best responds to the

[13] This is because the fourth level of iterated reasoning and the choice of *L0* cannot be distinguished. Tennis teams were defined as consisting of 4 rather than 3 players because in previous experimental studies of other games, the fourth level of iterated reasoning was rarely observed, whereas the third level was more commonly observed. This finding justifies the assumption that *L3* is the highest type.

[14] In many other games appearing in the literature (for example, Costa-Gomes et al. 2001), the definition of level-*k* types would be affected dramatically by a transition to this assumption.

belief that choices are uniformly distributed except for one strategy L_{k-1} which is (even slightly) more frequent. This also means that the presence of a uniform noise, which may be interpreted as errors that players make, should not affect the behavior of higher level-*k* types.

Another strategy to be considered as an anchor for iterated reasoning is 4321 ($S_0(2)$). Allocating the players in the reverse order can be viewed as a salient strategy, though a weaker one than 1234. It is likely that non-strategic types would choose this strategy while strategic types might treat it as an anchor for iterative reasoning. Thus, the choice of $S_k(2)$ is considered as a possible outcome of another level-*k* decision rule, based on a different anchor. Clearly, allowing for another kind of level-0 type can only improve the fit of the level-*k* model. [15]

The experimental results will be analyzed in light of the above specification, thus allowing for two possible anchors and two possible types that use each level of reasoning. In other words, all the strategies in Cycle 1 and Cycle 2 are associated with level-*k* reasoning. This specification will turn out to be the best for accounting for the experimental data.

Alternative specifications of *L0*

There are other intuitively appealing specifications of level-0 types. For example, consider a non-strategic type who chooses each strategy in the game randomly and equally often, excluding the strategy 1234 which he chooses more frequently. Given this alternative specification, *L1*, who best responds to *L0*, would choose 4123 as before and hence higher types would also behave as before. Note that from *L1*'s point of view, the interpretation of this *L0* is the same as in the original model, under the assumption that type *Lk* best responds to the belief that the most frequent choice is L_{k-1} and that the rest of the choices are uniformly distributed. The non-strategic type could be specified in a similar manner under the assumption that the strategy 4321 is chosen more frequently than the rest or under the assumption that both 1234 and 4321 are chosen more frequently than other strategies. In this last case, as long as 1234 receives more weight than 4321, the best response to this type would be 4123. Allowing the existence of two non-strategic types, one who gives more weight to 1234 and another who gives more weight to 4321, implies that the two types who use the first step of iterated reasoning (based on the two possible anchors) choose L_1 and

[15] Note that a special property of Cycle 1 and Cycle 2 is that the strategy that defeats *S* is created by "shifting" each tennis player to the position to his right (left) and the last (first) tennis player to the first (last) position. Thus, it is cognitively easy to find *D*(*S*) for any *S* in these cycles.

$S_1(2)$, respectively. Note that the alternative specifications of *L0* above would not change the typical behavior of higher types and hence should not affect the explanatory power of the model. The only possible change that could result is an increase or decrease in the proportion of behavior that can be explained by the level-0 types.

Taking *L0* to be a type who chooses a strategy randomly and uniformly (ignoring the framing) also could be considered intuitively appealing; however, it does not produce any constraint on the *k*-level types for any $k>0$. In fact, all 24 strategies are best responses to this strategy and thus, for any strategy *S* and for any *k*, one can say that *S* is the choice of a level-*k* type. Since this specification does not produce any prediction, I do not treat the uniform randomization decision rule as an anchor of level-*k* thinking.[16] However, I allow for an almost identical *L0* type who chooses randomly and nearly uniformly, that is, he assigns equal probability to all strategies and (even slightly) higher probability to L_0.

3. Experiments

Two studies were designed to test whether the adapted level-*k* model can explain behavior in the game, to ascertain the depth of iterated reasoning in this context and to provide insights on the subjects' reasoning. Study 1 explores behavior in the original coach problem whereas Study 2 investigates a variant of the game which was introduced in Section 2.5.

3.1 Study 1

3.1.1 Experimental Design

I report three different experiments that are based on the Tennis Coach problem introduced in Section 1. The original text used for each experiment appears in the appendix.

Experiment 1

The experiment was carried out in a number of economics courses at Tel Aviv University and at the College for Management in Israel. The students were asked to participate in an experiment at the beginning of the lesson. In keeping with the theoretical framework, the experiment was carried out in the form of a tournament. A total of 113 subjects participated in 7 tournaments of about 16 participants each. Each Subject chose one strategy and then the

[16] Similar reasons lead Crawford and Iriberri (2007a) to avoid specifying *L0* as a type who chooses a strategy randomly and uniformly.

strategy was played **against his classmates' strategies**.[17] The tournament winner was the one whose total score was the highest. The winner in each tournament received a prize of 200 NIS (around $50).

Since there are 24 strategies in this game, a large sample is required in order to obtain a meaningful distribution of strategies. Experiments 2 and 3 were conducted online[18] which enabled collecting 772 additional observations. In the first part of those experiments, each subject participated in a tournament, as in Experiment 1. In addition, Experiment 2 investigated the subjects' understanding of the best response function by testing whether they could optimally respond to simple given beliefs. In Experiment 3, subjects were asked to explain their choices and their decision time was measured. This data may contribute to our understanding of the behavior in the game.

Experiment 2

Students from three undergraduate economics courses in Israel (at Tel Aviv University, Haifa University and Ben-Gurion University) were invited by email to take part in the online experiment within the next few days. 279 students responded and were randomly assigned to play either the original game denoted as Version 1 or a variant of the game denoted as Version 2. Here I report the choices of 131 subjects who were assigned to Version 1 (the results of Version 2 are reported in Study 2). The winner of the tournament in each class won NIS 200. After choosing a strategy subjects answered three questions that tested their understanding of the best response function. They were asked to provide an optimal response to each of the following beliefs: "All other subjects will choose (A, B, A+, B+)", "All other subjects will choose (B+, B, A+, A)" and "Most of the subjects will choose (B, A+, A, B+)". Subjects were told that there is at least one correct answer to each question and that those who answered the questions correctly would win some CD's.

Since the number of students who entered the website and only then decided not to participate was negligible, I conclude that a subject's decision to participate in the experiment

[17] In all the experiments, strategies were not presented in a list in order to avoid order effects. Subjects faced a matrix with four columns representing the different positions and four rows representing players' levels. They allocated the tennis players on their team by marking one box in each row.

[18] All the online experiments reported in this paper were conducted through the website: http://gametheory.tau.ac.il, which was created by Ariel Rubinstein and provides tools for conducting choice and game theoretic experiments.

was no different in character than the decision to participate in a laboratory experiment. Therefore, there is no reason to think that the recruiting method used here attracted a subject pool different from that of conventional laboratory experiments.

Experiment 3

The subjects in this study consisted of 641 students in 14 different courses of game theory and economics, originating from 7 countries.[19] The lecturers in these courses assigned the Tennis Coach problem as a homework task. The website's server recorded the time each subject spent on making the decision (response time) together with the strategy that he chose. Following the decision, subjects were asked to explain why they had chosen the strategy they did. The subjects did not know in advance that they would be asked to explain their choice or that their response time would be recorded. Lecturers were not able to observe the individual decisions made by their students. They did have access to the distribution of choices made, the three winning strategies and the identities of the three winning students. The winners in the tournament did not receive a monetary prize. Nevertheless, they had an incentive to treat the tournament seriously in order to have the honor of being announced in class as one of the winners.

Some may consider the lack of monetary incentives to be a disadvantage of the experimental method since monetary incentives can increase the engagement of the subjects and reduce the noise in the experimental results. However, the evidence in Camerer and Hogarth (1999) suggest that monetary incentives typically have little effect in experiments of this kind.[20] Moreover, the results here turned out to be very similar to those of the other two experiments.

The experimental method used here has several advantages. In particular, the use of the didactic website is a convenient and inexpensive way to collect a large number of observations. The large sample is important in this game in order to obtain a meaningful distribution of choices. It also facilitates the comparison of response times for different strategies and the analyses of the subjects' own explanations of their choices.

---

[19] The US, the UK, Colombia, the Slovak Republic, Argentina, Canada and Brazil.

[20] "The data show that incentives sometimes improve performance, but often don't… In games, auctions, and risky choices the most typical result is that incentives do not affect mean performance, but incentives often reduce variance in responses". Camerer and Hogarth (1999), page 34.

Non-choice data: response time and explanations

As the analysis in Costa-Gomes et al. (2001) suggests, it is possible to draw incorrect conclusions concerning the frequencies of types based on observed choice alone. They used subjects' patterns of information search to interpret their choices in normal-form games. The approach in this paper is to use subjects' response time and explanations to interpret their observed choices.

A subject's explanation of his choice may reveal the decision rule he used and in particular whether it was based on iterative reasoning. Recall that subjects were asked to explain their choices only after making the decision and therefore their choices could not have been affected.

Response time (RT) is defined as the number of seconds from the moment that the server receives the request for the problem until the moment that an answer is returned to the server. This additional information is used to classify strategies in the game as intuitive choices or as an outcome of cognitive deliberation. This method is discussed in Rubinstein (2007), whose main claim is that the RT of choices made using cognitive reasoning is longer than that of choices made instinctively, i.e. on the basis of emotional response. This approach is in line with dual-system theories, such as that in Kahneman and Frederick (2002).

3.1.2 Experimental Results

Table 1 in the appendix presents the data for Experiments 1-3. Recall that each experiment consists of a number of tournaments. I focus on analyzing the aggregate data in each of the experiments. By and large, the main features of the distribution of choices are preserved in the individual tournaments. The following table summarizes the findings concerning level-*k* reasoning in the three experiments.

Choice Percentages	**Strategies**								
	L_0	L_1	L_2	L_3	$S_0(2)$	$S_1(2)$	$S_2(2)$	$S_3(2)$	***Other***
Experiment 1 (n=113) classes, with monetary reward	18.6	9.7	4.4	3.5	12.4	1.8	5.3	0	44.3
Experiment 2 (n=131) online, with monetary reward	10.7	19.1	4.6	2.3	7.6	5.3	6.1	0.8	43.5
Experiment 3 (n=641) online, no monetary rewards	22	10.1	3.3	3.6	8.7	3.6	2.8	2.7	43

In all three experiment, about 56% of the subjects' choices were strategies in the first two cycles, where 37-41% of the subjects chose one of the following three strategies: L_0, L_1 or $S_0(2)$. Strategies in other cycles were chosen far less frequently – almost always by less than 4% of the subjects. The main difference between the results of the three experiments is the "switch" of about 10% in Experiment 2 from strategy L_0 to strategy L_1 relative to Experiment 1 and 3. When this difference is neutralized, the three distributions are no longer significantly different.[21] More importantly, the qualitative findings in the three experiments are similar. Thus: (1) As will be shown below, the distribution of strategies is far removed from equilibrium. (2) The strategy L_0 was one of the most popular choices, which confirms its salience and its role as a potential anchor for iterated reasoning. (3) A significant proportion of choices (43- 44%) was outside the first two cycles and thus cannot be attributed to level-*k* thinking. (4) High levels of iterated reasoning were uncommon. In particular, level-2 strategies were chosen much less often than in other games reported in the literature and level-3 strategies were almost totally absent.

Best responding to a given belief (Part 2 of Experiment 2)

In the second part of Experiment 2, subjects were asked to provide an optimal response to each of the following (possibly noisy) single-point beliefs: 1. All other subjects will choose (A, B, A+, B+) 2. All other subjects will choose (B+, B, A+, A) and 3. Most of the subjects will choose (B, A+, A, B+). 125 out of the 131 subjects participated in this part of the experiment. The following table summarizes the results.

Questions	**1**	**2**	**3**	**1 & 2 & 3**	**At least two out of three**
% that answered correctly:	93	90	89	81	93

Comment: Among those who chose 1234 or 4321, only 12.5% (3 students out of 24) did not answer the three best response questions perfectly. In other words, their possibly naive choice does not indicate that they did not understand the game or did not know how to best respond to simple beliefs.

[21] Applying the chi-square test with respect to eight categories - one for L_0 and L_1 as a category, one for each of the other strategies in cycles 1 and 2 and another for the rest, it was found that there is not significant difference between the frequencies of categories in pairwise comparisons (Experiments 1 and 2: chi-squares=4.81, p=0.68; Experiments 2 and 3: chi-square=7.46, p=0.38; Experiments 1 and 3: chi-squares=8.06, p=0.33).

Explanations (Part 2 of Experiment 3)

A total of 368 subjects in Experiment 3 provided an explanation of their choices (70% out of 526 who were asked). Each of the explanations is classified according to one of the following categories and the proportion of each category is estimated:

1. ***Intuitive choice*** (18%)

This category includes explanations such as: "It was a guess"; "I don't know why"; "It felt right" and "Intuition". 45% of subjects who provided intuitive explanations chose L_0.

2. ***Random choice*** (18%)

This category includes explanations that mentioned the word "random" or similar words. Some of them explained the randomization as an attempt to choose a different strategy from that of other players or to surprise their opponent. The category also includes explanations such as: "It does not matter what I choose because the distribution of choices is practically uniform if I don't know it". Among subjects in this category, 10% chose L_0 and explained that it did not matter what they chose. The other 90% said that they randomized and 19 strategies were chosen by them.[22]

3. ***First step of iterated reasoning*** (10%)

This category includes explanations that describe best responding to the belief that most of the choices will be X (primarily L_0 or $S_0(2)$). 80% of the subjects in this category chose L_1 and 8% chose $S_1(2)$.

4. ***Second step of iterated reasoning*** (three subjects, less than 1%)

This category includes explanations that describe best responding to the belief that most of the choices will be L_1.

5. ***Other strategic decision rules*** (53%)

This category includes explanations such as: "I am mixing good and bad players", "I am sacrificing the weak player in order to win in other positions", "My choice was based on my life experience", "The best players of my opponent were likely to be in the middle positions and therefore I put mine on the edges" and "The player in the first position should be the best one since my opponent will put A in the first position" (or something similar based on some other partial belief). It also includes explanations based on incorrect reasoning (such as "I am trying to achieve a tie") or irrelevant considerations (such as taking into account order effects). Interestingly, only four subjects mentioned the concept of Nash equilibrium in their

[22] Unchosen strategies: 4231, 2143, 1243, 1432 and 1342. Most frequently chosen strategies: 1234 (24%), 2413 (14%), 4321 (13%) and 1324 (6%).

explanation, although many of the subjects had studied game theory. Each of the 24 strategies was chosen by subjects in this category.

The following observations can be made based on the explanations:

Among those who chose L_0, 28% did not explain their choice. Of those who did (83), 65% belong to the intuitive and random choice categories and the rest belong to the category of other strategic rules. These frequencies suggest that L_0 is typically an instinctive choice or the outcome of a low level of sophistication.

In contrast, only 11% of those who chose L_1 did not provide an explanation. Among those who did (41), 75% of the explanations suggest that the subject indeed utilized the first step of iterated reasoning with L_0 as an anchor.

None of the explanations provided for the choice of L_3 included a process of iterated reasoning. Only three explanations (out of 15) for the choice of L_2 explicitly described the use of two levels of iterated reasoning. Thus, the subjects' explanations suggest that many (though not all) of those who chose this category used alternative decision rules rather than high levels of iterated reasoning. Another important finding is that no one who chose a strategy other than L_2 or L_3 explained that he had used two levels of iterated reasoning or higher.

3.1.3 Discussion of Study 1

Can equilibrium explain the results?

In each of the three experiments, the distribution of strategies is not consistent with any equilibrium prediction. Since there is a large set of equilibria in this game, I show that the experimental results do not satisfy the general properties of equilibrium.

One way to see it is by examining the expected score of the strategies presented in Table 1. In Experiment 1, the strategies L_1 and $S_1(4)$ are the clear leaders and the only strategy that comes close to the highest score is L_2. These three strategies were chosen by only 17% of the subjects. Thus, the vast majority of the subjects could have significantly improved their chances of winning by deviating to L_1 or $S_1(4)$. Similarly, the strategy L_1 is the clear leader in Experiment 3 and the only strategy that comes close to it is $S_1(3)$. These strategies were chosen by only 14% of the subjects. In Experiment 2, the best response to the distribution is clearly L_2, which was chosen by less than 7% of the subjects.

In order to demonstrate that minor changes in subjects' choices would not turn the distribution into equilibrium, the following exercise was carried out: In each study, an equal number of subjects was subtracted from each pair of non-sequential strategies in a cycle, thus leaving the choices of 57 subjects in Experiment 1 (50% of the population), 51 subjects in Experiment 2 (39% of the population) and 283 subjects in Experiment 3 (44% of the population). The resulting distributions appear in Table 2. As a consequence of Property 2, subtracting an equal number of subjects from the choices of both S and $D^2(S)$ leaves the best response to the distribution unchanged. Hence, in each of the studies, the resulting distribution is an equilibrium if and only if the original distribution is as well.

In the resulting distribution in Experiment 1, $P(L_0)$=0.28 and $P(L_2)$=0. In order for L_1 to earn an equilibrium score, the total frequency of $S_2(3)$, $S_2(4)$, $S_2(5)$ and $S_2(6)$ has to be 0.56- twice as much as the frequency of L_0. However, these strategies' frequency is less than 0.11. Thus, L_1 would earn more than equilibrium score even if some subjects' changed their choices. In the resulting distribution in Experiment 2, $P(L_1)$=0.43 and $P(L_3)$=0. This implies that L_2 would earn more than the equilibrium score even if all other choices were concentrated around $S_3(i)$, for i=3,4,5,6. However, $P(S_3(\text{i}))$=0, for i=4,5,6 and $P(S_3(3))$ is just 0.06. In the resulting distribution in Experiment 3, $P(L_0)$=0.42 and $P(L_2)$=0. This implies that L_1 would earn more than the equilibrium score even if all other choices were concentrated around $S_2(i)$, for i=3,4,5,6. The argument is strengthened by the fact that $P(S_2(\text{i}))$=0, for i=3,4,5,6.

Level-*k* thinking

As stated by Crawford and Iriberri (2007b): "The estimated distribution tends to be stable across games, with most of the weight on *L1* and *L*2. Thus, the anchoring *L0* type exists mainly in the minds of higher types." The results of the three experiments reflect a low level of sophistication in terms of level-*k* reasoning. Moreover, many choices do not reflect level-*k* reasoning at all and are the result of other types of deliberations.

Generally speaking, the frequency of non-strategic types (level-0) is much higher and the frequency of level-1 types lower than in other related studies; higher types are in fact almost totally absent. The proportion of subjects that actually use a high level of iterated reasoning might be even smaller than that indicated by observed choice since subjects who chose randomly or used decision rules other than iterated reasoning also must have chosen L_2, L_3, $S_2(2)$ or $S_3(2)$. Note that I do not consider the choice of L_0 to be an outcome of four steps of iterated reasoning since in previous studies this level of reasoning was not evident. This is also supported by L_0's low response time (median=125s) compared to the other strategies in Cycle 1 (and compared to strategies outside Cycle 1) and the fact that no one who made this choice explained it as being a best response to L_3.

The subjects' explanations indicate that the only common starting point for iterated reasoning in players' minds was L_0. A secondary and much less common anchor for iterated reasoning was $S_0(2)$. Furthermore, strategies outside Cycle 1 or Cycle 2 have much lower response times than L_1's (median=194s),[23] suggesting that there are no other pure strategies with the same role as L_1.

In the second part of Experiment 2, 81% of the subjects answered all three questions perfectly and 93% answered correctly at least two questions out of three. This indicates that the game is not complex; subjects understand the game and are cognitively able to best respond to a single-point belief, such as the belief that all other choices will be *S*. Thus, we should not expect participants in the game to make many errors when trying to best respond to a belief of this kind.

The high percentage of correct answers to Question 3 implies that subjects also have the correct intuition regarding the optimal response to the belief that most of the subjects (rather than all) will choose *S*. This result is important since the belief that most of the subjects will choose *S* sounds more plausible than the belief that all of them will choose *S*.

In answering Question 1 and 2, almost all subjects chose the best response that **defeats** the strategy (i.e., wins 3 out of 4 matches) assumed to be chosen by other coaches. Only a few chose one of the four pure strategies that score 2.5 points. These findings provide further support for the plausibility of the definition of iterated reasoning used in this game (i.e., that the typical choice of *Lk* **defeats** the strategy chosen by *Lk-1*).

The above results concerning the subjects' ability to respond optimally to a given (possibly noisy) single-point belief suggest that level-*k* choices are not prevalent in the coach problem since subjects do not hold a belief of this kind. The subjects' explanations support this conjecture and suggest that a vast majority of them do not hold any belief at all on the distribution of the other subjects' strategies. However, subjects do attempt to forecast features of their opponents' strategies (such as: A+ is not likely to be assigned to Position 1) and respond according to that partial belief. This finding suggests an extension of the level-*k* approach that allows for iterative reasoning given partial beliefs. Arad and Rubinstein (2010)

[23] An exception is the strategy $S_1(4)$ whose median response time is 181s (chosen by 2.8% of the subjects).

introduce a decision procedure that implements iterative reasoning in several different "dimensions" of the strategy.

3.1.4 Is There an Alternative Level-*k* Model that Account for the Data?

Replacing the specification of *L0* could not provide a better explanation for the experimental results in Study 1. In order to account for the pattern observed in the first cycle, the typical choice of *L0* must be in the first cycle as well. Consider, for example, the specification of the strategy L_3 as the typical choice of the *L0* type (in this case, the strategy L_0 will be interpreted as the typical choice of the *L1* type and the strategy L_1 will be interpreted as the choice of the *L2* type). This alternative specification explains the same fraction of experimental behavior but is less plausible for the following reasons: (1) L_3 it is not salient and is not likely to be a starting point for iterative reasoning, (2) the low response time associated with the choice of L_0 suggests it is an instinctive one and (3) none of the subjects who chose L_0 mentioned in his explanation an attempt to respond to L_3.

Maintaining the specification of *L0* and constructing the hierarchy of types using alternative assumptions on the belief held by type *Lk* does not help in explaining the data. One alternative assumption is that *Lk* believes that all other subjects are lower types than he is, but not necessarily *Lk-1*. This is the assumption made in the cognitive hierarchy model suggested by Camerer, Ho and Chong (2004). In the context of the coach problem, it may be possible to interpret the choice of 4123 (L_1) as a best response to the belief of *L2* that the population consists of *L1* and *L0* types (rather than just *L1*) and that *L0* is more frequent than *L1*. Note, however, that if *L2*'s belief is assumed to even roughly describe the observed distribution of lower types in the experiment, then *L2* would typically choose 3412 (L_2) in Experiment 2, since *L1* is more frequent than *L0*.

In order to apply the cognitive hierarchy model, we also need to assume that the distribution of types follows a one-parameter Poisson distribution. A parameter that produces a distribution with around 20% *L0* choices would imply extremely high percentages of *L1* and *L2* choices and hence cannot account for the data in Experiments 1 and 3. A parameter that produces 10% *L0* choices and 20% *L1* choices would imply an even larger percentage of *L2* choices (more than 20%). Even if we consider *Lk* types in both Cycle 1 and 2, a Poisson distribution does not provide a good fit to the distribution obtained in this experiment.

Comment: Other solution concepts, such as quantal response equilibrium (introduced in McKelvey and Palfrey (1995)), which require that better strategies given the empirical data are chosen more frequently, cannot explain the results. In Experiments 1 and 3, L_1 yields much higher expected payoff than L_0, but L_0 is chosen much more often. Similarly, in Experiment 2, L_1 is significantly more frequent than L_2 but its score is much lower.

3.2 Study 2

3.2.1 Experimental Design

Subjects in this study played a variant of the game denoted by Version 2, which is presented in Section 2.5. Recall that the only difference between the two versions is in the system of scoring. In Version 2, a team scores 1 point only if it wins three matches out of four against another team. This system of scoring makes Version 2 cognitively simpler than Version 1. More importantly, defeating is equivalent to winning here, which implies that L_1 is the **unique** best response to the belief that "**all** other coaches choose L_0". This property of Version 2 makes the above definition of level-*k* types more plausible here. Another nice property of this version is that *D*(*S*) is the unique best response to the belief that *S* is the **most frequent choice**, regardless of the "noise" distribution. Thus, the adapted level-*k* model is even more appropriate in Version 2 than it is in Version 1. Comparing behavior in the two versions would determine whether these differences increase the use of iterative reasoning.

The data reported in this study was collected in two experiments. The subjects in Experiment A consisted of 148 undergraduate economics students in Israel who were assigned to Version 2 in the same online experiment that was reported in Experiment 2. As in Experiment 2, the winner in each tournament received 200 NIS. Experiment B used the method described in Experiment 3 to collect data from 704 students in 14 countries[24] who were studying in 22 different courses.

3.2.2 Experimental Results and Discussion

Table 3 in the appendix presents the distribution of strategies in the two experiments. The following table summarizes the main finding concerning level-*k* reasoning.

[24] The US, Mexico, Brazil, Chile, India, Switzerland, Moldova, Ecuador, France, Brunei Darussalam, Germany, Portugal, Spain and Israel.

Choice Percentages	Strategies								
	L_0	***L_1***	***L_2***	***L_3***	***S_0(2)***	***S_1(2)***	***S_2(2)***	***S_3(2)***	***Others***
Experiment A (N=148) online, with monetary reward	10.1	18.9	6.8	2	4.1	10.8	2.7	1.4	43.3
Experiment B (N=704) online, no monetary reward	18.6	12.9	5.3	2.8	5.8	5.4	2.4	1.7	45

It is straightforward to show that neither distribution of chosen strategies is consistent with any of this version's equilibria. The suggested specification of the level-*k* model is the best possible one to account for the data in the experiment. However, it cannot explain a large proportion of the choices.

Roughly speaking, the distribution of strategies in this version resembles that in Version 1. In fact, there are no significant differences between the distribution of choices in the two versions (chi-square=25.85 and p=0.31 were obtained in a test comparing the data collected through the didactic website and chi-square=31.75 and p=0.11 were obtained in a test comparing the data from the experiment that used monetary incentives). As to the explanations of the subjects, the proportions of the various categories were similar to those in Study 1. Thus, the "guidance" provided by the scoring system and the relative cognitive simplicity of this version of the game did not significantly increase the use of iterated reasoning. Recall that in Version 2, *D(S)* was the unique best response to the belief that "the most frequent choice will be *S*" which includes the belief that "all other coaches will choose *S*". Hence, the resemblance of the results in the two versions provides further support for the plausibility of the definition of types in Version 1.

4. Concluding Remarks

The Tennis Coach problem captures various strategic real-life interactions. Examples include: allocating troops among a number of battlefields, choosing the order of R&D projects to be undertaken, promises in election campaigns, assigning workers to projects in a competitive environment and, of course, assigning players in sports games. The paper's theoretical analysis provides a complete characterization of equilibria in the coach problem. In an attempt to explain the experimental behavior in the game, the equilibrium solution concept is replaced by an adapted level-*k* model, which is based on a natural specification of iterated reasoning in this setting.

Although level-*k* thinking seems to be highly appropriate in the coach problem, the adapted model explains only part of the experimental results and many of the choices seem to be the result of other decision rules not based on level-*k* thinking. Perhaps the most striking result is the low frequency of types that use high levels of iterated reasoning. Even the first step of iterated reasoning is not very common in the two versions of the game and higher steps of reasoning are almost totally absent. These findings are supported by the subjects' explanations. Furthermore, their explanations hint that many of them do not hold a concrete belief over other subjects' choices and certainly do not best respond to the belief that most of the subjects are level-*k* types.

The results in this paper differ from those obtained in previous studies, which found high frequencies of level-*k* reasoning among subjects in various games. I suggest two reasons for this: First, the pure strategies attributed to level-*k* reasoning in the coach problem are only a small fraction of the possible choices in the game. Second, the rich structure of the game triggers other kinds of strategic thinking based on partial beliefs over the opponents' strategies. Further research is needed in order to more clearly identify the circumstances in which the level-*k* approach is successful at explaining the data.

Although iterated reasoning was not triggered as often in the coach problem as in other games reported in the literature, the level-*k* concept may have important roles in this context. Thus: (1) understanding the empirical features of level-*k* reasoning in this game makes it possible to predict the optimal strategy. The robust findings that L_0 and L_1 are the most frequent strategies and that the rest of the strategies are significantly less common, suggest that also in other samples the winning strategy will be L_1 or L_2 (depending on the relative frequency of L_0 and L_1). This prediction does not rely on the exact frequencies of level-*k* strategies and other choices. It is due to the game's structure and the qualitative properties of the empirical distribution.

(2) Even these low frequencies of level-*k* choices may affect dramatically the long run play of the game. Although the distributions of strategies in both studies was far from equilibrium, if subjects were to play the game repeatedly and in each round would internalize the distribution of strategies in the previous round, they might converge to one of the equilibria of the one-shot game. Since subjects may notice the patterns based on iterated reasoning in earlier rounds, they might modify their choices in later rounds accordingly. In particular, I conjecture that in later rounds subjects’ choices would be concentrated in the first cycle. Thus, level-*k* reasoning may turn out to influence not only outcomes of one-shot games, but also the selection of equilibrium in the long run.

References

Arad, A, and Rubinstein A., 2009. Colonel Blotto's Top Secret Files. Working paper.

Avrahami, J, and Kareev, Y., 2009. Do the Weak Stand a Chance? Distribution of Resources in a Competitive Environment. Cognitive Science, 33, 940-950.

Bacharach, M, and Stahl, D., 2000. Variable-Frame Level-n Theory. Games and Economic Behavior 32, 220–246. Previous Version: http://www.eco.utexas.edu/~stahl/vflnt4a.pdf

Borel, E., 1921. La th´eorie du jeu les ´equations int´egrales `a noyau sym´etrique. Comptes Rendus de l'Acad´emie 173, 1304–1308; English translation by Savage, L. 1953. The Theory of Play and Integral Equations with Skew Symmetric Kernels. Econometrica 21, 97–100.

Camerer, C, F. and Hogarth, R, M., 1999. The Effects of Financial Incentives in Experiments: A Review and Capital-Labor-Production Framework. Journal of Risk and Uncertainty 19, 7-42.

Camerer, C, F., Ho, T, H. and Chong, J, K., 2004. A Cognitive Hierarchy Model of Games. Quarterly Journal of Economics 119 (3), 861–898.

Chowdhury, S. M., Kovenock, D. and Sheremeta, R. M., 2009. An Experimental Investigation of Colonel Blotto Games. CESifo Group Munich, Working Paper Series, 2688.

Costa-Gomes, M. A., Crawford, V. P. and Broseta, B., 2001. Cognition and Behavior in Normal-Form Games: an Experimental Study. Econometrica,69 (5): 1193–1235.

Costa-Gomes, M. A. and Crawford, V. P., 2006. Cognition and Behavior in Two-Person Guessing Games: An Experimental Study. American Economic Review 96, 1737-1768.

Crawford, V. P., 2003. Lying for Strategic Advantage: Rational and Boundedly Rational Misrepresentation of Intentions. American Economic Review 93, 133–149.

Crawford, V. P. and Iriberri, N., 2007 [a]. Fatal Attraction: Salience, Naivete, and Sophistication in Experimental Hide-and-Seek Games. American Economic Review 97, 1731–1750.

Crawford, V. P. and Iriberri, N., 2007 [b]. Level-k Auctions: Can a Non-Equilibrium Model of Strategic Thinking Explain the Winner's Curse and Overbidding in Private-Value Auctions? Econometrica 75, 1721–1770.

Dekel, E., Jackson, M. and Wolinsky, A., 2008. Vote Buying: General Elections. Journal of Political Economy 116 (2), 351-380.

Fershtman, C. and Rubinstein, A., 1997. A Simple Model of Equilibrium in Search Procedures. Journal of Economic Theory 72, 432-441.

Hamilton, J. and Romano, R. E.,1998. Equilibrium Assignments in Pairwise Team Contests: How to Form Political Slates and Tennis Teams. European Journal of Political Economy 14 (1), 101-114.

Hart, S., 2008. Discrete Colonel Blotto and General Lotto Games. International Journal of Game Theory, 36(3), 441-460.

Ho, T, H. Camerer, C. and Weigelt, K., 1998. Iterated Dominance and Iterated Best Response in Experimental 'p-Beauty Contests'. American Economic Review 88 (4), 947–969.

Kahneman, D. and Frederick, S., 2002. Representativeness Revisited: Attribute Substitution in Intuitive Judgment. Heuristics & Biases: The Psychology of Intuitive Judgment, Eds: T. Gilovich, T. , Griffin, D. and Kahneman, D, New York: Cambridge University Press, 49-81.

Laslier, J, F. and Picard, N., 2002. Distributive Politics and Electoral Competition. Journal of Economic Theory 103 (1),106-130.

McKelvey, R. and Palfrey, T., 1995. Quantal Response Equilibria for Normal Form Games. Games and Economic Behavior 10, 6–38.

Myerson, R.B., 1993. Incentives to Cultivate Minorities under Alternative Electoral Systems. The American Political Science Review 87, 856–869.

Nagel, R., 1995. Unraveling in Guessing Games: An Experimental Study. American Economic Review 85 (5), 1313-1326.

Rey-Biel, P., 2009. Equilibrium play and best response to (stated) beliefs in normal form games. Games and Economic Behavior 65(2), 572-585.

Roberson, B., 2006. The Colonel Blotto Game. Economic Theory 29, 1–24.

Rubinstein, A. Tversky, A. and Heller, D., 1996. Naive Strategies in Competitive Games. Understanding Strategic Interaction - Essays in Honor of Reinhard Selten, Eds: Albers, W. et al, Springer-Verlag, 394-402.

Rubinstein, A., 2007. Instinctive and Cognitive Reasoning: Response Times Study. Economic Journal 117, 1243-1259.

Stahl, D. and Wilson, P., 1994. Experimental Evidence on Players' Models of Other Players. Journal of Economic Behavior and Organization 25 (3), 309-327.

Stahl, D. and Wilson, P., 1995. On Players' Models of Other Players: Theory and Experimental Evidence. Games and Economic Behavior 10 (1), 218-254.

Appendix

The score matrix

Strategies by cycles (score of the row player)	Cycle 1				Cycle 2				Cycle 3				Cycle 4				Cycle 5				Cycle 6			
	1234	**4123**	**3412**	**2341**	**4321**	**3214**	**2143**	**1432**	**1324**	**4213**	**3142**	**2431**	**4231**	**3124**	**2413**	**1342**	**1243**	**4132**	**3421**	**2314**	**2134**	**1423**	**4312**	**3241**
1234	2	1	2	3	2	2	2	2	2	1.5	2	2.5	2	1.5	2	2.5	2	1.5	2	2.5	2	1.5	2	2.5
4123	3	2	1	2	2	2	2	2	2.5	2	1.5	2	2.5	2	1.5	2	2.5	2	1.5	2	2.5	2	1.5	2
3412	2	3	2	1	2	2	2	2	2	2.5	2	1.5	2	2.5	2	1.5	2	2.5	2	1.5	2	2.5	2	1.5
2341	1	2	3	2	2	2	2	2	1.5	2	2.5	2	1.5	2	2.5	2	1.5	2	2.5	2	1.5	2	2.5	2
4321	2	2	2	2	2	1	2	3	2	1.5	2	2.5	2	1.5	2	2.5	2	2.5	2	1.5	2	2.5	2	1.5
3214	2	2	2	2	3	2	1	2	2.5	2	1.5	2	2.5	2	1.5	2	1.5	2	2.5	2	1.5	2	2.5	2
2143	2	2	2	2	2	3	2	1	2	2.5	2	1.5	2	2.5	2	1.5	2	1.5	2	2.5	2	1.5	2	2.5
1432	2	2	2	2	1	2	3	2	1.5	2	2.5	2	1.5	2	2.5	2	2.5	2	1.5	2	2.5	2	1.5	2
1324	2	1.5	2	2.5	2	1.5	2	2.5	2	1	2	3	2	2	2	2	1.5	2	2.5	2	2.5	2	1.5	2
4213	2.5	2	1.5	2	2.5	2	1.5	2	3	2	1	2	2	2	2	2	2	1.5	2	2.5	2	2.5	2	1.5
3142	2	2.5	2	1.5	2	2.5	2	1.5	2	3	2	1	2	2	2	2	2.5	2	1.5	2	1.5	2	2.5	2
2431	1.5	2	2.5	2	1.5	2	2.5	2	1	2	3	2	2	2	2	2	2	2.5	2	1.5	2	1.5	2	2.5
4231	2	1.5	2	2.5	2	1.5	2	2.5	2	2	2	2	2	1	2	3	2.5	2	1.5	2	1.5	2	2.5	2
3124	2.5	2	1.5	2	2.5	2	1.5	2	2	2	2	2	3	2	1	2	2	2.5	2	1.5	2	1.5	2	2.5
2413	2	2.5	2	1.5	2	2.5	2	1.5	2	2	2	2	2	3	2	1	1.5	2	2.5	2	2.5	2	1.5	2
1342	1.5	2	2.5	2	1.5	2	2.5	2	2	2	2	2	1	2	3	2	2	1.5	2	2.5	2	2.5	2	1.5
1243	2	1.5	2	2.5	2	2.5	2	1.5	2.5	2	1.5	2	1.5	2	2.5	2	2	1	2	3	2	2	2	2
4132	2.5	2	1.5	2	1.5	2	2.5	2	2	2.5	2	1.5	2	1.5	2	2.5	3	2	1	2	2	2	2	2
3421	2	2.5	2	1.5	2	1.5	2	2.5	1.5	2	2.5	2	2.5	2	1.5	2	2	3	2	1	2	2	2	2
2314	1.5	2	2.5	2	2.5	2	1.5	2	2	1.5	2	2.5	2	2.5	2	1.5	1	2	3	2	2	2	2	2
2134	2	1.5	2	2.5	2	2.5	2	1.5	1.5	2	2.5	2	2.5	2	1.5	2	2	2	2	2	2	1	2	3
1423	2.5	2	1.5	2	1.5	2	2.5	2	2	1.5	2	2.5	2	2.5	2	1.5	2	2	2	2	3	2	1	2
4312	2	2.5	2	1.5	2	1.5	2	2.5	2.5	2	1.5	2	1.5	2	2.5	2	2	2	2	2	2	3	2	1
3241	1.5	2	2.5	2	2.5	2	1.5	2	2	2.5	2	1.5	2	1.5	2	2.5	2	2	2	2	1	2	3	2

Proofs

Proposition 1: *A probability distribution P constitutes an equilibrium if and only if:*

$$\begin{pmatrix} \delta_7 \\ \delta_8 \\ \delta_9 \\ \delta_{10} \\ \delta_{11} \\ \delta_{12} \end{pmatrix} = \begin{pmatrix} -(\delta_1 + \delta_3 + \delta_5) \\ -(\delta_2 + \delta_4 + \delta_6) \\ 0.5 \cdot (-\delta_2 - \delta_4 - 2\delta_6 + \delta_3 - \delta_1) \\ 0.5 \cdot (\delta_1 + \delta_3 + 2\delta_5 + \delta_4 - \delta_2) \\ 0.5 \cdot (\delta_2 + \delta_4 + 2\delta_6 + \delta_3 - \delta_1) \\ 0.5 \cdot (-\delta_1 - \delta_3 - 2\delta_5 + \delta_4 - \delta_2) \end{pmatrix}$$

Proof: The following system of 12 linear equations confirms that in each of the 6 cycles, two adjacent strategies in the cycle, *S* and *D*(*S*), both score 0 points. This implies that each of the other two adjacent strategies earns 0 points as well. Therefore, the system characterizes the game's set of equilibria in the game.

$$\begin{cases} \delta_5 + \delta_7 + \delta_9 + \delta_{11} + 2\delta_1 = 0 & [L_1 \text{ scores } 0] \\ \delta_6 + \delta_8 + \delta_{10} + \delta_{12} + 2\delta_2 = 0 & [L_2 \text{ scores } 0] \\ \delta_5 + \delta_7 - \delta_9 - \delta_{11} + 2\delta_3 = 0 & [S_1(2) \text{ scores } 0] \\ \delta_6 + \delta_8 - \delta_{10} - \delta_{12} + 2\delta_4 = 0 & [S_2(2) \text{ scores } 0] \\ \delta_1 + \delta_3 - \delta_{10} + \delta_{12} + 2\delta_5 = 0 & [S_1(3) \text{ scores } 0] \\ \delta_2 + \delta_4 + \delta_9 - \delta_{11} + 2\delta_6 = 0 & [S_2(3) \text{ scores } 0] \\ \delta_1 + \delta_3 + \delta_{10} - \delta_{12} + 2\delta_7 = 0 & [S_1(4) \text{ scores } 0] \\ \delta_2 + \delta_4 - \delta_9 + \delta_{11} + 2\delta_8 = 0 & [S_2(4) \text{ scores } 0] \\ \delta_1 - \delta_3 + \delta_6 - \delta_8 + 2\delta_9 = 0 & [S_1(5) \text{ scores } 0] \\ \delta_2 - \delta_4 - \delta_5 + \delta_7 + 2\delta_{10} = 0 & [S_2(5) \text{ scores } 0] \\ \delta_1 - \delta_3 - \delta_6 + \delta_8 + 2\delta_{11} = 0 & [S_1(6) \text{ scores } 0] \\ \delta_2 - \delta_4 + \delta_5 - \delta_7 + 2\delta_{12} = 0 & [S_2(6) \text{ scores } 0] \end{cases}$$

The solution of this system is the 6-dimension space that appears in the proposition. ■

Proposition 2. *If 1>P(S)>0.5 for some S, then D(S) is the best response to P.*

Proof:

Assume without loss of generality that $1>P(L_0)>0.5$. We need to show that no strategy earns as high a score as L_1. It is sufficient to show that for any $X \neq L_1$, $score(X,L_0)=(3-t)$ implies that $score(X,Y)-score(L_1,Y)\leq t$ for any *Y*. In other words, *X* cannot compensate for its inferiority to L_1 against L_0 by its superiority when playing against some other strategies. The proof continues by considering all the possible strategies *X* and confirms that the condition on the payoffs is satisfied for all of them:

(I) The case of $X=L_3$ is straightforward: $score(L_3,L_0)=1$, and $score(L_3,Y)-score(L_1,Y)\leq 2$ since the lowest possible score is 1 point and the highest is 3 points.

(II) If $score(X,L_0)=2$, assume to the contrary that $score(X,Y)-score(L_1,Y)>1$. This implies that $score(L_1,Y)=1$ or 1.5 and thus *Y* can only be L_2 or S_2(i), for i=3,4,5,6. However, the only strategies that score 2.5 or 3 points against L_2 or S_2(i) are S_3(i) and L_3, which do not tie with L_0, a contradiction.

(III) In the case of $X=S_1$(i), for *i*=3,4,5,6, $score(X,L_0)=2.5$. Since S_1(i) is parallel to L_1, it scores at most half a point more than L_1 against $Y \in Cycle\ i$ or *Cycle 1*. $S_1(i)$ can score at most 2.5 points against $Y \notin cycle\ i\ or\ cycle\ 1$, while L_1 scores at least 1.5 points. $Score(L_1,Y)=1.5$ only if $Y=S_2(j)$ for *j*=3,4,5,6, and $score(S_1(i),S_2(j))<2.5$.

(IV) In the case of $X=S_3$(i), for *i*=3,4,5,6, $score(X,L_0)=1.5$. S_3(i) cannot score 2 points more than L_1 against some other strategy *Y*: $score(S_3(i),Y)=3$ only for $Y=S_2(i)$, and $score(L_1,S_2(i))=1.5$ and not 1. ■

Experimental results

Table 1 below presents the aggregate quantitative data for each of the experiments 1-3 in Study 1. In each experiment, the columns from left to right are: the number and then proportion of subjects who chose the strategy, and the average score of that strategy in the general tournament.

Table 1

Strategies & Notation		**Experiment 1 (N=113)**			**Experiment 2 (N=131)**			**Experiment 3 (N=641)**		
		#	%	**Score**	#	%	**Score**	#	%	**Score**
1234	L_0	21	18.6%	1.87	14	10.7%	1.87	141	22%	1.94
4123	L_1	11	9.7%	2.12	25	19.1%	2.1	65	10.1%	2.22
3412	L_2	5	4.4%	2.10	6	4.6%	2.23	21	3.3%	2.06
2341	L_3	4	3.5%	1.84	3	2.3%	2	23	3.6%	1.78
4321	$S_0(2)$	14	12.4%	1.94	10	7.6%	2.02	56	8.7%	1.96
3214	$S_1(2)$	2	1.8%	2.03	7	5.3%	2.06	23	3.6%	2.06
2143	$S_2(2)$	6	5.3%	2.03	8	6.1%	2.07	18	2.8%	2.04
1432	$S_3(2)$	0	0%	--	1	0.8%	2.04	17	2.7%	1.94
1324	$S_0(3)$	2	1.8%	1.90	2	1.5%	1.98	33	5.2%	1.94
4213	$S_1(3)$	7	6.2%	2.03	2	1.5%	2.06	25	3.9%	2.16
3142	$S_2(3)$	7	6.2%	2.06	4	3.1%	2.12	11	1.7%	2.06
2431	$S_3(3)$	0	0%	--	5	3.8%	2.04	10	1.6%	1.84
4231	$S_0(4)$	7	6.2%	1.90	3	2.3%	1.91	23	3.6%	1.95
3124	$S_1(4)$	3	2.7%	2.12	4	3.1%	2.1	18	2.8%	2.12
2413	$S_2(4)$	5	4.4%	2.06	3	2.3%	2.18	23	3.6%	2.05
1342	$S_3(4)$	1	0.9%	1.85	4	3.1%	2	14	2.2%	1.88
1243	$S_0(5)$	1	0.9%	1.90	0	0%	----	15	2.3%	2.01
4132	$S_1(5)$	5	4.4%	2.03	5	3.8%	2.04	13	2%	2.08
3421	$S_2(5)$	2	1.8%	2.06	3	2.3%	2.14	8	1.3%	1.99
2314	$S_3(5)$	2	1.8%	1.93	2	1.5%	2.06	27	4.2%	1.92
2134	$S_0(6)$	4	3.5%	1.99	7	5.3%	1.99	19	3%	1.96
1423	$S_1(6)$	2	1.8%	2.03	4	3.1%	2.1	11	1.7%	2.07
4312	$S_2(6)$	0	0%	--	5	3.8%	2.11	10	1.6%	2.04
3241	$S_3(6)$	2	1.8%	1.93	4	3.1%	2	17	2.7%	1.93

Table 2 below presents the distribution following the normalization discussed on page 26 for each of the experiments 1-3.

Table 2

Strategies & Notation		**Normalization for equilibrium analysis**		
		Experiment 1 (n=57)	**Experiment 2 (n=51)**	**Experiment 3 (n=283)**
1234	L_0	28.1%	15.7%	42.4%
4123	L_1	12.3%	43.1%	14.8%
3412	L_2	0%	0%	0%
2341	L_3	0%	0%	0%
4321	$S_0(2)$	14%	3.9%	13.4%
3214	$S_1(2)$	3.5%	11.8%	2.1%
2143	$S_2(2)$	0%	0%	0%
1432	$S_3(2)$	0%	0%	0%
1324	$S_0(3)$	0%	0%	7.8%
4213	$S_1(3)$	12.3%	0%	5.3%
3142	$S_2(3)$	8.8%	3.9%	0%
2431	$S_3(3)$	0%	5.9%	0%
4231	$S_0(4)$	3.5%	0%	0%
3124	$S_1(4)$	3.5%	0%	1.4%
2413	$S_2(4)$	0%	0%	0%
1342	$S_3(4)$	0%	0%	0%
1243	$S_0(5)$	0%	0%	2.5%
4132	$S_1(5)$	5.3%	5.9%	0%
3421	$S_2(5)$	1.8%	5.9%	0%
2314	$S_3(5)$	0%	0%	5%
2134	$S_0(6)$	7%	3.9%	3.2%
1423	$S_1(6)$	0%	0%	0%
4312	$S_2(6)$	0%	0%	0%
3241	$S_3(6)$	0%	0%	2.1%

Table 3 presents the data from Experiments A and B (**Version 2**, with and without monetary incentives, respectively).

Table 3

Strategies	**Notation**	**Experiment A (N=148)**		**Experiment B (N=704)**	
		#	%	#	%
1234	L_0	15	10.1%	131	18.6%
4123	L_1	27	18.2%	91	12.9%
3412	L_2	10	6.8%	37	5.3%
2341	L_3	3	2.0%	20	2.8%
4321	$\mathrm{S}_0(2)$	6	4.1%	41	5.8%
3214	$S_1(2)$	16	10.8%	38	5.4%
2143	$S_2(2)$	4	2.7%	17	2.4%
1432	$S_3(2)$	2	1.4%	12	1.7%
1324	$S_0(3)$	2	1.4%	24	3.4%
4213	$S_1(3)$	10	6.8%	33	4.7%
3142	$S_2(3)$	3	2.0%	15	2.1%
2431	$S_3(3)$	4	2.7%	11	1.6%
4231	$S_0(4)$	5	3.4%	30	4.3%
3124	$S_1(4)$	8	5.4%	17	2.4%
2413	$S_2(4)$	12	8.1%	22	3.1%
1342	$S_3(4)$	3	2.0%	8	1.1%
1243	$S_0(5)$	1	0.7%	18	2.6%
4132	$S_1(5)$	1	0.7%	17	2.4%
3421	$S_2(5)$	0	0%	8	1.1%
2314	$S_3(5)$	4	2.7%	33	4.7%
2134	$S_0(6)$	5	3.4%	28	4%
1423	$S_1(6)$	0	0%	21	3%
4312	$S_2(6)$	4	2.7%	11	1.6%
3241	$S_3(6)$	3	2.0%	21	3%

<u>The form in Experiment 1 (translated from Hebrew)</u>

You are a tennis team coach, planning to send your team to a tournament. Each of your classmates is a coach of a team participating in this tournament.

Each team has four players: one of level **A+** (the highest level), one of level **A**, one of level **B+**, and one of level **B** (the lowest level).

The coach's task is to assign his players to **"position 1"**, **"position 2"**, **"position 3"** and **"position 4"** (one player in each position).

Each team will play against each of the other teams in the tournament. A game between two teams includes four matches: a tennis player that was assigned by his coach to "position X" will play once against the player in "position X" of the other team. You don't know how the other coaches assign their players.

In any match between two tennis players of different levels, the one with the higher level wins. When two players with the same level play, the outcome is a tie.

A winner in a match brings his team **1 point**, and a player who ends the match with a tie brings his team **½ a point**. A loss yields **0 points**.

The team's score at the end of the tournament is the number of points it gained in the games against all other teams (of the other participants in the experiment). The winning team is the one with the highest score.

The coach whose team gained the highest score will win 200 NIS. (In case of several winning teams, a lottery will determine who wins the prize.)

How will you allocate the players in your team? (one player in each position)

	Position 1	**Position 2**	**Position 3**	**Position 4**
A+				
A				
B+				
B				

Version 1 and 2 in Experiments 2 and A, respectively (translated from Hebrew)

You are a tennis team coach, planning to send your team to a tournament. Each team in the tournament has four players: one of level **A+** (the highest level), one of level **A**, one of level **B+**, and one of level **B** (the lowest level).

The coach's task is to assign his players to **"position 1"**, **"position 2"**, **"position 3"** and **"position 4"** (one player in each position).

Each team will play against each of the other teams in the tournament. A game between two teams includes four matches: a player that was assigned by his coach to "position X" will play once against the player in "position X" of the other team. You don't know how the other coaches assign their players.

In any match between two players of different levels, the one with the higher level wins. When two players with the same level play, the outcome is a tie.

[In version 1 - A winner in a match brings his team **1 point**, and a player who ends the match with a tie brings his team **½ a point**. A loss yields **0 points**.]

[In version 2 - **At the end of any game between two teams, a team gets 1 point only if it won three matches out of four**. In such a case, the other team gets 0 points. In case of any other result, none of the teams gets points.]

The team's score at the end of the tournament is the number of points it gained in all the games against other teams.

The winning team is the one with the highest score, and the prize for the winner is 200 NIS. [In case of several winners, one of them will be selected randomly to receive the prize].

How will you allocate your players? **

	Position 1	**Position 2**	**Position 3**	**Position 4**
A+				
A				
B+				
B				

** Note that other students in your class play the role of other coaches in the tournament, so your total score in this game will be your team's total score, after playing against each of the other students' teams.

<u>Version 1 and 2 as they appear in Experiments 3 and B, respectively</u>

Games and Behavior - The Problems

The Problem

You are a tennis team coach, planning to send your team to a tournament. Each team in the tournament has four players: one of level **A+** (the highest level), one of level **A**, one of level **B+**, and one of level **B** (the lowest level).

The coach's task is to assign his players to "**position 1**", "**position 2**", "**position 3**" and "**position 4**" (one player in each position).
Each team will play against each of the other teams in the tournament. A game between two teams includes four matches: a player that was assigned by his coach to "position X" will play once against the player in "position X" of the other team. You don't know how the other coaches assign their players.

In any match between two players of different levels, the one with the higher level wins. When two players with the same level play, the outcome is a tie.

[<u>In version 1</u> - A winner in a match brings his team **1 point**, and a player who ends the match with a tie brings his team **½ a point**. A loss yields **0 points**.]

[<u>In version 2</u> - **At the end of any game between two teams, a team gets 1 point only if it won three matches out of four**. In such a case, the other team gets 0 points. In case of any other result, none of the teams gets points.]

The team's score at the end of the tournament is the number of points it gained in all the games against other teams.
The winning team is the one with the highest score, and the prize is $10,000. [In case of several winning teams, the prize is divided between them.]
The only goal of players and coaches (including you) is to have their team getting the highest score among the teams.

How will you allocate your players in order to achieve this goal?**

** Note that other students in your class play the role of other coaches in the tournament, so your total score in this game will be your team's total score, after playing against each of the other students' teams.

	pos 1	pos 2	pos 3	pos 4
A+	☐	☐	☐	☐
A	☐	☐	☐	☐
B+	☐	☐	☐	☐
B	☐	☐	☐	☐