Coordinating with a "Problem Solver"

Jacob Glazer
Tel Aviv University and
University of Warwick

and

Ariel Rubinstein
Tel Aviv University and
New York University

Abstract

We introduce a new type of agent called a "problem solver" (PS). The PS interacts with conventional players in a modified coordination game. The PS wishes to coordinate with at least one other player whereas the other players would like to coordinate with the PS but not with the other players. The PS receives partial information about the other players' moves . Unlike a regular player, the PS does not "put himself in the shoes" of other players. Rather, he solves the puzzle of finding all possible configurations of moves consistent with what he observes and takes an action that is optimal given a uniform distribution over those configurations. Equilibria are characterized. Some of them have the feature that the PS always coordinates with some players but may be uncertain that he will successfully do so.

The second author acknowledges financial support from ERC grant 269143.
Our thanks to Noga Alon, Ayala Arad, Kfir Eliaz, Gil Kalai and Rani Spiegler for helpful conversations.

1. **Introduction**

According to the standard game theoretic approach, a player "puts himself in the shoes" of the other players, forms beliefs about their behavior and responds optimally to these beliefs. However, a player's deliberations often do not involve forming beliefs about the strategic considerations of the other players, even if he has invested significant cognitive resources.

Imagine that you are solving a crossword puzzle. It is unlikely that you are even thinking about its author.

Or, imagine the following scenario: There are two cards facedown in front of you. On each of them is the name of a different country. Your task is to guess one of the two names. You are told that both names have four unique letters and that they have three letters in common. You will eventually come to the conclusion that the only possible configuration is {Iran, Iraq} and thus each of these two names will be a successful guess. However, what would you guess if you are told that the two names share at least two letters. In that case, there are four possible pairs that fit this description: {Chad,Cuba}, {Iran, Iraq}, {Iran, Mali}, {Iraq,Mali}. If you aren't thinking strategically about the motives of the person who chose the two countries, then your guess will probably be Iran, Iraq or Mali, but not Chad or Cuba since each of the first three names appears in two of the possible configurations whereas the other two appear in only one.

One final example: You are involved in a coordination game with another person. The two of you would like to coordinate on a location A,B or C. The other person chooses his location first; you have to choose a location after observing only a signal, A,B or C, about his choice. The signal is somewhat noisy but there is a high probability that it is accurate. Suppose that you receive the signal C. It is likely that you will choose the location C; however, according to the standard approach, ignoring the signal C and choosing A is consistent with equilibrium.

These three examples demonstrate that players sometimes do not behave strategically but rather deliberate on their decision as if they were solving a puzzle. Accordingly, we introduce a new type of economic agent called a problem solver (PS). The PS interacts with conventional players but does not observe their moves, about which he receives only partial information. We assume that the PS calculates the set of possible configurations of the other players' moves that are consistent with the information he receives and chooses a best response to the uniform distribution over that set. For the PS, finding all the configurations of moves that are consistent with what he observes is like solving a puzzle.

The platform we use to demonstrate the idea is a new version of the coordination game. The players in the game are labelled $0, 1, .., n$. Each player chooses an alternative from a set X. Player 0 is interested in coordinating his choice with at least one of the other players. Each of the players $1, .., n$ would like to coordinate with player 0 but not with any other player. Players $1, .., n$ make their choices simultaneously. Player 0 receives some information about their choices and then make his own. The game has trivial equilibria, some of which do not make much sense. In fact, regardless of the information player 0 receives, every profile of n distinct choices, one of which is chosen by player 0 with probability 1, is consistent with a pure sequential equilibrium. According to this approach, player 0 thinks strategically and the success of the coordination is due to his knowledge of the equilibrium. Indeed, the information player 0 receives regarding the other n players is superfluous. We suggest a more realistic approach according to which player 0 is a problem solver. He identifies all the profiles of the players' choices that are consistent with the data he receives and treats them as equally likely. He then chooses an action in order to maximize the chance that his choice will match that of at least one of the other players.

We analyze the equilibria of the modified game assuming that the set X is a large matrix and player 0 (the PS) observes only the number of players located in each row and in each column of the matrix. The equilibrium of the model in the presence of a problem solver differs significantly from that of the above coordination game in which player 0 behaves like a regular player. We show that in all of the game's equilibria, the PS coordinates with one of the other players with certainty. However, an interesting phenomenon emerges: There exist equilibria in which the PS chooses a position in the matrix which he believes might be vacant even though it is occupied with certainty. Such a phenomenon would not arise in the case that player 0 is just a regular player.

2. **The model**

There are $n + 1$ agents. We refer to agent 0 as a PS and to agents $1, .., n$ as players.

Each agent chooses a position, which is a pair of characteristics in the set $X = A \times B$ (where A and B are disjoint sets). For simplicity, assume that both A and B contain at least n elements. We will often refer to the elements of A as columns and to the elements of B as rows. A product set of columns and rows is called a *box*.

The players make their choices simultaneously. The PS does not observe the players' choices but only the number of occupied entries in each column and row. With that information he chooses his position.

The PS gets utility 1 if his choice coincides with the choice of one of the players, and 0 otherwise. Each of the players $1,..,n$ gets utility of 1 if his choice coincides only with the choice of the PS; he gets utility of -1 if his choice coincides with at least one of the other players and 0 otherwise.

Given these preferences we can confine ourselves to strategy profiles where no two players occupy the same position. An outcome of the n players' choices is a *matrix* $M = (M_{a,b})$, where $M_{a,b} \in \{0,1\}$, with n 1's. The notation $M_{a,b} = 1$ signifies that the position (a,b) is occupied and $M_{a,b} = 0$ signifies that it is vacant. The PS observes only the *data vector* $d(M) = (d(M)(x))_{x \in A \cup B}$ where $d(M)(x)$ is the number of players occupying entries with a characteristic x in the matrix M (i.e., $d(M)(a) = \Sigma_l M_{al}$ and $d(M)(b) = \Sigma_k M_{kb}$). A vector d is consistent if there is a matrix M such that $d = d(M)$. By definition, the PS observes only consistent data vectors. We refer to $d(a) + d(b)$ as the *score* of the entry (a,b).

The PS assumes that all matrixes consistent with the data he observes are equally likely and randomly picks one entry from the set of entries with the highest probability of being occupied. Formally, a matrix M is said to be d-consistent if $d(M) = d$. Let $\lambda(d,x)$ be the proportion of matrixes consistent with d , in which x is occupied. If $\lambda(d,x) = 1$, we say that x is *revealed to be occupied* by d. If $\lambda(d,x) = 0$, we say that x is *revealed to be vacant* by d (hereafter, we use the term "revealed" to mean "revealed to be occupied"). We say that the matrix M is *revealed* if it is the only matrix consistent with $d(M)$. Denote by $C(M) = \{x \mid \lambda(d(M),x)$ is maximal$\}$ the set of all entries with the highest probability of being occupied, given $d(M)$. The probability that the PS picks x is $\mu(M,x) = 1/|C(M)|$ for each $x \in C(M)$ and $\mu(M,x) = 0$ for each $x \notin C(M)$. In other words, $C(M)$ is the PS's choice set given that he observes $d(M)$ (and thus depends on $d(M)$ only).

An *equilibrium* is a matrix M such that no player can increase his probability of being picked by moving to a vacant entry. Formally, let $M(x \to y)$ be the matrix derived from M after switching the values of entries x and y. The matrix M is an *equilibrium* if, for each occupied entry x, $\mu(M,x) \geq \mu((M(x \to y),y)$ for any entry y that is not occupied in M.

Example 1: Consider the following matrixes with $n = 5$ (when describing a matrix, vacant rows and columns are not depicted).

$$M_1 = \begin{array}{|c|c|c|} \hline 1^* & 1 & 1 \\ \hline 1 & 0^{**} & 0 \\ \hline 1 & 0 & 0 \\ \hline \end{array} \quad M_2 = \begin{array}{|c|c|c|} \hline 1^* & 1 & 1 \\ \hline 1 & 1 & 0 \\ \hline 0 & 0^{**} & 0 \\ \hline \end{array} \quad M_3 = \begin{array}{|c|c|c|c|c|} \hline 1^* & 1 & 1 & 1 & 1 \\ \hline * * & * * * & & & \\ \hline \end{array}$$

The matrix M_1 is revealed. Thus, the probability of each occupied entry being picked is $1/5$. However, M_1 is not an equilibrium. If the occupier of $*$ moves to $* *$, the new data vector will be: $((2,2,1),(2,2,1))$, which is consistent with 5 matrixes, such that each of the four entries (a,b) for which $(d(a),d(b)) = (2,2)$ is occupied in 4 of the 5 matrixes. Each of the 4 entries (a,b) for which $(d(a),d(b))$ is $(2,1)$ or $(1,2)$ is occupied in 2 of the 5 matrices. The unique entry (a,b), for which $(d(a),d(b)) = (1,1)$, is occupied in only one of the d-consistent matrixes. Thus, by moving from $*$ to $* *$ the mover increases his probability of being picked from $1/5$ to $1/4$.

The matrix M_2 is revealed but is not an equilibrium since the player who occupies $*$ can increase his probability of being picked from $1/5$ to $1/3$ by moving to $* *$.

The matrix M_3 is revealed and is an equilibrium. If the player at $*$ moves to an entry such as $* *$ (which does not share any characteristics with the other four occupied entries), then the new data vector will be consistent with 5 matrixes and in only one of them is $* *$ occupied. Each of the other occupied entries is occupied in 4 of the 5 matrices consistent with the new data and thus the mover reduces his probability of being picked from $1/5$ to 0. If the occupier of $*$ moves to an entry such as $* * *$ (which shares one characteristic with one other occupied entry), then the new matrix will also be revealed and the player gains nothing by moving.

3. The Problem Solver's Behavior

In this section, we present some properties of the set of matrixes which are consistent with a given data set. These properties determine the Problem Solver's "response function". In particular, we will show that either:

(1) The set of entries that are revealed to be occupied forms a step set and the PS randomly picks one of these entries (every other entry has a strictly lower score than at least one revealed entry);

or

(2) No entry is revealed to be occupied and the PS randomly picks one of the entries with the maximal score.

Claim 1: Let d be a consistent data vector such that $d(1) > d(2)$, where 1 and 2 are elements of A. Then, for any $b \in B$, $\lambda(d,(1,b)) \geq \lambda(d,(2,b))$. Furthermore, if there is a d-consistent matrix M such that $M_{1,b} = 1$ and $M_{2,b} = 0$, then the inequality is strict.

Proof: Let $b \in B$. Fix the values for all entries other than those in columns 1 and 2. Partition the class of all d-consistent matrixes with these fixed values outside columns 1 and 2 into (up to) four cells, denoted by $M(\alpha,\beta)$, $\alpha \in \{0,1\}$, $\beta \in \{0,1\}$, such that $M(\alpha,\beta)$ is the cell in this partition that consists of the matrixes for which $M_{1,b} = \alpha$ and $M_{2,b} = \beta$. We will show that for each class $|M(1,0)| \geq |M(0,1)|$ and if $M(1,0)$ is not empty, then the inequality is strict. This is sufficient to prove the claim since If there is a d-consistent matrix M such that $M_{1,b} = 1$ and $M_{2,b} = 0$, then for at least one set of entries in columns $B - \{1,2\}$ we have $M(1,0) \neq \emptyset$.

We first show that if $M(1,0) = \emptyset$, then $M(0,1) = \emptyset$. If $M(0,1)$ is not empty, then there is a d –consistent matrix M with $M_{1,b} = 0$ and $M_{2,b} = 1$. By $d(1) > d(2)$, there is a row b' where $M_{1,b'} = 1$ and $M_{2,b'} = 0$. Switching all values in $\{1,2\} \times \{b,b'\}$, we get another d-consistent matrix which is in $M(1,0)$. Therefore, if $M(1,0)$ is empty, then the number of matrixes in this class in which $M_{1,b} = 1$ is the same as the number of matrixes in the class in which $M_{2,b} = 1$.

We next show that if $M(1,0) \neq \emptyset$, then the number of elements in $M(0,1)$ is strictly smaller than in $M(1,0)$. Define $L_{1,1}$ to be the set of rows in which the data regarding the rows implies that the missing values in columns 1 and 2 are $(1,1)$, and define $L_{0,0}$ in a similar manner. For the rows in $B - L_{11} - L_{00}$, the data dictates that the missing values in columns 1 and 2 be either $(0,1)$ or $(1,0)$. It must be that in any $\delta(1) = d(1) - |L_{1,1}|$ of these rows the values in the two columns are $(1,0)$ and in the other $\delta(2) = d(2) - |L_{1,1}|$

rows the values must be $(0,1)$. Thus,
$|M(1,0)| = C(\delta(1)-1, \delta(1)+\delta(2)-1) > |M(0,1)| = C(\delta(1), \delta(1)+\delta(2)-1)$ where the strict inequality follows from the fact that $\delta(1) > \delta(2)$. ■

Claim 2: Let d be a consistent data set. Assume that a^* and b^* maximize d over A and B, respectively. If, for some a, $d(a^*) > d(a)$ and $1 > \lambda(d,(a^*,b^*)) > 0$, then $\lambda(d,(a^*,b^*)) > \lambda(d,(a,b^*))$.

Proof: By claim 1, it is sufficient to show that there exists a matrix M consistent with d, such that $M_{a^*,b^*} = 1$ and $M_{a,b^*} = 0$.

Since $1 > \lambda(d,(a^*,b^*))$, there exists a d-consistent matrix M with $M_{a^*,b^*} = 0$.

If $M_{a,b^*} = 1$, then since $d(a^*) > d(a)$, there must be some $b \in B$ with $M_{a^*,b} = 1$ and $M_{a,b} = 0$. By switching the values in the four entries $\{a^*, a\} \times \{b^*, b\}$, we obtain a d-consistent M' with $M'_{a^*,b^*} = 1$ and $M'_{a,b^*} = 0$.

b*	0		1		
b	1		0		
	a*		**a**		

If $M_{a,b^*} = 0$, then by $d(a^*) > d(a)$ there exists a row b with $M_{a^*,b} = 1$ and $M_{a,b} = 0$. However, since b^* maximizes d over B, there also exists an a' such that $M_{a',b^*} = 1$ and $M_{a',b} = 0$. By switching the values in $\{a^*, a'\} \times \{b^*, b\}$, we obtain a d-consistent matrix M' with $M'_{a^*,b^*} = 1$ and $M'_{a,b^*} = 0$.

b*	0		0		1
b	1		0		0
	a*		**a**		**a′**

■

Claim 3: Let M^* be a matrix such that the box $C \times R$ is occupied and the "dual" box $C^c \times R^c$ is vacant. Then, each element in $C \times R$ is revealed to be occupied and each element in the dual box is revealed to be vacant.

R	1	1	1	?	?
R	1	1	1	?	?
	?	?	?	0	0
	?	?	?	0	0
	?	?	?	0	0
	C	**C**	**C**		

Proof: Given a data vector d and a set E of rows or columns, let $d(E) = \sum_{e \in E} d(e)$. It must be that $d(M^*)(C^c) + d(M^*)(R^c) + |C| \times |R| = n$. On the other hand, for every matrix M such that $M_{a,b} = 0$ for some $(a,b) \in C \times R$, it must be that $d(M)(C^c) + d(M)(R^c)$ $+|C| \times |R| > n$ and thus M is inconsistent with $d(M^*)$. An analogous argument can be used to show that the positions in the dual box are revealed to be vacant. ■

From Claim 1, it follows that the set of revealed entries is a "step set". That is, there are two sequences of sets $B(1) \supset B(2) \supset \ldots \supset B(I)$ and $A(1) \subset A(2) \subset \ldots \subset A(I)$ such that the set of revealed (to be occupied) entries is $\cup(A(i) \times B(i))$ (in the figure below it is the set $A(1) \times B(1) = \{a_1, a_2\} \times \{b_1, b_2, b_3, b_4, b_5\}$, $A(2) \times B(2) = \{a_1, a_2, a_3\} \times \{b_1, b_2, b_3, b_4\}$, etc.) and the d-value of any element in $A(i)$ ($B(i)$) is larger than the d-value of any element in $A(j)$ ($B(j)$) where $j > i$.

$\mathbf{b}_1$	Y	Y	Y	Y	Y	Y	X
$\mathbf{b}_2$	Y	Y	Y	Y	Y	X	Z
$\mathbf{b}_3$	Y	Y	Y	Y	Y	X	Z
$\mathbf{b}_4$	Y	Y	Y	X	X	Z	Z
$\mathbf{b}_5$	Y	Y	X	Z	Z	Z	Z
$\mathbf{b}_6$	X	X	Z	Z	Z	Z	Z
$\mathbf{b}_7$	X	X	Z	Z	Z	Z	Z
	$\mathbf{a}_1$	$\mathbf{a}_2$	$\mathbf{a}_3$	$\mathbf{a}_4$	$\mathbf{a}_5$	$\mathbf{a}_6$	$\mathbf{a}_7$

In what follows, we denote the step set of revealed entries by $Y = \cup A(i) \times B(i)$, the dual step set by $Z = \cup A(i)^c \times B(i)^c$ and the union of the boxes between Y and Z by $X = \cup(A_i - A_{i-1}) \times (B_{i-1} - B_i)$.

Claim 4: Let d be a consistent data vector. Assume that the set of revealed elements $Y = \cup A(i) \times B(i)$ is not empty. Then, any element in the dual set $Z = \cup A(i)^c \times B(i)^c$ is revealed to be vacant.

Proof: Given a consistent data vector d, the Gale-Ryser algorithm (see Ryser (1963) and Kraus (1996)) ends with a d-consistent matrix. The algorithm is sequential and starts with a certain initial matrix. In each step of the algorithm, a permissible pair of entries that are positioned in the same row – one occupied (the one with the higher d) and the other vacant – is selected and their values are swapped. An important property of the algorithm is that any choice of a chain of permissible pairs leads to a d-consistent matrix.

In order to describe the algorithm precisely, order the elements in the set A according to their d-values $a_1, \ldots, a_{|A|}$ and order the elements in B according to their d-values $b_1, \ldots, b_{|B|}$. The algorithm starts with a matrix M_0 in which, for any row b, 1's are assigned to the $d(b)$ entries in this row at the columns $a_1, \ldots, a_{d(b)}$. For M_0, the number of 1's in column a_k is $z(k) = |\{b \mid \text{the number } d(b) \geq k\}|$. Obviously, (a) $\sum_{i=1,..,k} z(i) \geq \sum_{i=1,..,k} d(a_i)$ for all k and (b) if $\sum_{i=1,..,k} z(i) = \sum_{i=1,..,k} d(a_i)$, then in any d-consistent matrix all entries in $\{a_1,..,a_k\} \times \{b_1,..,b_{z(k)}\}$ are occupied and all entries in $\{a_{k+1}, \ldots, a_{|A|}\} \times \{b_{z(k)+1}, \ldots, b_{|B|}\}$ are vacant. In each step of the algorithm, a "1" in the lowest index column for which the number of $1'$s is strictly greater than $d(a)$ is moved to the first column a' in which the number of $1'$s is strictly less than $d(a')$.

In his proof that the algorithm ends with a d-consistent matrix, Kraus (1996) shows that the algorithm works by starting with any matrix having the following two properties: (i) The number of 1's in each $b \in B$ is $d(b)$ and (ii) for each k, the sum of the 1's in the first k columns is at least as large as $\sum_{i=1,..,k} d(a_i)$.

To prove Claim 4, it is sufficient to show that, for all l, $\sum_{i=1,..,|A(l)|} z(i) = \sum_{i=1,..,|A(l)|} d(a_i)$. Assume not. Then, for some l we have $\sum_{i=1,..,|A(l)|} z(i) > \sum_{i=1,..,|A(l)|} d(a_i)$. Let k^* be the lowest $k > |A(l)|$ for which $\sum_{i=1,..,k} z(i) = \sum_{i=1,..,k} d(a_i)$. It must be that $d(b_{|B(l)|}) < k^*$; otherwise, by (b) above, the set $A(l)$ would be larger. Now start the Gale -Ryser algorithm from a matrix that modifies M_0 by moving a single "1" from the entry $(a_{|A(l)|}, b_{|B(l)|})$ to the entry in the same row in column k^*. The matrix satisfies properties (i) and (ii) above and the algorithm leads to a d-consistent matrix in which one of the first $|A(l)|$ entries in row $b_{|B(l)|}$ is 0, violating the assumption that this entry is revealed to be occupied.

4. Equilibrium

In this section, we show that in all equilibria the PS picks an occupied entry with probability 1 (Proposition 1). However, in some equilibria, the PS assigns a strictly positive probability to the possibility that the occupied position he is picking is vacant. We also classify the structure of all equilibria (Proposition 2).

Proposition 1: In equilibrium the PS picks an occupied entry with probability 1.

Proof: Let M be an equilibrium. If there are revealed positions, then the PS obviously chooses one of them. If no entry is revealed, then, by Claim 2, $C(M)$ is the box of all entries that maximize the $d(M)$ score. It is left to show that all entries in $C(M)$ are occupied.

If an entry (a,b) outside $C(M)$ is occupied and $(a^*,b^*) \in C(M)$ is vacant, then the move from (a,b) to (a^*,b^*) is beneficial: the score of (a^*,b^*) in $M((a,b) \to (a^*,b^*))$ increases by at least 1 relative to the score in M and the score of any other entry increases by at most 1. Thus, (a^*,b^*) maximizes the score after the move, and hence it will be picked with a positive probability.

If all the occupied entries are in $C(M)$ and $C(M)$ contains a vacant entry, then any vertical move from an occupied entry (a^*,b) to a vacant entry (a^*,b^*) is beneficial since (i) (a^*,b^*) maximizes the score after the move, (ii) $C(M((a^*,b) \to (a^*,b^*)))$ $\subseteq C(M)$ and (iii) the entry (a^*,b) is excluded from $C(M((a^*,b) \to (a^*,b^*)))$. ■

Proposition 2: If M is an equilibrium, it must have one of the following three structures:

(1) The matrix M is revealed and forms a "step set". The PS picks one of the occupied positions.

(2) There is a "step set" of entries that are revealed as occupied and the PS picks one of them. The "dual step set" is revealed to be vacant. Each box lying between these two sets contains at least three occupied entries that are not picked by the PS.

(3) None of the positions are revealed to be occupied. All positions with maximal score are occupied and the PS picks one of them.

Proof: The fact that the set of revealed entries forms a step set follows from Claim 1. By Claim 4, the dual step set is revealed to be vacant.

Assume that M is an equilibrium in which some, but not all, occupied entries are revealed. We will show that any box $A' \times B'$ in area X (i.e., the area consisting of all the positions between Y, the set of entries that are revealed as occupied, and Z, the set of entries that are revealed to be vacant) contains at least three occupied entries that do not share any characteristic (namely, they are positioned in three different rows and three different columns). If $A' \times B'$ is entirely vacant, then a move from an unrevealed occupied entry into this empty box will be beneficial to the mover since it will reveal this entry to be occupied (by Claim 3). It is impossible that all occupied entries in this box lie in the same row or the same column since (again, by Claim 3) they would then be revealed. It is also impossible that in equilibrium all occupied entries in this box lie in exactly two rows (or two columns) since if (a_1, b_1) and (a_2, b_2) are occupied then a move from (a_1, b_1) to (a_1, b_2) is beneficial since the deviator will be revealed (once again, by Claim 3).

Finally, if none of the positions are revealed as occupied, then, by claim 2, $C(M)$ contains all of the positions with the highest score and, by the proof of Proposition 1, all of the entries in $C(M)$ must be occupied. ■

An interesting feature of the model is that although in equilibrium the PS always picks an occupied position, there are equilibria in which he may not know that the position is occupied. Such a phenomenon could not occur under the conventional equilibrium assumption but may emerge in equilibria of the third structure. The matrix M_3 demonstrates such an equilibrium for $n = 10$. The example can be extended to any $n > 10$ by "extending" the "arms of the L". In M_3, the PS chooses the top-left position and assigns a probability of about 85% to this entry being occupied (since it is occupied in 2400 of the 2850 matrixes that are consistent with the data that the PS observes).

$$M_3 = \begin{array}{|c|c|c|c|c|c|c|} \hline 1 & 1 & 1 & 1 & 0 & 0 & 0 \\ \hline 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ \hline 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ \hline 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ \hline 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ \hline 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ \hline 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ \hline \end{array}$$

To summarize:

Proposition 3: For $n \geq 10$ there exists an equilibrium in which the PS picks an occupied entry although the probability he assigns to it being occupied is strictly less than one.

5. Equilibrium when players can share entries

In this section we study a variant of the model in which the PS' utility is as before but a player's utility is modified to be $1/k$ if he shares the PS's chosen cell with $k-1$ other players, and 0 otherwise. The rationale for this utility function is that once the PS has chosen an occupied cell he randomly picks one of the occupiers and a player is interested in increasing the probability of him being picked. A matrix is an assignment of non-negative integers (not necessarily zeros or ones) to all entries, such that the sum of the numbers in all entries is n. For each row (column), the PS receives information about the total number of *players* occupying this row *(column)*. The PS identifies the matrixes that are consistent with the data. He then randomly picks, an entry with the largest number of consistent matrixes in which this entry is occupied (by at least one player). For example, the data $d(a_1) = d(b_1) = 3$ and $d(a_2) = d(b_2) = 1$ is consistent with two matrixes:

$\begin{array}{|c|c|} \hline 2 & 1 \\ \hline 1 & 0 \\ \hline \end{array}$ and $\begin{array}{|c|c|} \hline 3 & 0 \\ \hline 0 & 1 \\ \hline \end{array}$.

Only the entry (a_1, b_1) is occupied in both matrices and therefore it is picked by the PS.

An equilibrium is a matrix such that a player who occupies one entry cannot increase his utility by moving to another. Formally, there is no occupied entry x and entry y such that $\frac{\mu((M(x \to y),y)}{M(y)+1} > \frac{\mu(M,x)}{M(x)}$ where $M(x \to y)$ is the matrix where x is reduced by one and y is increased by 1.

Obviously, any matrix in which the n players occupy n entries in one row or in one column is an equilibrium (since the move of a player to another occupied entry reduces his utility from $1/n$ to $1/(2n-2)$).

Proposition 4: In every equilibrium, the n players occupy the same row or the same column. If $n > 2$ each player occupies a different entry.

Proof: Let M be an equilibrium and let $d = d(M)$. If all players are located in the same row (or column), then it must be that each occupies an exclusive entry since otherwise he could increase his utility by moving to a vacant entry in the same row.

Assume, for the purpose of contradiction, that M is an equilibrium in which the occupied entries are located in at least two rows and at least two columns.

Step 1: There are entries in M that are not revealed.

Proof: There must be two entries (a_1, b_1) and (a_2, b_2) such that $a_1 \neq a_2$ and $b_1 \neq b_2$ and $M_{a_1,b_1} \geq M_{a_2,b_2} > 0$. The matrix obtained by starting from M and subtracting the number M_{a_2,b_2} from the entries (a_1, b_1) and (a_2, b_2) and adding M_{a_2,b_2} to the entries (a_1, b_2) and (a_2, b_1) is also consistent with d, although (a_2, b_2) is not occupied. Thus, (a_2, b_2) is not revealed.

Step 2: Let M be a matrix such that $M_{a,b} = 0$, $d(M)(a) > 0$ and $d(M)(b) > 0$. Then there is a matrix M', consistent with $d(M)$, for which $M'_{a,b} = 1$.

Proof: There must be a' and b' such that $M_{a',b} > 0$ and $M_{a,b'} > 0$. The matrix M' obtained by subtracting 1 from $M_{a',b}$ and $M_{a,b'} > 0$ and adding 1 to $M_{a,b}$ and $M_{a',b'}$ is consistent with d. Thus (a, b) is not revealed as vacant.

Step 3: Assume that $d(a_1) > d(a_2)$. Let b_1 be a row. If (a_2, b) is revealed to be occupied then so is (a_1, b).

Proof: Assume that there is a matrix M' consistent with $d(M)$ such that $M'_{a_1,b} = 0$. Since $d(a_1) > d(a_2)$, there is a set of rows B such that $\sum_{b \in B} M'_{a_1,b} > M'_{a_2,b_1} = m > 0$. Construct another matrix M'' where $M''_{a_2,b_1} = 0$, $M''_{a_1,b_1} = m$, $M''_{a_1,b} = M'_{a_1,b} - m_b$, $M''_{a_2,b} = M'_{a_2,b} + m_b$ and $\sum m_b = m$. The new matrix is consistent with $d(M)$, contradicting the assumption that (a_2, b) is revealed.

Step 4: Assume that $d(a_1) > d(a_2)$. Let b_1 be a row. Then, either both (a_1, b_1) and (a_2, b_1) are revealed to be occupied or $\lambda(d,(a_1,b_1)) > \lambda(d,(a_2,b_1))$.

Proof: If (a_2, b_1) is revealed, then by Step 3 so is (a_1, b_1). If not, notice first that by Step 2 there is a $d(M)$ –consistent matrix where (a_2, b_1) is vacant and (a_1, b_1) is occupied.

Consider any assignment of values to the elements in $(A - \{a_1, a_2\}) \times B$ which can be extended to a d-consistent matrix. These assigned values determine, for every row b, the number n_b of players that should occupy the entries (a_1, b) and (a_2, b). Let $K = n_{b_1}$. The number of d-consistent matrixes in which (a_2, b_1) is occupied by K players and (a_1, b_1) is vacant is equal to the number of vectors $(x_b)_{b \neq b_1}$ of non-negative numbers that sum to $d(a_2) - K$, such that $x_b \le n_b$ for all $b \neq b_1$. This number is the coefficient of $x^{d(a_2)-K}$ in the polynomial $\Pi_{b \neq b_1}(1 + x + x^2 + \ldots + x^{n_b})$ (which is a product of $|B|-1$ polynomials). Similarly, the number of matrixes that are consistent with (a_1, b_1) being occupied by K players and (a_2, b_1) being vacant is the coefficient of $x^{d(a_1)-K}$ in this polynomial. As shown in Stanley (1989), the sequence of this polynomials' coefficients is symmetric around the "center" ($\sum_{b \neq b_1} n_b/2$ which must be positive since $d(a_1) > d(a_2)$), strictly increasing to the left of it and strictly decreasing to the right of it. Since $\sum_{b \neq b_1} n_b = d(a_1) - K + d(a_2)$ and $d(a_1) > d(a_2)$, it must be that $d(a_1) - K$ is closer to the center ($[d(a_1) - K + d(a_2)]/2)$ than $d(a_2) - K$.

Step 5: Each occupied position in M is picked with positive probability.

Suppose not. Then a player who is located at a position that is not picked with positive probability can switch to a position with a maximum score and, by steps 3 and 4, will be picked with positive probability.

Step 6: Deriving the contradiction.

By steps 1 and 5, no occupied positions are revealed. By step 4, the PS picks, with equal probability, any one of the positions in the box of entries that maximizes the score. By Step 5, all players are positioned inside the box. Not all positions in the box are occupied since the set would then be revealed. Thus, there is a vacant entry (a,b) inside the box. There is a player whose probability of winning is strictly less than $1/n$ (since (a,b) is picked as well). This player's move to the vacant entry will either reveal him and increase his utility to at least $1/n$ or it will make him the sole maximizer of the rank and therefore he will be picked with probability 1.

References

Kraus, Manfred (1996). "A Simple Proof of the Gale-Ryser Theorem". *The American Mathematical Monthly*, *103*, 335-337.

Ryser, Herbert. J. (1963). *Combinatorial Mathematics*, Wiley.

Stanley, Richard P. (1989) "Log-Concave and Unimodal Sequences in Algebra, Combinatorics, and Geometrya". *Annals of the New York Academy of Sciences 576.1*: 500-535.