Nash Equilibria

Vladimir Gurvich

1 Preliminary talks

What is common in Go, Chess, Checkers, and Gomoku? They all are finite positional games with perfect information and without positions of chance. The last means that "all players know everything" and hence, they know the same. This is not the case with cards or domino. A player does (or at least supposed to) not see the opponents' cards. In Backgammon, there are positions of chance, when the dices are tossed. Word "positional" means that moves lead between positions. Word "finite" means that there are finitely many positions.

Any configuration of stones is a position in Go. There are very (but finitely) many of them. A play starts with some initial position and terminates in a final one (which is called a terminal). For example, a mate or a stale mate are terminals in Chess. It is important to notice that positions can be repeated and a play may cycle.

All games considered above are zero-sum two-person games. The word "zero-sum" means that the winning of one is the losing of the other. And if one gets some number of points then the other loses the same number. A pair of strategies is optimal if they form an equilibrium, that is, the corresponding result can be improved by neither of the players.

Yet, everything becomes much more complicated when there are more than two players (or two but the game is not zero-sum). It may happen that each player can guarantee only a very poor result, because it is difficult to fight against the coalition of all other players.

An exercise: There are 10 matches, 3 players take them in a cyclic order, 1,2,3,4, or 5 for one move. One that take the last one should do the dishes. Prove that any two can always force the third one to do the dishes.

Let doing the dishes costs 2, while the winners get +1 each (that is, their cost is (-1) for each). The game is zero-sum. We want to define an equilibrium in this case too, that is, to suggest 3 strategies such that if one player changes his strategy, he cannot profit whenever two opponents keep their old strategies. This concept was introduced by John Nash in 1950 and it is called Nash equilibrium.

An exercise: Suggest some Nash equilibria in the considered game.

Our main problem is to understand which games are Nash-solvable (that is, always have Nash equilibria) and which may have none.

First, let us demonstrate that any acyclic case (in which no position can be repeated) is Nash-solvable.

Let us assign to a game a directed graph (digraph), whose vertices are the positions and the directed edges (arcs) are moves. The positions with no moves are called terminals. Let us assign to each terminal and player a number which this player gets in this terminal. For any non-terminal position we define a player who makes a move in it. Let all moves from a position v lead to terminals. Then the corresponding player chooses a best his move, say to a terminal v'. Then we transfer to v all points defined for v'. Now v "becomes terminal", etc. This procedure is called Backward Induction. It defines strategies of all players.

An exercise: Prove that these strategies form a Nash equilibrium for any acyclic games.

Let us define the concept of a strategy more accurately. A stationary strategy of a player is a rule choosing his move in each his position. For example in the above "match-game" the greedy strategy suggests to take the maximum number of matches each time. Let us note that a stationary strategy does not depend on history, that is, on the preceding positions or moves. In particular, the play must cycle whenever a position appears twice. In Chess this means a draw. However, cycles may have some other costs, for example, all players lose.

Up to now, we considered only the terminal cost functions, that is, the result depended only on the terminal (or cycle). More generally, each player may pay for each (not only for his own) move and the final result is defined as the sum of all these local costs.

An exercise: There are 5 matches. 3 players take them in a cyclic order. One who takes the last match gets a bonus: 3 matches. The result of a player is the number of collected matches. Construct the graph and find an equilibrium.

If a play cycles then we assume that the cycle is passed infinitely many times. In other words, the cost of such a play is either + or $-\infty$ For convenience we will assume that it is $+\infty$ whenever the sum of all local costs ≥ 0.

2 A mini-survey "before accurate definitions"

We will consider the following question:
Which positional games (with perfect information and without moves of chance) have Nash equilibria (in pure stationary strategies)? In some cases, the answer is well known. We start with a mini-survey postponing the precise definitions till the next section. Nash equilibria (NE) exist for the following classes of games:

A. Acyclic games, in which no position can appear twice. In this case, NE always exist. Yet, already in Chess, or even Go, positions can be repeated.

B. Zero-sum two-person games. Both Go and Chess are included. Yet, what if the conflict is not antagonistic? or there are more than two participants?

C. NE exist if players' moves may depend on the preceding moves.
Yet, we will restrict ourselves (and the players) to the pure stationary strategies. In other words, a move in the present position depends only on it, not on preceding positions or moves. Furthermore, this move is chosen deterministically, without any randomization. For example, Backgammon will not be considered.

Let us note, however, that in all of the above cases, A,B, and C, a NE exists even if we allow positions of chance. It is also known that NE may fail to exist in games with imperfect information (like card games or Domino). Yet, we will not consider (and will not even define) them.

Let us summarize:
We will restrict ourselves to games with perfect information, without random moves, and to pure stationary strategies. Yet, the number of players may be greater than two and, even a two-person game, may be not zero-sum.

Somewhat surprisingly, not much is known about the considered games. There are several solution-concepts, among which the NE is certainly most popular. (It will be defined in the next section.) Although, five Nobel Prizes in Economics were granted for works on NE, but in my opinion, the "simplest" and most challenging related mathematical problems are open yet.

Here we will consider two such problems-conjectures. By computers, they are verified for sufficiently (but not too) large examples. So, I hope for affirmative solutions but a counterexample would not be too surprising to me, either.

These conjectures are verified in some special cases, which will serve as exercises. In contrast, some other special cases are difficult, or even open.

3 Main definitions

They all are pretty natural but still so much formalities may scare some.
In this case, I recommend to skip this section during the first reading and use it as a dictionary, only when necessary.

Graphs, positions, and moves. Given a finite directed graph (digraph) $G = (V, E)$, a vertex $v \in V$ is interpreted as a *position*, while a directed edge $e = (v, v')$ is a move from the position v in the corresponding game.

The positions $V_T \subset V$ without any moves are called *terminal.*

Let us choose also an initial position $v_0 \in V \setminus V_T$.

Players. Let us assign to each non-terminal position $v \in V \setminus V$ a *player* $i \in I = \{1, \ldots, n\}$ who will make a move in v, say that v is controlled by i, and write $i = \phi(v)$. In other words, the mapping $\phi : V \setminus V_T \to I$ distributes the non-terminal positions among players.

Triplet $\{G, \phi, v_0\}$ is called a *positional structure.*

Strategies and situations. A *strategy* x_i of player $i \in I$ is a plan choosing a move $e = (v, v')$ in every position $v \in \phi^{-1}(i)$ controlled by i. In other words, the mapping x_i assigns a move $e = (v, v')$ from v to v' for each position $v \in \phi^{-1}(i)$.

These are the so-called *pure stationary strategies*. As we already mentioned, we will not consider, nor even define, any others.

Let each player i choose a strategy x_i. The obtained n-tuple $x = (x_1, \dots, x_n)$ will be called a strategy profile or a *situation*.

Plays. Eeach situation x uniquely defines a play $p(x)$, since each player i knows what to do in every position $v \in \phi^{-1}(i)$ (to move according to x_i).

The play $p(x)$ begins in v_0 and either it terminates an a $v \in V_T$ or "cycles", that is, passes a directed cycle C infinitely many times. (Let us notice that $p(x)$ cannot leave C, since all strategies of x are stationary.)

Thus, we obtain a mapping $g : X \to P$ that assigns a play $p = p(x) \in P$ to every situation $x \in X$. Such a mapping is called a *game form*.

Local and effective cost functions. Each player $i \in I$ pays the value $c(i, e) \in \mathbb{R}$ for each (not only for his own) move $e \in E$. This real number is the *local cost*. (Of course, if $c(i, e) < 0$ then i gets $|c(i, e)|$ rather than pays it.)

A positional structure and a local cost function define a positional game.

The *effective cost* $c(i, p)$ of a play $p = p(x)$ for a player $i \in I$ is defined as follows. If p ends in a terminal position $v \in V$ then $c(i, p) = \sum_{\in p} c(i,)$, that is, the cost is additive and is equal to the sum of the costs of all moves of p for i.
If p cycles on C, we have to compute the cost $c(i, C) = \sum_{\in C} c(i, e)$ of C first. If $c(i, C) \geq 0$ then $c(i, p) = \infty$ and $c(i, p) = -\infty$ when $c(i, C) < 0$.

Such definition is natural. Indeed, play p repeats C infinitely and the local costs are summed up. Yet, when C is a "zero-cycle" for i, that is, $c(i, C) = 0$, we still set $c(i, p) = \infty$. This is just a helpful convention.

A game form $g : X \to P$ together with an effective cost function $c : I \times P \to \mathbb{R}$ define a game (g, c) in the normal form.

Naturally, each player $i \in I$ is trying to *minimize* his effective cost $c(i, p)$.
In particular, all players should avoid non-negative cycles. We will see, however, that it might be not that easy to do.

Terminal moves, costs, and games. A move $e = (v, v')$ is called *terminal* if $v' \in V_T$ is a terminal position. Let us notice that a terminal move cannot belong to a directed cycle. A local cost function c (and the obtained game) are called *terminal* if $c(i, e) \equiv 0$ for each player i and every *non-terminal* move e. In this case, the effective cost of a terminal play p depends only on its last move and if p cycles then its cost is $+$ or $-\infty$, by definition.

Zero-sum games. We say that a local cost function c (as well as the corresponding game) are *zero-sum* if $\sum_{i \in I} c(i, e) = 0$ for every move $e \in E$. Zero-sum two-person, $n = 2$, games play very important role.

Any n-person game can be easily converted into a zero-sum $(n+1)$-person game. Let us introduce a new $(n+1)$st player who will be a dummy (in control of no positions) and define his local cost function $c(n+1, e) = -\sum_{i=1}^{n} c(i, e)$.

Games in normal form; general definition. Let $I = \{1, \dots, n\}$ be a set of players, $_i$ be a finite set of strategies of $i \in I$, and $X = X_1 \times \ldots \times X_n$ be the direct product of these n sets, that is, X is the set of situations.

Furthermore, let P denote an arbitrary set of outcomes (plays, in our case). An arbitrary mapping $g : X \to P$ is called a *game form*.

Finally, given an arbitrary cost function $c : I \times P \to \mathbb{R}$, its real values $c(i, p)$ show how much player $i \in I$ must pay for the play $p \in P$.

Finally, the pair (g, c) is called a *game in normal form*.

Nash equilibria and saddle points.

A situation $x = (x_1, \dots, x_n) \in X_1 \times \ldots \times X_n = X$ is called a *Nash equilibrium* (NE), if no player $i \in I$ can profit by replacing his strategy (x_i by x'_i) provided all other players keep their old strategies. Formally, we can write this as follows:

$c(i, g(x)) \leq c(i, g(x'))$ for every player $i \in I$ and for each situation $x' \in X$ such that all components (strategies) of x' are the same as in x, except for, maybe, the ith one; that is, only x'_i may differ from x_i, while $x'_j = x_j$ for all $j \in I \setminus \{i\}$.

This concept was introduced by John Nash in 1950. For the two-person zero-sum games, NE are *saddle points*, which concept is about 200 years older.

In contrast to saddle points, the concept of NE is vulnerable for criticism. Indeed, two players might change simultaneously their strategies and profit both. Moreover, sometimes all n players may do the same. (In other words, NE are not necessarily Pareto-optimal.) Furthermore, NE (in pure strategies) may fail to exist (as well as the saddle points, yet). Then, NE may be numerous, moreover NE costs might be not unique, either. However, we are here not for criticizing but for studying the concept of NE. (Remember also about five Nobel Prizes :-)

Uniform Nash equilibria. A situation $x \in X$ is called a *uniform* NE , if it is a NE not only with respect to the given initial position v_0 but with respect to any other initial position $v_0' \in V$, as well.

4 Conjectures and problems

We are interested in the NE existence theorems, or in other words in *Nash-solvability* (NS), of the positional games defined above.

The number of points in parentheses measure the complexity of the problems.

Conjecture 1 (500). Prove or disprove that any two-person positional game is Nash-solvable (NS).

Problem 1 (10). Show that, proving Conjecture 1, one can assume, without loss of generality, that the considered game contains no "zero-cycles"; more precisely, $\sum_{e \in C} c(i,e) \neq 0$ for every directed cycle C and player $i \in I = \{1,2\}$.

Let us recall that the effective cost of any cycling play is either + or $-\infty$.

This is a brand new conjecture. It was verified by a recent computer code written by Vladimir Oudalov for many digraphs with 10-18 vertices.

Problem 2 (25). Give an example showing that Conjecture 1 cannot be extended for the three-person case, $n = 3$.

The conjecture does hold for the two-person zero-sum case but all known proofs are difficult.

Moreover, in this case we can introduce a *finite* effective cost function for all plays resulting in zero-cycles such that a saddle point always exists. (Let us recall that we have set $c(i,p) = +\infty$ in the considered case.)

Problem 3 (70). Show that the following "natural redefinings" fail:
$c(i,p) = 0$ or $c(i,p) = \sum_{e \in p} c(i,e)$. In both cases there may be no saddle points. Give examples and try to find the "right" definition.

Conjecture 2 (500). Prove (or disprove) that an n-person positional game is Nash-solvable whenever all its local costs are non-negative.

Problem 4 (5). Show that proving Conjecture 2, without loss of generality, one can restrict himself to the strictly positive costs.

This conjecture remains unproved even in the following very special cases:

Conjecture 2a (300). Does Conjecture 2 hold for the games with a terminal cost function? (In this case the effective cost of every cycling play is $+\infty$.)

Conjecture 2b (400). The same question but now we do not assume that every cycling play is the worst outcome for all n players. Instead, we assume that all such plays form the same outcome but each player can rank this cyclic and all terminal outcomes arbitrarily. Does NS hold?

Conjecture 2c (200). Does conjecture 2 hold for two-person games?

In other words, here we unite the conditions of Conjectures 1 and 2.

Problem 5 (100). Prove that Conjecture 2 holds, yet, for the case of two players and terminal cost functions.

This result can be derived from my old theorem of 1975. By definition, a general n-person game form $g : X_1 \times \ldots \times X_n \to P$ is NS when the corresponding game (g,c) has at least one NE for every cost function $c : I \times P \to \mathbb{R}$.

Here $c(i,p)$ is the cost of the outcome $p \in P$ for the player $i \in I$.

For the two-person case, $I = \{1, 2\}$, let us introduce the next two relaxations of the above condition. A two-person game form g will be called:

zero-sum solvable if it is solvable in the class of the zero-sum games;

± 1 *solvable* if it is solvable in the class of the zero-sum games whose cost function takes the values $+1$ and -1 only.

Problem 5a (100). Prove that all three above properties (solvability, zero-sum solvability, and ± 1 solvability) of the two-person game forms are equivalent.

The equivalence of the last two properties is not difficult to show and I proved it a bit earlier, in 1973. Yet, even earlier this was demonstrated by Jack Edmonds and Delbert Ray Fulkerson, in their paper "Bottleneck extrema", Journal of Combinatorial Theory 8:3 (1970) 299–306.

Unfortunately, the statement of Problem 5a cannot be extended to the three-person case, already. More precisely, let us assign to each n-person game form n two-person game forms in which i is playing against the complementary coalition $I \setminus \{i\}$, for all n players $i \in I$.

Problem 5b (50). Construct a three-person game form that is not NS, while all three corresponding two-person game forms are NS.

Problem 5c (20). Construct an "inverse" example: a three-person NS game form such that all three corresponding two-person game forms are not NS.

Problem 6 (20). Reduce problem 4 to Problem 5.

Problem 7 (15). Prove Nash-solvability for the games on acyclic digraphs.
(A digraph G is acyclic if it has no directed cycles.)
Hint: Make use of the dynamic programming (which is called "backward induction" in game theory.) This result was obtained by Harold Kuhn in 1951 and David Gale in 1953, soon after Nash coined his concept of equilibrium.

Problem 7a (20). Prove that NS holds for the acyclic case, even in the presence of positions of chance (for each of which a probabilistic distribution among possible moves is given).

Of course, solution of the last two problems costs 20 points rather than 35.

Problem 8 (40). Prove that a NE (that is, a saddle point) always exists in the zero-sum two-person positional games.

In particular, for Chess and Go. This theorem belongs to Ernst Zermelo:
"On an Application of Set Theory to the Theory of the Game of Chess",
Proceedings of the Fifth Congress of Mathematicians at Cambridge, 1912.

Let us note that this result also can be extended to allow positions of chance. Yet, this direction would lead us too far to the stochastic game theory. So, we will postpone it for the future.

Problem 9 (10). Let us agree to finish the game as soon as a position is repeated, the obtained cycling play being the outcome. Then, any finite digraph is reduced to a tree (that have no cycles at all, directed or not). Then, why Conjectures 1 and 2 do not result from Problem 7?

Problem 10 (15). Give an example of a terminal two-person game that has a unique directed cycle and no uniform NE. (Here, we do not assume that the cycle is the worst outcome for both players. Each of them may rank it and the terminals arbitrarily.)

Problem 11 (100). Give an example of a terminal two-person game that has a unique directed cycle and no uniform NE, now assuming that the cycle is the worst outcome for both players.

Problem 12 (25). Provide a similar three-person example with a unique directed cycle that is the worst outcome for all players and without uniform NE.

Such examples were obtained just recently: in 2003 for Problem 11 and in 2008 for Problem 12. Of course, in all three cases (of Problems 10, 11, and 12), a NE exists with respect to any fixed initial position. Otherwise, Conjecture 2 would be disproved.