Strategic full-information game

Game definition: $(N, \langle S_i \rangle, \langle u_i \rangle)$

1. $N = \{1, 2, \ldots, n\}$ is the finite set of n players.
2. S_i is the *strategy space* for each player i. Every $a \in S_i$ is a *pure strategy* or *move* for player i. The set of *strategy profiles* S is $S_1 \times S_2 \times \ldots \times S_n$. A strategy profile specifies a pure strategy for each player.
3. u_i is the utility function for each player, $u_i : S \to \mathbb{R}$
4. Game is finite iff S_i is finite for every player.

Note on notation: A strategy profile can be expressed as $s = (s_1, s_2, \ldots, s_n)$. The notation (s_i, s_{-i}) refers to the strategy profile where the ith player plays s_i, and the strategies for all other players are summarized as s_{-i}.

Important Assumptions

1. Each player knows the pure strategies available to all of the players.
2. All players are rational utility maximizers (i.e., risk neutral)
3. There is *common knowledge* about the players as utility maximizers. All players know that all other players are rational. And all players know that all players know this, *ad infinitum.*
4. Players are computationally unbounded.

Pure Nash Equilibrium

A profile $s \in S$ is a *pure Nash equilibrium* iff no player has a unilateral incentive to deviate from s, *i.e.*, for all $i \in N$ and for all $s'_i \in S_i, u_i(s) \geq u_i(s'_i, s_{-i})$. Given s_{-i} the *best response* set for i is $B_i(s_i) = \{x \in S_i | u_i(x, s'_{-i}) = \max_{y \in S_i} u_i(y, s'_{-i})\}$. s is a Nash equilibrium iff $s_i \in B_i(s_{-i})$ for all i.

Note: Player i has a *dominant strategy* a iff $a \in B_i(s_{-i})$, $\forall s_{-i}$. In other words, a strategy is dominant if it is a best response to any combination of opponents' strategies. By way of examples, there is no dominant strategy for either player in Battle of the Sexes or Pennies. In the Prisoner's Dilemma, Defect is a dominant strategy for either player.

Sufficient Condition for Pure Nash Equilibrium

Theorem: A game $(N, \langle S_i \rangle, \langle u_i \rangle)$ has a pure Nash equilibrium if $\forall i$, S_i is a compact convex set in $\mathbb{R}^d$ (for some finite d) and each u_i is a *quasi-concave* function in S_i. A function $f : S \rightarrow R$ is concave if $\forall x, y \in S, \forall \lambda \in [0, 1]$

$$f(\lambda x + (1 - \lambda)y) \geq \lambda f(x) + (1 - \lambda)f(y)$$

The function f is quasi-concave if it satisfies the weaker condition $\forall x, y \in S, \forall \lambda \in [0, 1]$:

$$f(\lambda x + (1 - \lambda)y) \geq min\{f(x), f(y)\}$$

The requirement that u_i be quasi-concave in S_i means that $\forall i$, $u_i : S_i \rightarrow R$ and $\forall s_{-i}$, it is the case that $u_i(\cdot, s_{-i})$ is quasi-concave.

Simple Two-player Games

↓Boy Girl→	Baseball	Softball
Baseball	(3,2)	(1,1)
Softball	(0,0)	(2,3)

In the Battle of the Sexes Game, the girl wants to go to the softball game, while the boy wants to go the baseball game. Both parties would rather be together at their less-prefered event than to go alone. The payoff matrix shows that, for example, if they both go to the baseball game, the utility is 3 for the boy and 2 for the girl. There are two pure Nash Equilibria here: $< softball, softball >$ and $< baseball, baseball >$.

↓P1 P2→	Heads	Tails
Heads	(1,-1)	(-1,1)
Tails	(-1,1)	(1,-1)

In the Pennies game, each player flips a penny. If the outcome is heads-heads or tails-tails, player 1 wins the pennies. If the coins turn up different, player 2 wins the pennies. There is no pure Nash equilibrium. No matter the outcome, the losing player can improve his payoff by switching.

↓P1 P2→	Cooperate	Defect
Cooperate	(-1,-1)	(-3,0)
Defect	(0,-3)	(-2,-2)

In the Prisoner's Dilemma game, the players have been accused of a serious crime, but the evidence is not sufficient to convict. They are interrogated separately and enticed to tesify against each other. If they both cooperate (i.e. with each other, by remaining silent), they will receive light sentences on charges of a lesser crime $(-1, -1)$. If one defects (i.e. by incriminating the other), and the other cooperates, the one who defects goes free and the one who cooperates serves a long sentence $(0, -3), (-3, 0))$. If both players defect, they both serve medium sentences $(-2, -2)$. There is a single pure Nash equilibrium: both players defect. Note that if both cooperate, the payoff for both is higher. But this is not an equilibrium, since either player could improve his payoff by defecting.

Mixed Strategies

A mixed strategy for player i is a probability distribution p_i over S_i. A mixed strategy profile $p = (p_1, p_2, \ldots, p_n)$ specifies a mixed strategy for each player. Given a finite game $(N, \langle S_i \rangle, \langle u_i \rangle)$, a *derived game* is $(N, \langle \Delta(S_i) \rangle, \langle U_i \rangle)$, where

$$\Delta(S_i) = \{\mu \in [0,1]^{|S_i|} | \sum_{j=1}^{|S_i|} \mu_j = 1\}$$

and

$$U_i(p_i, p_{-i}) = \sum_{a \in S_i} p_i(a) u(a, p_{-i}).$$

Note that $\mathcal{P} = \Delta(S_1) \times \Delta(S_2) \times \ldots \times \Delta(S_n)$ is the strategy profile space. Also, in this scenario, the goal of each agent is to maximize its *expected* utility.

Theorem 0.1 *(Nash, 1951): Every finite game has a mixed equilibrium.*

Exercise: Nash's theorem requires S_i's to be finite. There exist games where the S_i are infinite and they do not have a mixed equilibrium.

Theorem 0.2 *(Glicksburg 1953): Every game with S_i compact, convex and u_i continuous on S has a mixed equilibrium.*

Lemma 0.3 *A mixed strategy profile p is an equilibrium iff $\forall i \in N$ and $\forall a \in S_i$ such that $p_i(a) > 0$* i.e.*, a has non-zero support in $p_i, a \in B(p_{-i})$.*

Proof: (Every pure strategy with non-zero support in p_i is a best response to $p_{-i} \Rightarrow$ Nash equilibrium). If s_i is not a best response, obtain p'_i by setting $p'_i(s_i) = 0$ and $\forall s'_i \in B(p_{-i})$

$$p'_i(s'_i) = \frac{p_i(s'_i)}{1 - p_i(s_i)}$$

This redistribution of probability mass ensures that $u_i(p'_i, p_{-i}) > u_i(p_i, p_{-i})$, which contradicts the claim that p was a Nash equilibrium.

($\Leftarrow$) Suppose that every pure strategy with support in p_i is a best response to p_{-i} but that p is not a Nash equilibrium. That means that at least one player can improve its utility by changing its strategy. Suppose $\exists p'_i$ such that $u_i(p'_i, p_{-i}) > u_i(p_i, p_{-i})$. Then $\exists s'_i$ such that $p'_i(s'_i) > p_i(s'_i)$ and $U_i(p'_i, p_{-i}) > U_i(p_i, p_{-i})$. But this contradicts the claim that s_i was a best response to p_{-i}, so again we have a contradiction. $\square$

Interesting fact: Each player is indifferent to pure strategies in its support. The motivation for randomization is to force other players to randomize. Examples of Mixed equilibria are $\langle (3/4, 1/4), (1/4, 3/4) \rangle$ in the Battle of the Sexes game and $\langle (1/2, 1/2), (1/2, 1/2) \rangle$ in the Pennies game. Note that the value of the game in Battle of the Sexes under the mixed equilibrium is lower than the value of the game under the pure strategy equilibria. There is no mixed strategy equilibrium for the Prisoner's Dilemma.

Finite 2-player Bimatrix Games

If $|S_1|$ and $|S_2|$ are finite then the utility functions u_1 and u_2 can be specified by two $|S_1| \times |S_2|$ matrices A,B which specify the payoffs for each player given the pure strategy for player 1 (row index) and player 2 (column index).

Theorem 0.4 *Every finite 2-player game with A,B (payoff matrices) rational has a rational equilibrium x^*, y^* such that x^*, y^* are of size $poly(size(A, B))$.*

Proof: Linear Programming concepts are key to the proof. Fix x, y to be *some* Nash equilibrium. (From the above, we know there must be at least one.) Let $X \subseteq S_1, Y \subseteq S_2$ be the sets of *supports* for x, y. In other words, $X = \{a | a \in S_1 \text{ and } x(a) > 0\}$. We will define *another* Nash equilibrium that satisfies the size properties specified by the theorem. Let us look at all Nash equilibria p, q such that the support of p is X and the support of q is Y. We know that

1. $\sum_{a \in X} p(a) = 1$
2. $\forall a \notin X, p(a) = 0$
3. $\forall a \in X, p(a) > 0$

Similarly,

1. $\sum_{b \in Y} p(b) = 1$
2. $\forall b \notin Y, p(b) = 0$
3. $\forall b \in Y, p(b) > 0$

From the lemma, we know that $\forall a, a' \in X$

$$\sum_{b \in Y} A[a, b]q(b) = \sum_{b \in Y} A[a', b]q(b)$$

Similarly, we know that $\forall b, b' \in Y$

$$\sum_{a \in X} A[a, b]p(a) = \sum_{a \in X} A[a, b']p(a)$$

Using these constraints, we can create a linear program, the solution for which is a Nash equilibrium for the game. The feasible region of the linear program is a polytope. From LP theory, we know that if a linear program has a feasible solution, then there is at least one solution at an extreme point. Such an extreme point solution will be a Nash equilibrium of the appropriate size.

□