Cooperative games– mixed strategies

Moty Katzman

October 24, 2017

Example: matching pennies

Alice and Bob both turn a penny with either heads or tails. If they match, Alice takes them, otherwise Bob takes them.

	H	T
H	1, -1	-1, 1
T	-1, 1	1, -1

Alice hires you as her Game Theory adviser; what do you advise?
Game Theory cannot possibly advise Alice to play H or T; Bob's adviser is just as knowledgable, and if, say, the correct move for Alice is H, Bob will be instructed to play T and win.
There is no correct action here, unless ... we expand the notion of a strategy: Alice is advised to choose her move *randomly*, say, play H with probability α and T with probability $1 - \alpha$. Similarly Bob will choose H with probability β and T with probability $1 - \beta$.

Expected utility

Assumption: Given the choice between two uncertain outcomes, players will choose the one with highest expected utility. (The Von-Neumann, Morgenstein expected utility hypothesis.)

Mixed strategies and mixed-strategy Nash equilibria

Consider a game $(S_1, \ldots, S_n, u_1, \ldots, u_n)$.

Definition. A *mixed strategy* for player i is a function $p_i : S_i \to [0,1]$ such that $\sum_{s_i \in S_i} p(s_i) = 1$. (We interpret $p(s_i)$ as the probability that player i plays s_i.) Given mixed strategies $p_1, \ldots, p_n$ for players $1, \ldots, n$ we define

$$u_i(p_1, \ldots, p_n) = \sum_{(s_1,\ldots,s_n) \in S_1 \times \cdots \times S_n} p_1(s_1) \ldots p_n(s_n) u_i(s_1, \ldots, s_n)$$

(which is the expected payoff for player i when these mixed strategies are played.)

Definition. Mixed strategies $p_1, \ldots, p_n$ for players $1, \ldots, n$ are a *mixed strategy Nash equilibrium* of the game if for all $1 \leq i \leq n$ and any mixed strategy p_i' for player i we have

$$u_i(p_1, \ldots, p_{i-1}, p_i, p_{i+1} \ldots p_n) \geq u_i(p_1, \ldots, p_{i-1}, p_i', p_{i+1} \ldots p_n).$$

(Each p_i gives a response with highest expected value to the other mixed strategies.)

Example: matching pennies

Alice plays p, $p(H) = \alpha$, $p(T) = 1 - \alpha$;
Bob plays q, $q(H) = \beta$, $q(T) = 1 - \beta$. We have

$$u_1(p, q) = \alpha\beta + (1-\alpha)(1-\beta) - \alpha(1-\beta) - (1-\alpha)\beta = (4\beta-2)\alpha + 1 - 2\beta$$

$u_2(p, q) = -u_1(p, q)$.
To maximize u_1 for a given β, we take $\alpha = 1$ if $\beta > 1/2$, $\alpha = 0$ if $\beta < 1/2$ and any α when $\beta = 1/2$. Similarly, u_2 is maximized by $\beta = 0$ if $\alpha > 1/2$, $\beta = 1$ if $\alpha < 1/2$ and any β when $\alpha = 1/2$.
The only mixed strategy Nash equilibrium occurs with $\alpha = \beta = 1/2$.

Example: The battle of the sexes

Alice and Bob decide to go on a date together: Alice suggests going to the pub and Bob suggests watching a film. Before they can agree, the MAS348 lecture is over and they rush to different lectures. Realizing they don't have a way to contact each other before their date, they have to decide independently where to go. They are playing the following game:

	Pub	Film
Pub	3,1	0,0
Film	0,0	1,3

This game has two Nash equilibria in pure strategies, namely (Pub, Pub) and $(Film, Film)$. Are there any other Nash equilibria?

	Pub	Film
Pub	3,1	0,0
Film	0,0	1,3

Suppose Alice goes to the pub with probability p and Bob goes to the film with probability q. Alice's expected utility is $3p(1-q) + (1-p)q = p(3-4q) + q$ and Bob's expected utility is $p(1-q) + 3(1-p)q = q(3-4p) + p$. To find Alice's best responses we distinguish between three cases:

$q < 3/4$ and Alice's expected utility is maximized when $p = 1$,

$q > 3/4$ and Alice's expected utility is maximized when $p = 0$,

$q = 3/4$ now Alice's expected utility is $3/4$, regardless of p.

Similarly, to find Bob's best responses we distinguish between three cases:

$p < 3/4$ and Bob's expected utility is maximized when $q = 1$,

$p > 3/4$ and Bob's expected utility is maximized when $q = 0$,

$p = 3/4$ now Bob's expected utility is $3/4$, regardless of q.

The strategy profile corresponding to $p = q = 3/4$ is a Nash equilibrium!

Note: in a Nash equilibrium each player had an expected utility independent of the strategy adopted by the other player– *this is not a coincidence.*

Notation. For any action s of a player in a game, $\widehat{s}$ denotes the mixed strategy which plays s with probability 1.

Theorem. (The Indifference Principle) Consider a game (S_1, S_2, u_1, u_2). Let (p, q) be a mixed strategy profile which is a Nash equilibrium. For any $s \in S_1$ with $p(s) > 0$ we have $u_1(\widehat{s}, q) = u_1(p, q)$.

Proof. Let $\mathrm{Supp}(p) = \{s \in S_1 \mid p(s) > 0\}$ (the "support" of p). We have $u_1(p, q) = \sum_{a \in \mathrm{Supp}(p)} p(a) u_1(\widehat{a}, q)$.

If the theorem does not hold, there must be an $a \in \mathrm{Supp}(p)$ for which $u_1(\widehat{a}, q) > u_1(p, q)$, contradicting the fact that p is a best response to q. $\square$

Example: Rock-Scissors-Paper

Consider the game

	Rock	Scissors	Paper
Rock	0, 0	1, -1	-1, 1
Scissors	-1, 1	0, 0	1, -1
Paper	1, -1	-1, 1	0, 0

There are no dominant strategies, nor are there Nash equilibria in pure strategies. To find mixed strategy profiles which are Nash equilibria, find a mixed strategy (p_1, p_2, p_3) for Alice which gives Bob equal utility for all his pure strategies: this happens in $p_2 - p_3 = -p_1 + p_3 = p_1 - p_2 = 0$, which together with $p_1 + p_2 + p_3 = 1$ yields the mixed strategy $(1/3, 1/3, 1/3)$. Similarly, Bob's strategy is $(1/3, 1/3, 1/3)$.

Example: Alice and Bob play tennis

Bob is at the net and Alice needs to decide to hit the ball to Bob's right or left; and Bob needs to decide to jump to left or right. The probabilities of Bob responding are as follows:

	left	right
left	50%	20%
right	10%	80%

and we assume Alice and Bob are playing the following game

	left	right
left	0.5, 0.5	0.8, 0.2
right	0.9, 0.1	0.2, 0.8

	left	right
left	0.5, 0.5	0.8, 0.2
right	0.9, 0.1	0.2, 0.8

There are no dominant strategies, nor are there Nash equilibria in pure strategies.
When is $((p, 1-p), (q, 1-q))$ a Nash equilibrium?
this happens when $0.5p + 0.1(1-p) = 0.2p + 0.8(1-p)$, which gives $p = 0.7$ and $0.5q + 0.8(1-q) = 0.9q + 0.2(1-q)$, which gives $q = 0.6$.

Example: The attrition game

Consider the game

	Stay	Leave
Stay	-2, -2	1, 0
Leave	0, 1	0, 0

There are two pure-strategy Nash equilibria: (Stay, Leave) and (Leave, Stay). There is also a mixed-strategy Nash equilibrium: ((1/3,2/3), (1/3, 2/3)).
Change this to

	Stay	Leave
Stay	-2, -2	x, 0
Leave	0, 1	0, 0

for $x > 1$. We preserve the pure-strategy Nash equilibria, but the mixed-strategy Nash equilibrium is now: ((1/3,2/3), (x/(2+x), 2/(2+x))).

Change this to

	Stay	Leave
Stay	-2, -2	1, 0
Leave	y, 1	y, 0

for $-2 < y < 1$. How does the mixed-strategy Nash equilibrium change as y increases?

Dominating mixed strategies

Even when no action strictly dominates another, a mixed strategy might.
Consider the game

	L	R
U	3, 1	0, 2
M	1, 1	1, 0
D	0, 0	3, 1

None of the actions is dominated, but playing U and D with probability 1/2 each dominates M.
After eliminating M, Bob's R dominates L, and we obtain a solution (D, R) for this game.

Existence of Nash equilibria

Example. (A game with no Nash equilibrium)
Two players choose an integer. The player with the biggest integer gets 1, the other 0, and in case of a tie both get 0.
Example. (Another game with no Nash equilibrium)
$n > 1$ players choose a number in the open interval $(0, 1)$ and the mean μ of these numbers is computed. The players closest to $\mu/2$ win, the others lose.
Theorem. (Nash's Theorem)
Any game $(S_1, \ldots, S_n, u_1, \ldots, u_n)$ with finite strategy sets has at least one Nash equilibrium.

The proof of this theorem is quite intricate!

Nash Bargaining

In this section we do something new: rather than finding equlibria of games, we look for equitable outcomes.

Alice and Bob play the following game in tabular form

	L	M	R
U	2, -1	-2, 1	1, 1
D	-1, 2	0, 2	1, -2

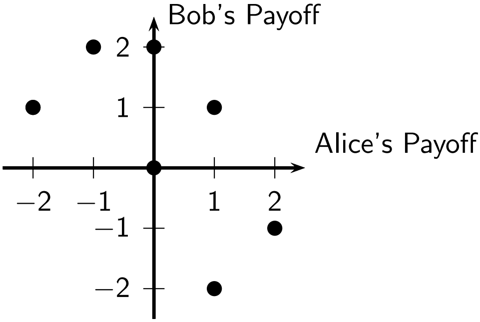

What are the possible outcomes?

Convex sets

Alice's mixed strategy (p_1, p_2) versus Bob's (q_1, q_2, q_3) yields expected payoffs $(A, B) = p_1 q_1(2, -1) + p_1 q_2(-2, 1) + p_1 q_3(1, 1) + p_2 q_1(-1, 2) + p_2 q_2(-0, 2) + p_2 q_3(1, -2)$. (Note that $p_1 q_1 + p_1 q_2 + p_1 q_3 + p_2 q_1 + p_2 q_2 + p_2 q_3 = p_1(q_1 + q_2 + q_3) + p_2(q_1 + q_2 + q_3) = p_1 + p_2 = 1$.) Another mixed strategy profile $((p_1', p_2'), (q_1', q_2', q_3'))$ yields (A', B') and for any $0 \leq t \leq 1$ $(t(p_1, p_2) + (1-t)(p_1', p_2'), t(q_1, q_2, q_3) + (1-t)(q_1', q_2', q_3'))$ yields $t(A, B) + (1-t)(A', B')$.

Definition. A subset S of $\mathbb{R}^d$ is *convex* if for any $v, w \in S$, $tv + (1-t)w \in S$ for all $0 \leq t \leq 1$.

Examples: feasible expected payoffs of games in normal form,vector-subspaces of $\mathbb{R}^d$, cones, intersections of convex sets, etc.

The convex hull

Definition. The *convex hull* of a set $A \subseteq \mathbb{R}^d$ is, equivalently,

(a) the intersection of all convex sets containing A,

(b) the smallest convex set containing A,

(c) the subset C of $\mathbb{R}^d$ defined as the set of all points of the form $\lambda_1 a_1 + \cdots + \lambda_k a_k$ where $k \geq 1$, $\lambda_1, \ldots, \lambda_k$ are non-negative real numbers, $\lambda_1 + \cdots + \lambda_k = 1$ and $a_1, \ldots, a_k \in A$ (*convex combinations* of points in A.)

The set of feasible expected payoffs of a game $G = (S, T, u_1, u_2)$ is a convex set containing $\{(u_1(s,t), u_2(s,t)) \mid s \in S, t \in T\} \subseteq \mathbb{R}^2$, and any feasible expected payoff is a convex combination of these points.

Definition. Let $G = (S, T, u_1, u_2)$ be a game. The *cooperative payoff region* of G is the convex hull of $\{(u_1(s,t), u_2(s,t)) \mid s \in S, t \in T\} \subseteq \mathbb{R}^2$. The points in this region are the possible expected utilities that the players can obtain by playing all possible mixed strategy profiles.

Nash Bargaining

Alice and Bob play (S, T, u_1, u_2) ask Charlie to suggest an outcome $\Psi : (P, (a_0, b_0)) = (a^*, b^*) \in P$ (the cooperative payoff region) and, if rejected, $(a_0, b_0) \in P$ will be imposed.

$\Psi : (P, (a_0, b_0)) = (a^*, b^*)$ must satisfy:

(1) *Individual Rationality:* $a^* \geq a_0$ and $b^* \geq b_0$,

(2) *Pareto Optimality:* if $(a, b) \in P$, $a \geq a^*$, $b \geq b^*$, then $(a, b) = (a^*, b^*)$,

(3) *Independence of Irrelevant Alternatives:* if $P' \subseteq P$ is convex, and contains (a_0, b_0) and (a^*, b^*), $\Psi : (P', (a_0, b_0)) = (a^*, b^*)$,

(4) *Invariance Under Linear Transformations:* If $P' = \{(\alpha a + \beta, \gamma b + \delta \,|\, (a, b) \in P\}$ for some constants $\alpha > 0, \beta, \gamma > 0, \delta$,
$\Psi : (P', (\alpha a_0 + \beta, \gamma b_0 + \delta)) = (\alpha a^* + \beta, \gamma b^* + \delta)$,

(5) *Symmetry:* If P is symmetric (i.e., $(x, y) \in P \Rightarrow (y, x) \in P$) and $a_0 = b_0$, $a^* = b^*$.

Does such arbitration procedure exist? If so, is it unique?

Theorem: (Nash) There exists a unique arbitration procedure Ψ which satisfies all five conditions above.

A quick review of calculus and a Lemma

Review: A continuous function $f : S \to \mathbb{R}$ defined on a closed, bounded set $S \subseteq \mathbb{R}^n$, attains its maximum in S, i.e., there exists an $s \in S$ such that $f(s) = \sup_{\sigma \in S} f(\sigma)$.

Lemma: Let $P \subset \mathbb{R}^2$ be closed, convex and bounded, $(a_0, b_0) \in P$, let $K = \{(a, b) \in P \mid a > a_0, b > b_0\}$ be non-empty. The function $f : K \to \mathbb{R}$ be defined as $f(a, b) = (a - a_0)(b - b_0)$ attains its maximum M on K at a unique point.

Proof: Let $\overline{K} = \{(a, b) \in P \mid a \geq a_0, b \geq b_0\}$. Since $\overline{K}$ is the intersection of two closed and convex sets, it is also closed and convex, and f attains its maximum on $\overline{K}$. Since f is an increasing function in a and b, the maximum is attained at K.

Assume there are distinct (a_1, b_1), (a_2, b_2) on which the maximum is attained; note that ($a_1 > a_2$ and $b_1 < b_2$) or ($a_1 < a_2$ and $b_1 > b_2$), and in either case $(a_1 - a_2)(b_2 - b_1) > 0$. Write $(a_3, b_3) = (1/2)(a_1, b_1) + (1/2)(a_2, b_2) \in K$.

The proof of Nash's Theorem

Now

$$
\begin{aligned}
f(a_3, b_3) &= ((a_1 + a_2)/2 - a_0)\,((b_1 + b_2)/2 - a_0) \\
&= (1/4)\,((a_1 - a_0) + (a_2 - a_0))\,((b_1 - b_0) + (b_2 - b_0)) \\
&= (1/4)[(a_1 - a_0)(b_1 - b_0) + (a_1 - a_0)(b_2 - b_0) \\
&\quad +(a_2 - a_0)(b_1 - b_0) + (a_2 - a_0)(b_2 - b_0)] \\
&= M + (1/4)[-(a_1 - a_0)(b_1 - b_0) + (a_1 - a_0)(b_2 - b_0) \\
&\quad +(a_2 - a_0)(b_1 - b_0) - (a_2 - a_0)(b_2 - b_0)] \\
&= M + (1/4)\,[(b_2 - b_0)(a_1 - a_2) + (b_1 - b_0)(a_2 - a_1)] \\
&= M + (1/4)\,[(a_1 - a_2)(b_2 - b_1)] > M,
\end{aligned}
$$

a contradiction.

The proof of Nash's Theorem

We describe Ψ explicitly as follows. Let $K = P \cap \{(a,b) \in \mathbb{R}^2 \mid a > a_0, b > b_0\}$, $f(a,b) = (a-a_0)(b-b_0)$. **Case 1:** $K \neq \emptyset$. Use previous Lemma to find unique $(a^*, b^*) \in K$ on which f attains its maximum, and declare $\Psi(P,(a_0,b_0)) = (a^*, b^*)$. **Case 2a:** $K = \emptyset$, $(a_0, b) \in P$ for some $b > b_0$. Let b^* be the largest b for which this occurs; since P is closed, $b^* \in P$. Declare $\Psi(P,(a_0,b_0)) = (a_0, b^*)$. **Case 2b:** $K = \emptyset$, $(a, b_0) \in P$ for some $a > a_0$. Let a^* be the largest a for which this occurs; since P is closed, $a^* \in P$. Declare $\Psi(P,(a_0,b_0)) = (a^*, b_0)$. **Case 3:** $K = \emptyset$, neither 2(a) or 2(b) hold. Declare $\Psi(P,(a_0,b_0)) = (a_0, b_0)$. Note that cases (2a) and (2b) cannot both hold, otherwise $K \neq \emptyset$.

We now verify that Ψ has the desired properties.

Individual Rationality is satisfied.
Pareto Optimality: If $(a, b) \in P$ with $a \geq a^*$ and $b > b^*$, in case (1) we have a contradiction $f(a, b) > f(a^*, b^*)$, and in case (2a) we have $a^* = a_0$, $b > b^*$, a contradiction. Similarly if $a > a^*$ and $b \geq b^*$.
Independence of Irrelevant Alternatives: In case (1), the maximum of f on P' is still attained at (a^*, b^*). Similarly in other cases.
Invariance Under Linear Transformations: If Case 1 holds for $P, (a_0, b_0)$ it also holds for $P', (\alpha a_0 + \beta, \gamma b_0 + \delta)$, and also, since $\alpha\gamma > 0$,
$(\alpha a + \beta - (\alpha a_0 + \beta))(\gamma b + \delta - (\gamma b_0 + \delta)) = \alpha\gamma(a - a_0)(b - b_0)$ attains its maximum at $(\alpha a^* + \beta, \gamma b^* + \delta)$. The other cases are similar.
Symmetry: In case (1), this follows from the symmetry of f, and the previous Lemma. Cases (2a) and (2b) are impossible, and Case (3) is easy.

uniqueness

Assume $\overline{\Psi}$ is a different procedure with the same properties, choose closed convex P and $(a_0, b_0) \in P$ such that $\Psi(P, (a_0, b_0)) \neq \overline{\Psi}(P, (a_0, b_0))$.
Assume first we are in Case (1), define the linear function $L(a, b) = ((a - a_0)/(a^* - a_0), (b - b_0)/(b^* - b_0))$, write $(a', b') = L(a, b)$, $P' = L(P)$, and note that $L(a_0, b_0) = (0, 0)$ and $L(a^*, b^*) = (1, 1)$. *Invariance Under Linear Transformations* implies that $\Psi(P', (0, 0)) = (1, 1)$ and $\overline{\Psi}(P', (0, 0)) \neq (1, 1)$.
Claim: for any $(a', b') \in P'$, $a' + b' \leq 2$. Assume not, and define $h(t) = f(t(a', b') + (1 - t)(1, 1)) = (ta' + (1 - t))(tb' + (1 - t))$.
We have $h(0) = 1$, $dh/dt(0) = a' + b' - 2 > 0$ by assumption so there is a small $t > 0$ such that $f(t(a', b') + (1 - t)(1, 1)) > 1 = f(1, 1)$, a contradiction.

uniqueness

Let $\hat{P}$ be the *symmetric convex hull* of P be the convex hull of $P' \cup \{(b,a) \,|\, (a,b) \in P'\}$. We also have $a' + b' \le 2$ for $(a', b') \in \widehat{P}$. Hence if $(a,a) \in \widehat{P}$, $a \le 1$; we deduce that $\overline{\Psi}(\widehat{P}, (0,0)) = (1,1)$, and the *Independence of Irrelevant Alternatives* implies that $\overline{\Psi}(P', (0,0)) = (1,1)$, a contradiction. <u>in Case 2(a),</u> write $\overline{\Psi}(P, (a_0.b_0)) = (\overline{a}, \overline{b})$. Individual Rationality implies that $\overline{a} = a^*$, and Pareto Optimality applied to $(a^*, \overline{b})$ and (a^*, b^*) implies $\overline{b} = b^*$. <u>Case 2(b)</u> is similar, and <u>Case 3</u> is easy. $\square$

An Example

Alice and Bob bargain over the following game in tabular form

	L	M	R
U	2, -1	-2, 1	1, 1
D	-1, -2	0, 2	1, -2

and if they fail to reach agreement they both get 0.

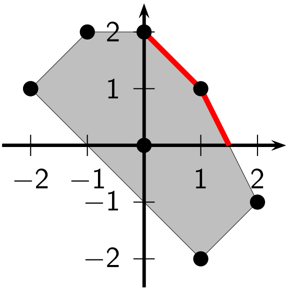

An Example

The Nash bargain will be somewhere on the red boundary above which is the union of two line segments given parametrically by
$t(0, 2) + (1 - t)(1, 1) = (1 - t, t + 1)$ and
$t(1, 1) + (1 - t)(3/2, 0) = (-t/2 + 3/2, t)$ with $0 \leq t \leq 1$.
Now $f(1 - t, t + 1) = (1 - t)(t + 1)$ attains its maximum value 1 at $t = 0$, and this corresponds to the point $(1, 1)$;
$f(-t/2 + 3/2, t) = (3 - t)t/2$ restricted to $0 \leq t \leq 1$ attains its maximum value 1 at $t = 1$, and this also corresponds to the point $(1, 1)$. We conclude that the Nash bargain results in expected utility 1 for both Alice and Bob.
If we modify the example so that failure to agree results in a payoff $(-2/5, 1)$, the possible bargain must lie on the first line segment above, and now
$f(1 - t, t + 1) = (1 - t + 2/5)(t + 1 - 1) = (7/5 - t)t$ attains its maximum value $t = 7/10$. We conclude that the Nash bargain results in expected utility $3/10$ for Alice and $17/10$ for Bob.