Game Theory and Algorithms[*]
Lecture 4: Mixed Strategies & Mixed Nash Equilibria

March 8, 2011

Summary: The ability for players to randomize their choices gives *mixed strategies*, in contrast to the *pure strategies* we have considered previously. To analyze mixed strategies we introduce a stronger assumption on players' preferences. In a later lecture we will prove a Nash equilibrium in mixed strategies (*mixed Nash equilibrium*) exists for every finite strategic game.

1 Mixed Strategies

We previously saw the example of Matching Pennies:

p1 \p2	H	T
H	$1, -1$	$-1, 1$
T	$-1, 1$	$1, -1$

A *mixed strategy* for player i means a fixed *probability distribution* from which player i will select their choice. We denote such a probability distribution by α_i, which is a vector containing a nonnegative real number $\alpha_i(a_i)$ for each $a_i \in A_i$, such that their sum is 1:

$$\alpha_i(a_i) \geq 0 \forall a_i \in A_i, \qquad \sum_{a_i \in A_i} \alpha_i(a_i) = 1.$$

A motivating example occurs in Matching Pennies: if you have no idea what your opponent will do, rather than commit to a fixed strategy where you could lose 1 dollar, if you make your choice uniformly at random (50% of the time heads and 50% tails) then the *expected value* of your gain/loss will be 0:

$$\text{if } \alpha_1(H) = \alpha_1(T) = 1/2, \text{ the expected value of } (u_1, u_2) \text{ is}$$

p1 \p2	H	T
α_1	$\frac{1}{2}1 + \frac{1}{2}(-1) = \mathbf{0}, \frac{1}{2}(-1) + \frac{1}{2}1 = \mathbf{0}$	$\frac{1}{2}(-1) + \frac{1}{2}1 = \mathbf{0}, \frac{1}{2}1 + \frac{1}{2}(-1) = \mathbf{0}$

[*] Lecture Notes for a course given by David Pritchard at EPFL, Lausanne.

(For a random variable x which takes on one of several values x_i with probabilities p_i, its *expected value* is $\sum_i x_i p_i$, i.e. its weighted mean value. This relates to repeated games since if we take many samples $x^{(1)}, x^{(2)}, \ldots$, the expected value equals the limit $\lim_{n\to\infty} \frac{1}{n} \sum_{i=1}^{n} x^{(i)}$.)

Assumption (von Neumann-Morgenstern preferences & Bernoulli payoffs)**.** *We assume for each player that a* randomized outcome with expected utility x *is equivalent to a* deterministic outcome with utility x *in terms of that player's preferences.*

Hence for a game with preferences and payoffs of the type described above, if we are given all values $u_i(a)$ for players i and outcomes a, by extension we compute the utility of a *mixed strategy profile* α by

$$u_i(\alpha) = \mathop{\mathsf{E}}_{\text{each } a_i \text{ drawn independently according to } \alpha_i} [u_i(a)] = \sum_{\text{all outcomes } a} u_i(a) \prod_{i=1}^{n} \alpha_i(a_i).$$

The above assumption is made almost universally; and in order to do any reasonable probabilistic analysis with a game, an assumption along the above lines needs to be made. However, keep in mind the following:

- Two games which are equivalent (in terms of players' preferences over outcomes) for pure strategies may not be equivalent under mixed strategies. For example, consider the original Taxpayers' game on the left, and version with actual dollar values on the right.

p1 \p2	A	DA
L	1,3	4,1
DL	2,2	3,3

p1 \p2	A	DA
L	-5000,0	15000,-15000
DL	-200,-1000	0,0

  So, a bad choice of utility function could make mixed-strategy modelling unrealistic.

- Modelling with dollar values is not von Neumann-Morgenstern in general. For example, consider the following two scenarios. In scenario 1, I give you \$1. In scenario 2, I flip a fair coin and give you \$1000002 if it lands heads, but you must give me \$1000000 if it lands tails. Are the scenarios equivalent from your perspective?

we will consider only situations where we restrict each player to randomize *independently* of one another. (The alternative where players can coordinate their randomization, called *correlated strategies/equilibria*, is also studied.)

A *pure* strategy means a deterministic strategy, i.e. with no randomness, the type we have studied up until today. Of course, every pure strategy a_i can also be viewed as a mixed strategy, where α_i has one component $\alpha_i(a_i)$ equal to 1 and all other components equal to 0.

1.1 Application: Strict Domination

One important possibility in games is that a mixed strategy can strictly dominate some pure strategy. Consider the following game, where *we only show the payoffs for player 1*:

p1 \p2	L	R
T	3	0
M	0	3
B	1	1

If player 1 seeks to maximize their own (von Neumann-Morgenstern) utility, then it is never a good idea to play B, since it is strictly dominated by the mixed strategy[1] $(T + M)/2$:

p1 \p2	L	R
(T+M)/2	3/2	3/2
B	1	1

Similarly, any mixed strategy for player 1 which assigns positive probability to B can be strictly improved (regardless of the opponents' actions) by setting the probability of B to zero and moving half of this probability each to T and M. *Iterated elimination* of strictly dominated strategies using mixed dominators enjoys the same properties as with pure dominators:

- The output of the algorithm does not depend on which choice we make when there are multiple dominated actions.
- Mixed Nash equilibria (see the next section) are never eliminated.

On the one hand, this new iterated elimination algorithm is better than the old algorithm in the sense that we always eliminate at least as much as before, and sometimes strictly more. But can it be efficiently implemented? (Previously, we could check for dominated strategies by checking all possibilities, but that is no longer possible.)

Exercise. Show, by using a polynomial-time algorithm to decide feasibility of linear inequality systems[2], that there is a polynomial-time algorithm to determine whether a given game has any pure strategy that is strictly dominated by any mixed strategy.

2 Mixed Nash Equilibria

The definition of Nash equilibria extends naturally to mixed strategies.

[1]This denotes the mixed strategy α_1 such that $\alpha_1(T) = \alpha_1(M) = 1/2$ and $\alpha_1(B) = 0$.

[2]I.e., a polynomial-time algorithm whose input is a series $\{\sum_j A_{ij}x_j \geq b_i\}_i$ of linear inequalities, and which either outputs a feasible solution x or determines that none exist.

Definition 1. A mixed strategy profile α is a *mixed Nash equilibrium* if for each player i and each of their mixed strategies α'_i,

$$u_i(\alpha) \geq u_i(\alpha'_i, \alpha_{-i}).$$

Equivalently, each player uses a best mixed response to their opponents.

It is easy to check that mixed Nash equilibria generalize pure Nash equilibria:

Theorem 2. *Let a be a strategy profile. Let α be the mixed strategy profile corresponding to a. Then α is a mixed Nash equilibrium if and only if a is a Nash equilibrium.*

Proof. One direction,"α is mixed NE" implies "a is NE" is trivial: if α is a mixed NE, since any pure deviation a_i by player i also can be thought of as a mixed deviation, there cannot be any profitable pure deviation from a.

The direction "a is NE" implies "α is mixed NE" is non-trivial: if player i has no profitable pure deviation, is it also true that they have no profitable mixed deviation? Yes: any random variable (here $u(\overline{a_i}, a_{-i})$ with $\overline{a_i}$ drawn from distribution α_i) whose value is always less than some constant ($u(a)$) also has its expected value less than that constant. In detail,

$$\begin{aligned} u(\alpha'_i, \alpha_{-i}) = u(\alpha'_i, a_{-i}) &= \sum_{\overline{a_i} \in A_i} \alpha'_i(\overline{a_i}) u(\overline{a_i}, a_{-i}) \\ &\leq \sum_{\overline{a_i} \in A_i} \alpha'_i(\overline{a_i}) u(a) \\ &\leq u(a) = u(\alpha) \end{aligned}$$

Thus as needed, no player has a profitable mixed deviation from α. □

One important fact about mixed Nash equilibria is Nash's Theorem:

Theorem 3 (Nash, 1950)**.** *Every strategic game with a finite number of players and a finite number of actions per player has a mixed Nash equilibrium.*

This is why they are called Nash equilibria. After some small examples, we will give prove Nash's Theorem in three increasingly general settings:

- First, we prove it for zero-sum two-player games, using linear program duality and maxminimization.
- Next, we prove it for all two-player games, using the existential non-polynomial-time algorithm of Lemke and Howson.
- Finally, we prove it for all finite games, using a topological fixed-point theorem.

Zero-sum two-player games are the most general class for which a polynomial-time algorithm is known to find a NE. For anything more general (non-zero-sum, or > 2 players) the problem is "PPAD-hard." We will make this more precise later; conventional (unproven) wisdom says there is no polynomial-time algorithm for PPAD-hard problems. This is most interesting: we know a Nash equilibrium exists, but we cannot efficiently find it!

3 Examples

Consider the Matching Pennies game again. It has no pure Nash equilibria. But we can easily check that the mixed action profile $((H+T)/2, (H+T)/2)$ is a mixed Nash equilibrium: it gives both players a utility of 0, and if either player deviates to another mixed strategy, they still get an expected payoff of 0. In fact, this is its *only* mixed Nash equilibrium.

Claim 4. *In Matching Pennies, the only mixed Nash equilibrium is* $((H+T)/2, (H+T)/2)$.

Proof. We will use a *mixed best response* argument. To simplify things, let x be the probability that player 1 picks H (so they pick T with probability $1-x$) and y be the probability that player 2 picks H. We abuse notation and identify the mixed strategies of player 1 with their single choice $x \in [0,1]$ and similarly for player 2. What x is a best response to any given y? The utility for player 1 is

$$xy - (1-x)y - (1-y)x + (1-x)(1-y) = 4xy - 2x - 2y + 1 = x(4y-2) + (1-2y).$$

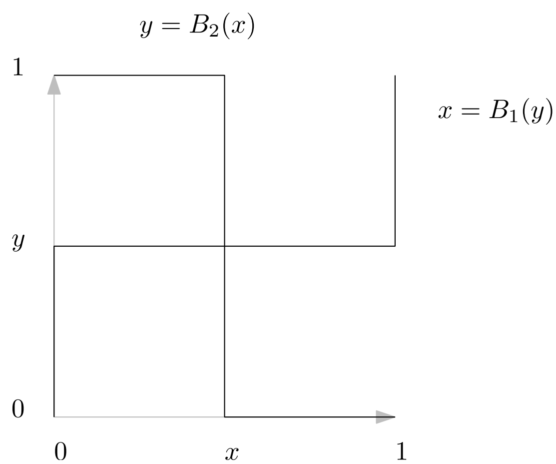

The last version makes it easy to compute the best response function B_1 for player 1: if $y > 1/2$ then the only best response is $x = 1$, if $y < 1/2$ the only best response is $x = 0$, and if $y = 1/2$ then all x are best responses. Similarly (but we skip the details), the best response for player 2 when $x < 1/2$ is $y = 1$, when $x > 1/2$ it is $y = 0$, and when $x = 1/2$ all y are best responses. If we plot both of the best response functions, we see that they have exactly one point in common, $x = y = 1/2$. This implies the claim. □

The same method from this proof works for any 2-player game with 2 actions per player:

Exercise. Find all mixed Nash equilibria of the Bach or Stravinksy game,

p1 \p2	B	S
B	2, 1	0,0
S	0,0	1, 2

By combining this approach with iterated elimination, you can solve the following question:

Exercise. Find all mixed Nash equilibria of the following game:

$p1\backslash p2$	L	C	R
T	3, 4	5, 3	2, 3
M	2, 5	3, 9	4, 6
B	3, 1	2, 5	7, 4

Here is a 2-player example where the number of choices per player is larger than 2:

Exercise. Consider *Matching n-ies*, the 2-player strategic game where both action sets A_i are $\{1, \ldots, n\}$; when $a_1 = a_2$ player 1 wins \$1 and player 2 loses \$1; when $a_1 \neq a_2$ player 1 loses \$1 and player 2 wins \$1. Find all mixed Nash equilibria of this game.

3.1 Support Characterization

Here is one tool that will be very useful to compute Nash equilibria in some games.

Definition 5. For a mixed strategy α_i, its *support* $\texttt{supp}(\alpha_i)$ is the subset of A_i receiving positive probability,

$$\texttt{supp}(\alpha_i) := \{a_i \in A_i \mid \alpha_i(a_i) > 0\}.$$

Theorem 6 (Support Characterization)**.** *Let α be a mixed strategy profile. Let $B_i(\alpha_{-i})$ denote the best pure responses of player i to their opponents,*

$$B_i(\alpha_{-i}) := \arg\max_{a_i \in A_i} u_i(a_i, \alpha_{-i}).$$

Then α is a mixed Nash equilibrium if and only if, for each player i, $\texttt{supp}(\alpha_i) \subseteq B_i(\alpha_{-i})$.

Note that this agrees with the characterization of best responses we saw in Matching Pennies.

Proof. The theorem is a combination of two observations:

- α is a mixed NE iff each player is using a best mixed response to their opponents (which follows from the definition of a mixed NE)
- α_i is a best mixed response for player i if and only if $\texttt{supp}(\alpha_i) \subseteq B_i(\alpha_{-i})$.

We only need to prove the second observation. Let $X = \max_{a_i \in A_i} u_i(a_i, \alpha_{-i})$ denote the value of player i's best pure response to α_{-i}. So

$$u_i(a_i', \alpha_{-i}) \begin{cases} = X, & \text{if } a_i' \in B_i; \\ < X, & \text{if } a_i' \notin B_i. \end{cases}$$

The first step is to show that player i using *any mixed strategy* against α_{-i} can obtain at most X (i.e., $u_i(\alpha_i', \alpha_{-i}) \leq X$ for all α_i') — this is similar to the proof of Theorem 2.

What mixed strategies are optimal for player i? First we claim that if $\texttt{supp}(\alpha_i) \subseteq B_i(\alpha_{-i})$, then α_i is a best response. Intuitively, this is clear since if player i uses a probability distribution over certain strategies, each of which attain a value of X, his expected value over this distribution is also X. We write it explicitly for clarity,

$$u_i(\alpha_i, \alpha_{-i}) = \sum_{a_i \in \alpha_i} \alpha_i(a_i) u_i(a_i, \alpha_{-i}) = \sum_{a_i \in \texttt{supp}(\alpha_i)} \alpha_i(a_i) u_i(a_i, \alpha_{-i}) = \sum_{a_i \in \texttt{supp}(\alpha_i)} \alpha_i(a_i) X = X \cdot 1$$

where in the last inequality we factored out the X and used the fact that the probabilities in α_i add up to one.

Finally, if $\texttt{supp}(\alpha_i) \nsubseteq B_i(\alpha_{-i})$ we need to show that α_i is *not* a best response. Note that $\texttt{supp}(\alpha_i) \nsubseteq B_i(\alpha_{-i})$ is the same as saying there is some a_i^* with $\alpha_i(a_i^*) > 0$ and $u_i(a_i^*, \alpha_i) < X$. So with respect to the probability distribution α_i, player i never obtains more than X, and has positive probability of obtaining strictly less than X. Similar to the other calculations we deduce $u_i(\alpha_i, \alpha_{-i}) < X$, i.e. α_i is not a best response, and we are done. $\square$

The Support Characterization exemplifies a theme in algorithmic game theory we will see again: we replace some condition on the player and their preferences with an equivalent combinatorial characterization. To practice using the Support Characterization, solve the following exercise.

Exercise. Consider the following game, the *Moose-Goose Hunt*. There are n players who are hunting a large moose. However, each one can either hunt the moose with the rest of the group, or choose to go off alone and hunt a goose instead. So $A_i = \{M, G\}$ for each player depending on what they choose to hunt. The utilities are given as follows, where $m > g > 0$ are fixed constants:

- any player who chooses G gets a utility of g
- if all players choose M, they all get a utility of m
- if player i chooses M, but at least one player chooses G, then player i gets 0 utility.

The idea is that hunting a moose is more profitable, but risky since it takes everyone to coordinate their efforts.

Find all *symmetric* mixed Nash equilibria of this game. (A mixed action profile is symmetric if each player assigns the same probability to M.)

Here is one consequence of the Support Characterization, and a warm-up to the Lemke-Howson algorithm:

Corollary 7. *Given a two-player strategic game with finite A_i and where the u_i are rational numbers, it has a mixed Nash equilibrium where all the probabilities are rational.*

Proof. Using Nash's theorem, we know that the game has at least one Nash equilibrium (α_1, α_2), but it might have irrational probabilities. We will use the following fact:

Fact 8. *If a linear inequality system is feasible, and all numbers A_{ij}, b_j in the program are rational, then the system has a solution which is rational.*

Continuing, for notational simplicity, we rename the actions of player 1 to $\{1, 2, \ldots, m\}$ and the actions of player 2 to $\{1, 2, \ldots, n\}$. Consider the following linear inequality system:

V there is a variable p_j for each action $j \in A_1$ of player 1 and another variable Y

C1 there are constraints $\forall j : p_j \geq 0$ and $\sum_j p_j = 1$ to ensure the p_j correspond to a probability distribution

C2 for all p_j in A_1 but not in $B_1(\alpha_2)$ there is a constraint $p_j = 0$

C3 for all k in $\mathtt{supp}(\alpha_2)$ there is a constraint $\sum_{i=1}^{m} p_j u_2(j, k) = Y$.

C4 for all k not in $\mathtt{supp}(\alpha_2)$ there is a constraint $\sum_{i=1}^{m} p_j u_2(j, k) \leq Y$.

The system is feasible: take each $p_j = \alpha_1(j)$ and let $Y = u_2(\alpha_1, \alpha_2)$ be player 2's utility in this equilibrium.

The key fact, more generally, is that whenever p is a feasible solution to this system, together with α_2 it makes a Nash equilibrium. To prove this we use the Support Characterization: the constraints C2 ensure player 1 only uses best responses to α_2; and the constraints C3, C4 ensure that each action used by player 2 is a best response to player 1's mixed strategy.

So by Fact 8, the system has a rational solution. Thus there is a mixed Nash equilibrium (α_1', α_2) where α_1' is rational. Repeating the argument for player 2, we are done. □

The Lemke-Howson algorithm will extend Corollary 7 to give a proof which is more satisfying, since it will not rely on Nash's theorem.

Exercise. Using the ideas in Corollary 7, give an (exponential-time) algorithm to find a Nash equilibrium of a 2-player strategic game.

Note that the rationality proven by Corollary 7 does not extend to 3-player games: the Moose-Goose hunt is an example.