Game Theory

2.1 Zero Sum Games (Part 2)

George Mason University, Spring 2018

The Nash equilibria of two-player, zero-sum games have various nice properties.

- **Minimax Condition** A pair of strategies is in equilibrium if (but not only if) the outcome determined by the strategies equals the minimal value of the row and the maximal value of the column.
- These solutions can be determined using the **maximinimizer** method.
- All maximinimizer strategy pairs are Nash equilibria.
- All maximinimizer strategy pairs have the same value, and this is the *value of the game*.

Let $min(a_j)$ designate the lowest utility obtainable by performing act $a_j \in \mathcal{A}_j$.

Def 2.1.4. A *maximinimizer* for player j is an action $a_j^* \in \mathcal{A}_j$ where $min(a_j^*) = max_{a_j \in \mathcal{A}_j}(min(a_j))$.

Thm 2.1.1. $\langle a_1^*, a_2^* \rangle$ is a Nash equilibrium of the two-player zero sum game $\mathcal{G}$ only if $a_1^* \in \mathcal{A}_1$ is a maximinimizer for player 1 and $a_2^* \in \mathcal{A}_2$ is a maximinimizer for player 2.

	a_1	a_2	a_3
a_1	8	8	7
a_2	0	-10	-4
a_3	9	0	-1

Recall that $\langle a_1, a_3 \rangle$ is the only Nash equilibrium of this game.

	a_1	a_2	a_3
a_1	8	8	7
a_2	0	-10	-4
a_3	9	0	-1

$a_1 \in \mathcal{A}_1$ is a maximinimizer for player 1.

	a_1	a_2	a_3
a_1	8	8	7
a_2	0	-10	-4
a_3	9	0	-1

$a_3 \in \mathcal{A}_2$ is a maximinimizer for player 2.

	a_1	a_2	a_3
a_1	0	8	3
a_2	0	1	10
a_3	-2	6	5

Recall that $\langle a_1, a_1 \rangle$ and $\langle a_2, a_1 \rangle$ are the Nash equilibria of this game.

	a_1	a_2	a_3
a_1	0	8	3
a_2	0	1	10
a_3	-2	6	5

$a_1 \in \mathcal{A}_1$ and $a_2 \in \mathcal{A}_1$ are maximinimizers for player 1.

	a_1	a_2	a_3
a_1	0	8	3
a_2	0	1	10
a_3	-2	6	5

$a_1 \in \mathcal{A}_2$ is a maximinimizer for player 2.

Proof of Thm 2.1.1. (If you're interested)

If $\langle a_1^*, a_2^* \rangle$ is a Nash equilibrium, then $u_2(g(\langle a_1^*, a_2^* \rangle)) \geq u_2(g(\langle a_1^*, a_2 \rangle))$ for each $a_2 \in \mathcal{A}_2$, so $u_1(g(\langle a_1^*, a_2^* \rangle)) \leq u_1(g(\langle a_1^*, a_2 \rangle))$ for each $a_2 \in \mathcal{A}_2$.

Hence, $min(a_1^*) = u_1(g(\langle a_1^*, a_2^* \rangle)) \leq max_{a_1 \in \mathcal{A}_1}(min(a_1))$.

If $\langle a_1^*, a_2^* \rangle$ is a Nash equilibrium, then $u_1(g(\langle a_1^*, a_2^* \rangle)) \geq u_1(g(\langle a_1, a_2^* \rangle))$ for each $a_1 \in \mathcal{A}_1$, so $u_1(g(\langle a_1^*, a_2^* \rangle)) \geq min(a_1)$ for each $a_1 \in \mathcal{A}_1$.

Hence, $u_1(g(\langle a_1^*, a_2^* \rangle)) \geq max_{a_1 \in \mathcal{A}_1}(min(a_1))$.

Thus, $min(a_1^*) = u_1(g(\langle a_1^*, a_2^* \rangle)) = max_{a_1 \in \mathcal{A}_1}(min(a_1))$. That is, a_1^* is a maximinimizer for player 1.

Similar reasoning establishes that
$min(a_2^*) = u_2(g(\langle a_1^*, a_2^* \rangle)) = max_{a_2 \in \mathcal{A}_2}(min(a_2))$.

Thm 2.1.2. If the two-player zero sum game $\mathcal{G}$ has a Nash equilibrium, then a_1^* and a_2^* are maximinimizers for players 1 and 2 respectively only if $\langle a_1^*, a_2^* \rangle$ is a Nash equilibrium of $\mathcal{G}$.

Proof of Thm 2.1.2. (If you're interested) Suppose that $\mathcal{G}$ has a Nash equilibrium.

Then from the Proof of Thm 2.1.1, we know that $max_{a_1 \in \mathcal{A}_1}(min(a_1)) = -max_{a_2 \in \mathcal{A}_2}(min(a_2))$.

Let $max_{a_1 \in \mathcal{A}_1}(min(a_1)) = v^*$. Then $max_{a_2 \in \mathcal{A}_2}(min(a_2)) = -v^*$.

Since a_1^* is a maximinimizer for 1, $u_1(g(\langle a_1^*, a_2 \rangle)) \geq v^*$ for all $a_2 \in \mathcal{A}_2$, so $u_1(g(\langle a_1^*, a_2^* \rangle)) \geq v^*$.

Since a_2^* is a maximinimizer for 2, $u_2(g(\langle a_1, a_2^* \rangle)) \geq -v^*$ for all $a_1 \in \mathcal{A}_1$, so $u_1(g(\langle a_1^*, a_2^* \rangle)) \leq v^*$.

Thus, $u_1(g(\langle a_1^*, a_2^* \rangle)) = v^*$ and $u_2(g(\langle a_1^*, a_2^* \rangle)) = -v^*$, so $\langle a_1^*, a_2^* \rangle$ is a Nash equilibrium of $\mathcal{G}$.

Thm 2.1.3. $\langle a_1^*, a_2^* \rangle$ and $\langle a_1^{**}, a_2^{**} \rangle$ are Nash equilibria of the two-player zero sum game $\mathcal{G}$ only if $u_1(g(\langle a_1^*, a_2^* \rangle)) = u_1(g(\langle a_1^{**}, a_2^{**} \rangle))$ (moreover, $u_2(g(\langle a_1^*, a_2^* \rangle)) = u_2(g(\langle a_1^{**}, a_2^{**} \rangle))$).

Less formally, all Nash equilibria of $\mathcal{G}$ have the same utilities. The equilibrium payoff v^* to player 1 is the *value* of the game.

Proof of Thm 2.1.3. (If you're interested) Suppose that $\langle a_1^*, a_2^* \rangle$ and $\langle a_1^{**}, a_2^{**} \rangle$ are Nash equilibria of $\mathcal{G}$.

From the Proof of Thm 2.1.1, we know that
$u_1(g(\langle a_1^*, a_2^* \rangle)) = u_1(g(\langle a_1^{**}, a_2^{**} \rangle)) = max_{a_1 \in \mathcal{A}_1}(min(a_1))$.

We also know that
$u_2(g(\langle a_1^*, a_2^* \rangle)) = u_2(g(\langle a_1^{**}, a_2^{**} \rangle)) = max_{a_2 \in \mathcal{A}_2}(min(a_2))$.

	a_1	a_2	a_3
a_1	8	8	7
a_2	0	-10	-4
a_3	9	0	-1

The value of this game is 7.

	a_1	a_2	a_3
a_1	0	8	3
a_2	0	1	10
a_3	-2	6	5

The value of this game is 0.

Thm 2.1.4 (Coordination Theorem for Zero Sum Games).
$\langle a_1^*, a_2^* \rangle$ and $\langle a_1^{**}, a_2^{**} \rangle$ are Nash equilibria of the two-player zero sum game $\mathcal{G}$ only if $\langle a_1^*, a_2^{**} \rangle$ and $\langle a_1^{**}, a_2^* \rangle$ are Nash equilibria of $\mathcal{G}$.

Proof of Thm 2.1.4. Suppose that $\langle a_1^*, a_2^* \rangle$ and $\langle a_1^{**}, a_2^{**} \rangle$ are Nash equilibria of $\mathcal{G}$.

By Thm 2.1.1, we know that $a_1^* \in \mathcal{A}_1$ and $a_1^{**} \in \mathcal{A}_1$ are maximinimizers for player 1, and $a_2^* \in \mathcal{A}_2$ and $a_2^{**} \in \mathcal{A}_2$ are maximinimizers for player 2.

Thus, by Thm 2.1.2, we know that both $\langle a_1^*, a_2^{**} \rangle$ and $\langle a_1^{**}, a_2^* \rangle$ are Nash equilibria of $\mathcal{G}$.

	a_1	a_2	a_3	a_4
a_1	1	2	3	1
a_2	0	5	0	0
a_3	1	6	4	1

$\langle a_1, a_1 \rangle$ and $\langle a_3, a_4 \rangle$ are Nash equilbria.

	a_1	a_2	a_3	a_4
a_1	1	2	3	1
a_2	0	5	0	0
a_3	1	6	4	1

So $\langle a_1, a_4 \rangle$ and $\langle a_3, a_1 \rangle$ are too.

Unfortunately, not every two-player zero sum game has an action profile in equilibrium.

	paper	rock	scissors
paper	0,0	1,-1	-1,1
rock	-1,1	0,0	1,-1
scissors	1,-1	-1,1	0,0

$$max_{a_1 \in \mathcal{A}_1}(min(a_1)) = -1$$

$$max_{a_2 \in \mathcal{A}_2}(min(a_2)) = -1$$

$$max_{a_1 \in \mathcal{A}_1}(min(a_1)) \neq -max_{a_2 \in \mathcal{A}_2}(min(a_2)).$$

Fortunately, once we allow for *mixed strategies*, every two-player zero sum game has a pair of pure or mixed strategies in Nash equilibirum.

A player's choices can now be *nondeterministic*.

		$[\frac{1}{4}]$ paper	$[\frac{1}{2}]$ rock	$[\frac{1}{4}]$ scissors
$[\frac{1}{3}]$	paper	0	1	-1
$[\frac{1}{3}]$	rock	-1	0	1
$[\frac{1}{3}]$	scissors	1	-1	0

Player 1's mixed strategy is to play paper with probability $\frac{1}{3}$, rock with probability $\frac{1}{3}$, and scissors with probability $\frac{1}{3}$.

Player 2's mixed strategy is to play paper with probability $\frac{1}{4}$, rock with probability $\frac{1}{2}$, and scissors with probability $\frac{1}{4}$.

A player's choices can now be *nondeterministic*.

		$[\frac{1}{4}]$	$[\frac{1}{2}]$	$[\frac{1}{4}]$
		paper	rock	scissors
$[\frac{1}{3}]$	paper	0	1	-1
$[\frac{1}{3}]$	rock	-1	0	1
$[\frac{1}{3}]$	scissors	1	-1	0

Viewing mixed strategies naïvely, we can think of player 1 and player 2 as committing themselves to randomized chance mechanisms that select actions with various probabilities.

Each $a_i \in \mathcal{A}_i$ is a *pure strategy* of player i.

Let $\triangle(\mathcal{A}_i)$ designate the set of probability measures over $\mathcal{A}_i$.

Each $\alpha_i \in \triangle(\mathcal{A}_i)$ is a *mixed strategy* of player i.

$\alpha_i(a_i) = p_i$ iff player i chooses $a_i \in \mathcal{A}_i$ with probability p_i.

The mixed strategy α_i for player $i \in \mathcal{N}$ with $|\mathcal{A}_i| = n$ where $a_1 \in \mathcal{A}_i$ is chosen with probability p_1, $a_2 \in \mathcal{A}_i$ is chosen with probability p_2, ..., and $a_n \in \mathcal{A}_i$ is chosen with probability p_n can be written as $\langle a_1[p_1], ..., a_n[p_n]\rangle$.

$\langle \text{paper}[\frac{1}{3}], \text{rock}[\frac{1}{3}], \text{scissors}[\frac{1}{3}]\rangle \in \triangle(\mathcal{A}_1)$.

$\langle \text{paper}[\frac{1}{4}], \text{rock}[\frac{1}{2}], \text{scissors}[\frac{1}{4}]\rangle \in \triangle(\mathcal{A}_2)$.

Note that pure strategies can be regarded as special cases of mixed strategies. For instance, rock $= \langle \text{paper}[0], \text{rock}[1], \text{scissors}[0]\rangle$.

Before we worked with action profiles in $\times_{i \in \mathcal{N}} \mathcal{A}_i$.

We will now work with mixed strategy profiles in $\times_{i \in \mathcal{N}} \triangle(\mathcal{A}_i)$.

For instance,
$\langle\langle \text{paper}[\frac{1}{3}], \text{rock}[\frac{1}{3}], \text{scissors}[\frac{1}{3}]\rangle, \langle \text{paper}[\frac{1}{4}], \text{rock}[\frac{1}{2}], \text{scissors}[\frac{1}{4}]\rangle\rangle$ is a mixed strategy profile in $\triangle(\mathcal{A}_1) \times \triangle(\mathcal{A}_2)$.

What is the expected utility of a mixed strategy profile $\langle \alpha_1, ..., \alpha_{|\mathcal{N}|} \rangle$ for each player $i \in \mathcal{N}$?

Note that each mixed strategy profile induces a probability distribution over the set $\times_{i \in \mathcal{N}} \mathcal{A}_i$ of action profiles.

	paper	rock	scissors
paper	0 [$\frac{1}{12}$]	1 [$\frac{1}{6}$]	-1 [$\frac{1}{12}$]
rock	-1 [$\frac{1}{12}$]	0 [$\frac{1}{6}$]	1 [$\frac{1}{12}$]
scissors	1 [$\frac{1}{12}$]	-1 [$\frac{1}{6}$]	0 [$\frac{1}{12}$]

In general, the probability of pure strategy profile $\langle a_1, ..., a_{|\mathcal{N}|} \rangle \in \times_{i \in \mathcal{N}} \mathcal{A}_i$ is $\Pi_{i \in \mathcal{N}} \alpha_i(a_i)$.

The expected utility of mixed strategy profile $\langle \alpha_1, ..., \alpha_{|\mathcal{N}|} \rangle \in \times_{i \in \mathcal{N}} \alpha_i$ for player $i \in \mathcal{N}$ is $\Sigma_{\langle a_1, ..., a_{|\mathcal{N}|} \rangle \in \times_{i \in \mathcal{N}} \mathcal{A}_i} \Pi_{i \in \mathcal{N}} \alpha_i(a_i) u_i(g(\langle a_1, ..., a_{|\mathcal{N}|} \rangle))$.

What is the expected utility of a mixed strategy profile $\langle \alpha_1, ..., \alpha_{|\mathcal{N}|} \rangle$ for each player $i \in \mathcal{N}$?

Note that each mixed strategy profile induces a probability distribution over the set $\times_{i \in \mathcal{N}} \mathcal{A}_i$ of action profiles.

	paper	rock	scissors
paper	0 $[\frac{1}{12}]$	1 $[\frac{1}{6}]$	-1 $[\frac{1}{12}]$
rock	-1 $[\frac{1}{12}]$	0 $[\frac{1}{6}]$	1 $[\frac{1}{12}]$
scissors	1 $[\frac{1}{12}]$	-1 $[\frac{1}{6}]$	0 $[\frac{1}{12}]$

For example,
$EU_1(\langle\langle \text{paper}[\frac{1}{3}], \text{rock}[\frac{1}{3}], \text{scissors}[\frac{1}{3}]\rangle, \langle \text{paper}[\frac{1}{4}], \text{rock}[\frac{1}{2}], \text{scissors}[\frac{1}{4}]\rangle\rangle) =$
$0 \times \frac{1}{12} + 1 \times \frac{1}{6} - 1 \times \frac{1}{12} - 1 \times \frac{1}{12} + 0 \times \frac{1}{6} + 1 \times \frac{1}{12} + 1 \times \frac{1}{12} - 1 \times \frac{1}{6} + 0 \times \frac{1}{12} = 0.$

What is the expected utility of a mixed strategy profile $\langle \alpha_1, ..., \alpha_{|\mathcal{N}|} \rangle$ for each player $i \in \mathcal{N}$?

Note that each mixed strategy profile induces a probability distribution over the set $\times_{i \in \mathcal{N}} \mathcal{A}_i$ of action profiles.

	paper	rock	scissors
paper	0 $[\frac{1}{12}]$	1 $[\frac{1}{6}]$	-1 $[\frac{1}{12}]$
rock	-1 $[\frac{1}{12}]$	0 $[\frac{1}{6}]$	1 $[\frac{1}{12}]$
scissors	1 $[\frac{1}{12}]$	-1 $[\frac{1}{6}]$	0 $[\frac{1}{12}]$

For example,
$EU_2(\langle\langle \text{paper}[\frac{1}{3}], \text{rock}[\frac{1}{3}], \text{scissors}[\frac{1}{3}]\rangle, \langle \text{paper}[\frac{1}{4}], \text{rock}[\frac{1}{2}], \text{scissors}[\frac{1}{4}]\rangle\rangle) =$
$0 \times \frac{1}{12} - 1 \times \frac{1}{6} + 1 \times \frac{1}{12} + 1 \times \frac{1}{12} + 0 \times \frac{1}{6} - 1 \times \frac{1}{12} - 1 \times \frac{1}{12} + 1 \times \frac{1}{6} + 0 \times \frac{1}{12} = 0.$

We can now extend the concept of Nash equilibrium to cover mixed strategies.

Def 2.1.6. A *mixed strategy Nash equilibrium* of $\mathcal{G}$ is a mixed strategy profile $\alpha^* \in \times_{i \in \mathcal{N}} \alpha_i$ such that for every player $j \in \mathcal{N}$, the following condition holds:

$$EU_j(\langle \alpha_j^*, \alpha_{-j}^* \rangle) \geq EU_j(\langle \alpha_j, \alpha_{-j}^* \rangle) \text{ for each } \alpha_j \in \triangle(\mathcal{A}_j).$$

In other words, given the other players' equilibrium mixed strategy profile α_{-j}^*, the equilibrium mixed strategy α_j^* of player j is optimal.

Osborne and Rubinstein: "No player can profitably deviate, given the actions of the other players."

Pure strategy Nash equilibria can be regarded as special cases of mixed strategy Nash equilibria.

		$[\frac{1}{4}]$	$[\frac{1}{2}]$	$[\frac{1}{4}]$
		paper	rock	scissors
$[\frac{1}{3}]$	paper	0	1	-1
$[\frac{1}{3}]$	rock	-1	0	1
$[\frac{1}{3}]$	scissors	1	-1	0

Is $\langle\langle \text{paper}[\frac{1}{3}], \text{rock}[\frac{1}{3}], \text{scissors}[\frac{1}{3}]\rangle, \langle \text{paper}[\frac{1}{4}], \text{rock}[\frac{1}{2}], \text{scissors}[\frac{1}{4}]\rangle\rangle$ a mixed strategy Nash equilibrium of Rock, Paper, & Scissors?

No. Player 1 does well to play $\langle \text{paper}[1], \text{rock}[0], \text{scissors}[0]\rangle$ instead.

$EU_1(\langle\langle \text{paper}[\frac{1}{3}], \text{rock}[\frac{1}{3}], \text{scissors}[\frac{1}{3}]\rangle, \langle \text{paper}[\frac{1}{4}], \text{rock}[\frac{1}{2}], \text{scissors}[\frac{1}{4}]\rangle\rangle) = 0.$

$EU_1(\langle\langle \text{paper}[1], \text{rock}[0], \text{scissors}[0]\rangle, \langle \text{paper}[\frac{1}{4}], \text{rock}[\frac{1}{2}], \text{scissors}[\frac{1}{4}]\rangle\rangle) = \frac{1}{4}.$

		$[\frac{1}{3}]$ paper	$[\frac{1}{3}]$ rock	$[\frac{1}{3}]$ scissors
$[\frac{1}{3}]$	paper	0	1	-1
$[\frac{1}{3}]$	rock	-1	0	1
$[\frac{1}{3}]$	scissors	1	-1	0

Is $\langle\langle \text{paper}[\frac{1}{3}], \text{rock}[\frac{1}{3}], \text{scissors}[\frac{1}{3}]\rangle, \langle \text{paper}[\frac{1}{3}], \text{rock}[\frac{1}{3}], \text{scissors}[\frac{1}{3}]\rangle\rangle$ a mixed strategy Nash equilibrium of Rock, Paper, & Scissors?

Yes. If either player plays the equiprobable mixed strategy, the expected utilities of both players are 0.

		$[\frac{1}{2}]$	$[\frac{1}{2}]$
		a_1	a_2
$[\frac{1}{2}]$	a_1	6	3
$[\frac{1}{2}]$	a_2	2	4

Is $\langle\langle a_1[\frac{1}{2}], a_2[\frac{1}{2}]\rangle, \langle a_1[\frac{1}{2}], a_2[\frac{1}{2}]\rangle\rangle$ a mixed strategy Nash equilibrium?

No. $EU_1(\langle\langle a_1[\frac{1}{2}], a_2[\frac{1}{2}]\rangle, \langle a_1[\frac{1}{2}], a_2[\frac{1}{2}]\rangle\rangle) = \frac{15}{4}$.

But $EU_1(\langle\langle a_1[1], a_2[0]\rangle, \langle a_1[\frac{1}{2}], a_2[\frac{1}{2}]\rangle\rangle) = \frac{18}{4}$.

Does this game even have any mixed strategy Nash equilibria?

Thm 2.1.5 (Maximin Theorem for Two-player Zero Sum Games). Every two-person zero sum game has at least one mixed strategy Nash equilibrium. Moreover, the expected utilities of each Nash equilibrium mixed strategy profile are the same.

Partial Proof of Thm 2.1.5. We show that any 2x2 zero sum game in *standard form* has a mixed strategy Nash equilibrium.

	a_1	a_2
a_1	a	b
a_2	c	d

$a, b, c, d \in \mathbb{R}[0, \infty)$ where $a > c$, $d > b$, and $d > c$.

Note that this game does not have any pure strategy Nash equilibria.

		$[q]$	$[1-q]$
		a_1	a_2
$[p]$	a_1	a	b
$[1-p]$	a_2	c	d

$$EU_1(\langle\langle a_1[p], a_2[1-p]\rangle, \langle a_1[q], a_2[1-q]\rangle\rangle) =$$
$$apq + bp(1-q) + c(1-p)q + d(1-p)(1-q) =$$
$$apq + bp - bpq + cq - cpq + d - dp - dq + dpq =$$

$$(a - c + d - b)pq - (d - b)p - (d - c)q + d =$$

$$Apq - Bp - Cq + D =$$
$$A((p - \tfrac{C}{A})(q - \tfrac{B}{A})) + \tfrac{DA-BC}{A}$$

where
$A = a - b + d - c > 0,\ B = d - b > 0,\ C = d - c > 0,\ D = d > 0.$

$EU_1(\langle\langle a_1[p], a_2[1-p]\rangle, \langle a_1[q], a_2[1-q]\rangle\rangle) = A((p-\frac{C}{A})(q-\frac{B}{A})) + \frac{DA-BC}{A}$.

$EU_2(\langle\langle a_1[p], a_2[1-p]\rangle, \langle a_1[q], a_2[1-q]\rangle\rangle) = -A((p-\frac{C}{A})(q-\frac{B}{A})) - \frac{DA-BC}{A}$.

**Player 1 can prevent EU_1 from falling below $\frac{DA-BC}{A}$ by setting $p = \frac{C}{A}$.

**Player 2 can prevent EU_2 from falling below $-\frac{DA-BC}{A}$ by setting $q = \frac{B}{A}$.

$\langle\langle a_1[\frac{C}{A}], a_2[1-\frac{C}{A}]\rangle, \langle a_1[\frac{B}{A}], a_2[1-\frac{B}{A}]\rangle\rangle$ is a mixed strategy equilibrium.

The value of the game is $\frac{DA-BC}{A}$.

	a_1	a_2
a_1	6	3
a_2	2	4

What is a mixed strategy Nash equilibrium of this game?

	a_1	a_2
a_1	6	3
a_2	2	4

$$A = a - b + d - c =$$
$$B = d - b =$$
$$C = d - c =$$
$$D = d =$$

$$p = \frac{C}{A} =$$
$$q = \frac{B}{A} =$$

$$EU_1(\langle\langle a_1[p], a_2[1-p]\rangle, \langle a_1[q], a_2[1-q]\rangle\rangle) =$$

	a_1	a_2
a_1	6	3
a_2	2	4

$$A = a - b + d - c = 5$$
$$B = d - b = 1$$
$$C = d - c = 2$$
$$D = d = 4$$

$$p = \frac{C}{A} = \frac{2}{5}$$
$$q = \frac{B}{A} = \frac{1}{5}$$

$$EU_1(\langle\langle a_1[\tfrac{2}{5}], a_2[\tfrac{3}{5}]\rangle, \langle a_1[\tfrac{1}{5}], a_2[\tfrac{4}{5}]\rangle\rangle) = \frac{2}{25} \times 6 + \frac{8}{25} \times 3 + \frac{3}{25} \times 2 + \frac{12}{25} \times 4 = \frac{90}{25}$$

$$EU_1(\langle\langle a_1[\tfrac{2}{5}], a_2[\tfrac{3}{5}]\rangle, \langle a_1[\tfrac{1}{5}], a_2[\tfrac{4}{5}]\rangle\rangle) = \frac{DA - BC}{A} = \frac{18}{5}$$

	a_1	a_2
a_1	-3	1
a_2	2	0

What is a mixed strategy Nash equilibrium of this game?

	a_1	a_2
a_1	0	4
a_2	5	3

First translation: add the constant $+3$.

	a_1	a_2
a_2	5	3
a_1	0	4

Second translation: switch rows.

	a_1	a_2
a_2	5	3
a_1	0	4

$$A = a - b + d - c =$$
$$B = d - b =$$
$$C = d - c =$$
$$D = d =$$
$$p = \frac{C}{A} =$$
$$q = \frac{B}{A} =$$
$$EU_1(\langle\langle a_1[p], a_2[1-p]\rangle, \langle a_1[q], a_2[1-q]\rangle\rangle) = \frac{DA-BC}{A} =$$

	a_1	a_2
a_2	5	3
a_1	0	4

$$A = a - b + d - c = 6$$
$$B = d - b = 1$$
$$C = d - c = 4$$
$$D = d = 4$$

$$p = \frac{C}{A} = \frac{2}{3}$$
$$q = \frac{B}{A} = \frac{1}{6}$$

$$EU_1(\langle\langle a_1[\tfrac{1}{3}], a_2[\tfrac{2}{3}]\rangle, \langle a_1[\tfrac{1}{6}], a_2[\tfrac{5}{6}]\rangle\rangle) = \frac{DA - BC}{A} = \frac{20}{6}$$

Before we're finished, we need to transform it back to the original utilities!

	a_1	a_2
a_1	-3	1
a_2	2	0

$$EU_1(\langle\langle a_1[\tfrac{1}{3}], a_2[\tfrac{2}{3}]\rangle, \langle a_1[\tfrac{1}{6}], a_2[\tfrac{5}{6}]\rangle\rangle) = \tfrac{20}{6} - 3 = \tfrac{1}{3}$$

And now for a shortcut.

We now have a proof that every two-by-two, two-person, zero-sum game has at least one pure or mixed-strategy Nash equilibrium. So we can construct a recipe for finding them:

1. First, use the maximin strategy to determine if the game has any pure strategy Nash equilibria. If it does, that's the *solution* to the game. The value of those equilibria is the *value* of the game.

2. If this strategy comes up empty, then, assume the game has the following structure:

		$[q]$	$[1-q]$
		C_1	C_2
$[p]$	R_1	a	b
$[1-p]$	R_2	c	d

Unlike the *standard form* game needed for the proof, there are no restrictions on a, b, c, and d in this strategy.

Our aim is to find values for p and q for which
$(pR_1,\ (1-p)R_2);\ (qC_1,\ (1-q)C_2)$

It turns out that:

$p = (d - c)/[(a + d) - (b + c)]$

and

$q = (d - b)/[(a + d) - (b + c)]$

The *value* of the game is $ap + c(1 - p)$.

Why does this work? When Row plays her half of an equilibrium mixed-strategy pair, she fixes the value of the game no matter what Col does. And the same goes for Col. If this is true, the EU of Row's equilibrium strategy against C_1 must be the same as the EU of that strategy against C_2, and the same holds for Col.

The EU for Row against C_1 is $ap + (1 - p)c$ and against C_2 is $bp + (1 - p)d$.

Since they are equal, we have

$ap + (1 - p)c = bp + (1 - p)d$.

Solving for p gives us the equation above, and we can do the same thing for q.

	a_1	a_2
a_1	-3	1
a_2	2	0

What is a mixed strategy Nash equilibrium of this game?

	a_1	a_2
a_1	-3	1
a_2	2	0

$p = (d - c)/[(a + d) - (b + c)] =$

and

$q = (d - b)/[(a + d) - (b + c)] =$

The *value* of the game is $ap + c(1 - p) =$

	a_1	a_2
a_1	-3	1
a_2	2	0

$p = (d - c)/[(a + d) - (b + c)] = \frac{(0-2)}{(-3+0)-(1+2)} = \frac{-2}{-6} = \frac{1}{3}$

and

$q = (d - b)/[(a + d) - (b + c)] = \frac{0-1}{(-3+0)-(1+2)} = \frac{-1}{-6} = \frac{1}{6}$

The *value* of the game is $ap + c(1 - p) = -3(\frac{1}{3}) + 2(1 - \frac{1}{3}) = \frac{1}{3}$