Static Games of Complete Information

Mixed Strategies

Let's Mix It

As pointed out before, neither of the solution concepts has to exist in pure strategies

Example: Rock-Paper-sCissors

	R	P	C
R	$0, 0$	$-1, 1$	$1, -1$
P	$1, -1$	$0, 0$	$-1, 1$
C	$-1, 1$	$1, -1$	$0, 0$

There are no strictly dominant pure strategies

No strategy is strictly dominated (IESDS removes nothing)

Each strategy is a best response to some strategy of the opponent (rationalizability removes nothing)

No pure Nash equilibria: No *pure* strategy profile allows each player to play a best response to the strategy of the other player

How to solve this?

Let the players randomize their choice of pure strategies ....

Probability Distributions

Definition 19

Let A be a finite set. A *probability distribution over A* is a function $\sigma : A \to [0,1]$ such that $\sum_{a \in A} \sigma(a) = 1$.

We denote by $\Delta(A)$ the set of all probability distributions over A.

We denote by $supp(\sigma)$ the *support* of σ, that is the set of all $a \in A$ satisfying $\sigma(a) > 0$.

Example 20

Consider $A = \{a, b, c\}$ and a function $\sigma : A \to [0,1]$ such that $\sigma(a) = \frac{1}{4}$, $\sigma(b) = \frac{3}{4}$, and $\sigma(c) = 0$. Then $\sigma \in \Delta(A)$ and $supp(\sigma) = \{a, b\}$.

Mixed Strategies

Let us fix a strategic-form game $G = (N, (S_i)_{i \in N}, (u_i)_{i \in N})$.

From now on, **assume that all S_i are finite!**

Definition 21

A *mixed strategy* of player i is a probability distribution $\sigma \in \Delta(S_i)$ over S_i. We denote by $\Sigma_i = \Delta(S_i)$ the set of all mixed strategies of player i.

We define $\Sigma := \Sigma_1 \times \cdots \times \Sigma_n$, the set of all *mixed strategy profiles*.

Recall that by Σ_{-i} we denote the set $\Sigma_1 \times \cdots \Sigma_{i-1} \times \Sigma_{i+1} \times \cdots \times \Sigma_n$

Elements of Σ_{-i} are denoted by $\sigma_{-i} = (\sigma_1, \ldots, \sigma_{i-1}, \sigma_{i+1}, \ldots, \sigma_n)$.

We identify each $s_i \in S_i$ with a mixed strategy σ that assigns probability one to s_i (and zero to other pure strategies).

For example, in rock-paper-scissors, the pure strategy R corresponds to σ_i which satisfies $\sigma_i(X) = \begin{cases} 1 & X = R \\ 0 & \text{otherwise} \end{cases}$

Mixed Strategies

Sometimes we assume $S_i = \{1, \ldots, m_i\}$, here $m_i \in \{1, 2, \ldots\}$, for all $i \in N$.

Then every mixed strategy σ_i is a vector $\sigma_i = (\sigma_i(1), \ldots, \sigma_i(m_i))^\top \in [0, 1]^{m_i}$ so that

$$\sigma_i(1) + \cdots + \sigma_i(m_i) = 1$$

Mixed Strategy Profiles

Let $\sigma = (\sigma_1, \ldots, \sigma_n)$ be a mixed strategy profile.

Intuitively, we assume that each player i *randomly* chooses his pure strategy according to σ_i and *independently* of his opponents.

Thus for $s = (s_1, \ldots, s_n) \in S = S_1 \times \cdots \times S_n$ we have that

$$\sigma(s) := \prod_{i=1}^{n} \sigma_i(s_i)$$

is the probability that the players choose the pure strategy profile s according to the mixed strategy profile σ, and

$$\sigma_{-i}(s_{-i}) := \prod_{k \neq i}^{n} \sigma_k(s_k)$$

is the probability that the opponents of player i choose $s_{-i} \in S_{-i}$ when they play according to the mixed strategy profile $\sigma_{-i} \in \Sigma_{-i}$.

(We abuse notation a bit here: σ denotes two things, a vector of mixed strategies as well as a probability distribution on S (the same for σ_{-i})

Mixed Strategies – Example

	R	P	C
R	$0,0$	$-1,1$	$1,-1$
P	$1,-1$	$0,0$	$-1,1$
C	$-1,1$	$1,-1$	$0,0$

An example of a mixed strategy σ_1: $\sigma_1(R) = \frac{1}{2}$, $\sigma_1(P) = \frac{1}{3}$, $\sigma_1(C) = \frac{1}{6}$.

Sometimes we write σ_1 as $(\frac{1}{2}(R), \frac{1}{3}(P), \frac{1}{6}(C))$, or only $(\frac{1}{2}, \frac{1}{3}, \frac{1}{6})$ if the order of pure strategies is fixed.

Consider a mixed strategy profile (σ_1, σ_2) where
$\sigma_1 = (\frac{1}{2}(R), \frac{1}{3}(P), \frac{1}{6}(C))$ and $\sigma_2 = (\frac{1}{3}(R), \frac{2}{3}(P), 0(C))$.
Then the probability $\sigma(R, P)$ that the pure strategy profile (R, P) will be chosen by players playing the mixed profile (σ_1, σ_2) is

$$\sigma_1(R) \cdot \sigma_2(P) = \frac{1}{2} \cdot \frac{2}{3} = \frac{1}{3}$$

Expected Payoff

... but now what is the suitable notion of payoff?

Definition 22
The *expected payoff* of player i under a mixed strategy profile $\sigma \in \Sigma$ is

$$u_i(\sigma) := \sum_{s \in S} \sigma(s) u_i(s) \qquad \left(= \sum_{s \in S} \prod_{k=1}^{n} \sigma_k(s_k) u_i(s) \right)$$

I.e., it is the "weighted average" of what player i wins under each pure strategy profile s, weighted by the probability of that profile.

Assumption: Every rational player strives to maximize his own expected payoff.
(This assumption is not always completely convincing ...)

Expected Payoff – Example

Matching Pennies:

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

Each player secretly turns a penny to heads or tails, and then they reveal their choices simultaneously. If the pennies match, player 1 (row) wins, if they do not match, player 2 (column) wins.

Consider $\sigma_1 = (\frac{1}{3}(H), \frac{2}{3}(T))$ and $\sigma_2 = (\frac{1}{4}(H), \frac{3}{4}(T))$

$$u_1(\sigma_1, \sigma_2) = \sum_{(X,Y)\in\{H,T\}^2} \sigma_1(X)\sigma_2(Y)u_1(X, Y)$$

$$= \frac{1}{3}\frac{1}{4}1 + \frac{1}{3}\frac{3}{4}(-1) + \frac{2}{3}\frac{1}{4}(-1) + \frac{2}{3}\frac{3}{4}1 = \frac{1}{6}$$

$$u_2(\sigma_1, \sigma_2) = \sum_{(X,Y)\in\{H,T\}^2} \sigma_1(X)\sigma_2(Y)u_2(X, Y)$$

$$= \frac{1}{3}\frac{1}{4}(-1) + \frac{1}{3}\frac{3}{4}1 + \frac{2}{3}\frac{1}{4}1 + \frac{2}{3}\frac{3}{4}(-1) = -\frac{1}{6}$$

"Decomposition" of Expected Payoff

Consider the matching pennies example from the previous slide:

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

together with some mixed strategies σ_1 and σ_2.

We prove the following important property of the expected payoff:

$$u_1(\sigma_1,\sigma_2) = \sum_{X\in\{H,T\}} \sigma_1(X)u_1(X,\sigma_2)$$

An intuition behind this equality is following:

- $u_1(\sigma_1,\sigma_2)$ is the expected payoff of player 1 in the following experiment: Both players simultaneously and independently choose their pure strategies X, Y according to σ_1, σ_2, resp., and then player 1 collects his payoff $u_1(X, Y)$.
- $\sum_{X\in\{H,T\}} \sigma_1(X)u_1(X,\sigma_2)$ is the expected payoff of player 1 in the following: Player 1 chooses his *pure* strategy X and then uses it against the mixed strategy σ_2 of player 2. Then player 2 chooses Y according to σ_2 *independently of* X, and player 1 collects the payoff $u_1(X, Y)$.

As Y does not depend on X in neither experiment, we obtain the above equality of expected payoffs.

"Decomposition" of Expected Payoff

Consider the matching pennies example from the previous slide:

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

together with some mixed strategies σ_1 and σ_2.

A formal proof is straightforward:

$$\begin{aligned} u_1(\sigma_1,\sigma_2) &= \sum_{(X,Y)\in\{H,T\}^2} \sigma_1(X)\sigma_2(Y)u_1(X,Y) \\ &= \sum_{X\in\{H,T\}} \sum_{Y\in\{H,T\}} \sigma_1(X)\sigma_2(Y)u_1(X,Y) \\ &= \sum_{X\in\{H,T\}} \sigma_1(X) \sum_{Y\in\{H,T\}} \sigma_2(Y)u_1(X,Y) \\ &= \sum_{X\in\{H,T\}} \sigma_1(X)u_1(X,\sigma_2) \end{aligned}$$

(In the last equality we used the fact that X is identified with a mixed strategy assigning one to X.)

"Decomposition" of Expected Payoff

Consider the matching pennies example from the previous slide:

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

together with some mixed strategies σ_1 and σ_2.

Similarly,

$$\begin{aligned} u_1(\sigma_1,\sigma_2) &= \sum_{(X,Y)\in\{H,T\}^2} \sigma_1(X)\sigma_2(Y)u_1(X,Y) \\ &= \sum_{X\in\{H,T\}}\sum_{Y\in\{H,T\}} \sigma_1(X)\sigma_2(Y)u_1(X,Y) \\ &= \sum_{Y\in\{H,T\}}\sum_{X\in\{H,T\}} \sigma_1(X)\sigma_2(Y)u_1(X,Y) \\ &= \sum_{Y\in\{H,T\}} \sigma_2(Y) \sum_{X\in\{H,T\}} \sigma_1(X)u_1(X,Y) \\ &= \sum_{Y\in\{H,T\}} \sigma_2(Y)u_1(\sigma_1,Y) \end{aligned}$$

Expected Payoff – "Decomposition" in General

Lemma 23
For every mixed strategy profile $\sigma \in \Sigma$ and every $k \in N$ we have

$$u_i(\sigma) = \sum_{s_k \in S_k} \sigma_k(s_k) \cdot u_i(s_k, \sigma_{-k}) = \sum_{s_{-k} \in S_{-k}} \sigma_{-k}(s_{-k}) \cdot u_i(\sigma_k, s_{-k})$$

Lemma 23 immediately implies that

- each $u_i(\sigma)$ is affine in each $\sigma_k(s_k)$,
- Also, $u_i(\sigma) = u_i(\sigma_1, \dots, \sigma_n)$ is linear in each σ_k.
  Indeed, assuming w.l.o.g. that $S_k = \{1, \dots, m_k\}$,

$$u_i(\sigma) = \sum_{s_k \in S_k} \sigma_k(s_k) \cdot u_i(s_k, \sigma_{-k}) = \sum_{\ell=1}^{m_k} \sigma_k(\ell) \cdot u_i(\ell, \sigma_{-k})$$

  is the scalar product of the vector $\sigma_k = (\sigma_k(1), \dots, \sigma_k(m_k))$ with the vector $(u_i(1, \sigma_{-k}), \dots, u_i(m_k, \sigma_{-k}))$.

Expected Payoff – Pure Strategy Bounds

Before proving Lemma 23, we prove the following simple corollary.

Corollary 24

For all $i, k \in N$ and $\sigma \in \Sigma$ we have that

- $\min_{s_k \in S_k} u_i(s_k, \sigma_{-k}) \leq u_i(\sigma) \leq \max_{s_k \in S_k} u_i(s_k, \sigma_{-k})$
- $\min_{s_{-k} \in S_{-k}} u_i(\sigma_k, s_{-k}) \leq u_i(\sigma) \leq \max_{s_{-k} \in S_{-k}} u_i(\sigma_k, s_{-k})$

Proof.

We prove $u_i(\sigma) \leq \max_{s_k \in S_k} u_i(s_k, \sigma_{-k})$ the rest is similar. Define $B := \max_{s_k \in S_k} u_i(s_k, \sigma_{-k})$. Then

$$
\begin{aligned}
u_i(\sigma) &= \sum_{s_k \in S_k} \sigma_k(s_k) \cdot u_i(s_k, \sigma_{-k}) \\
&= \sum_{s_k \in S_k} \sigma_k(s_k) \cdot (B - (B - u_i(s_k, \sigma_{-k}))) \\
&\leq \sum_{s_k \in S_k} \sigma_k(s_k) \cdot B \\
&= B
\end{aligned}
$$

Proof of Lemma 23

$$
\begin{aligned}
u_i(\sigma) = \sum_{s \in S} \sigma(s) u_i(s) &= \sum_{s \in S} \prod_{\ell=1}^{n} \sigma_\ell(s_\ell) u_i(s) \\
&= \sum_{s \in S} \sigma_k(s_k) \prod_{\ell \neq k}^{n} \sigma_\ell(s_\ell) u_i(s) \\
&= \sum_{s_k \in S_k} \sum_{s_{-k} \in S_{-k}} \sigma_k(s_k) \prod_{\ell \neq k}^{n} \sigma_\ell(s_\ell) u_i(s_k, s_{-k}) \\
&= \sum_{s_k \in S_k} \sum_{s_{-k} \in S_{-k}} \sigma_k(s_k) \sigma_{-k}(s_{-k}) u_i(s_k, s_{-k})
\end{aligned}
$$

Proof of Lemma 23 (cont.)

The first equality:

$$
\begin{aligned}
u_i(\sigma) &= \sum_{s_k \in S_k} \sum_{s_{-k} \in S_{-k}} \sigma_k(s_k)\sigma_{-k}(s_{-k})u_i(s_k, s_{-k}) \\
&= \sum_{s_k \in S_k} \sigma_k(s_k) \sum_{s_{-k} \in S_{-k}} \sigma_{-k}(s_{-k})u_i(s_k, s_{-k}) \\
&= \sum_{s_{-k} \in S_{-k}} \sigma_k(s_k)u_i(s_k, \sigma_{-k})
\end{aligned}
$$

Proof of Lemma 23 (cont.)

The second equality:

$$
\begin{aligned}
u_i(\sigma) &= \sum_{s_k \in S_k} \sum_{s_{-k} \in S_{-k}} \sigma_k(s_k)\sigma_{-k}(s_{-k})u_i(s_k, s_{-k}) \\
&= \sum_{s_{-k} \in S_{-k}} \sum_{s_k \in S_k} \sigma_k(s_k)\sigma_{-k}(s_{-k})u_i(s_k, s_{-k}) \\
&= \sum_{s_{-k} \in S_{-k}} \sigma_{-k}(s_{-k}) \sum_{s_k \in S_k} \sigma_k(s_k)u_i(s_k, s_{-k}) \\
&= \sum_{s_{-k} \in S_{-k}} \sigma_{-k}(s_{-k})u_i(\sigma_k, s_{-k})
\end{aligned}
$$

Solution Concepts

We revisit the following solution concepts in mixed strategies:

- strict dominant strategy equilibrium
- IESDS equilibrium
- rationalizable equilibria
- Nash equilibria

From now on, when I say a *strategy* I implicitly mean a

mixed strategy.

In order to deal with efficiency issues we assume that the size of the game G is defined by $|G| := |N| + \sum_{i \in N} |S_i| + \sum_{i \in N} |u_i|$ where $|u_i| = \sum_{s \in S} |u_i(s)|$ and $|u_i(s)|$ is the length of a binary encoding of $u_i(s)$ (we assume that rational numbers are encoded as quotients of two binary integers)

Note that, in particular, $|G| > |S|$.

Strict Dominance in Mixed Strategies

Definition 25

Let $\sigma_i, \sigma'_i \in \Sigma_i$ be (mixed) strategies of player i. Then σ'_i is *strictly dominated* by σ_i (write $\sigma'_i \prec \sigma_i$) if

$$u_i(\sigma_i, \sigma_{-i}) > u_i(\sigma'_i, \sigma_{-i}) \quad \text{for all } \sigma_{-i} \in \Sigma_{-i}$$

Example 26

	X	Y
A	3	0
B	0	3
C	1	1

Is there a strictly dominated strategy?

Question: Is there a game with at least one strictly dominated strategy but without strictly dominated *pure* strategies?

Strictly Dominant Strategy Equilibrium

Definition 27
$\sigma_i \in \Sigma_i$ is *strictly dominant* if every other mixed strategy of player i is strictly dominated by σ_i.

Definition 28
A strategy profile $\sigma \in \Sigma$ is a *strictly dominant strategy equilibrium* if $\sigma_i \in \Sigma_i$ is strictly dominant for all $i \in N$.

Proposition 2
If the strictly dominant strategy equilibrium exists, it is unique, all its strategies are pure, and rational players will play it.

Proof.
Let $\sigma^* = (\sigma_1^*, \ldots, \sigma_n^*) \in \Sigma_i$ be the strictly dominant strategy equilibrium.

By Corollary 24, for every $i \in N$ and $\sigma_{-i} \in \Sigma_{-i}$, there must exist $s_i \in S_i$ such that $u_i(\sigma_i^*, \sigma_{-i}) \leq u_i(s_i, \sigma_{-i})$.

But then $\sigma_i^* = s_i$ since σ_i^* is strictly dominant.

□

Computing Strictly Dominant Strategy Equilibrium

How to decide whether there is a strictly dominant strategy equilibrium $s = (s_1, \ldots, s_n) \in S$?

I.e. whether for a given $s_i \in S_i$, all $\sigma_i \in \Sigma_i \smallsetminus \{s_i\}$ and all $\sigma_{-i} \in \Sigma_{-i}$:

$$u_i(s_i, \sigma_{-i}) > u_i(\sigma_i, \sigma_{-i})$$

There are some serious issues here:

- Obviously there are uncountably many possible σ_i and σ_{-i}.
- $u_i(\sigma_i, \sigma_{-i})$ is nonlinear, and for more that two players even $u_i(s_i, \sigma_{-i})$ is nonlinear in probabilities assigned to pure strategies.

First, we prove the following useful proposition using Lemma 23:

Lemma 29

σ_i *strictly dominates* σ_i' **iff** *for all* pure strategy profiles $s_{-i} \in S_{-i}$*:*

$$u_i(\sigma_i, s_{-i}) > u_i(\sigma_i', s_{-i})$$

Proof: Simple application of the second equality from Lemma 23.

In other words, it suffices to check the strict dominance only with respect to all *pure* profiles of opponents.

Computing Strictly Dominant Strategy Equilibrium

How to decide whether for a given $s_i \in S_i$, all $\sigma_i \in \Sigma_i \setminus \{s_i\}$ and all $s_{-i} \in S_{-i}$ we have

$$u_i(s_i, s_{-i}) > u_i(\sigma_i, s_{-i})$$

Lemma 30

$u_i(s_i, s_{-i}) > u_i(\sigma_i, s_{-i})$ for all $\sigma_i \in \Sigma_i \setminus \{s_i\}$ and all $s_{-i} \in S_{-i}$

iff

$u_i(s_i, s_{-i}) > u_i(s_i', s_{-i})$ for all $s_i' \in S_i \setminus \{s_i\}$ and all $s_{-i} \in S_{-i}$.

Proof: Simple application of the first equality from Lemma 23.

Thus it suffices to check whether $u_i(s_i, s_{-i}) > u_i(s_i', s_{-i})$ for all $s_i' \in S_i$ and all $s_{-i} \in S_{-i}$.

This can easily be done in time polynomial w.r.t. $|G|$.

IESDS in Mixed Strategies

Define a sequence $D_i^0, D_i^1, D_i^2, \ldots$ of strategy sets of player i.
(Denote by G_{DS}^k the game obtained from G by restricting the pure strategy sets to D_i^k, $i \in N$.)

1. Initialize $k = 0$ and $D_i^0 = S_i$ for each $i \in N$.
2. For all players $i \in N$: Let D_i^{k+1} be the set of all pure strategies of D_i^k that are *not* strictly dominated in G_{DS}^k by *mixed strategies*.
3. Let $k := k + 1$ and go to 2.

We say that $s_i \in S_i$ *survives IESDS* if $s_i \in D_i^k$ for all $k = 0, 1, 2, \ldots$

Definition 31
A strategy profile $s = (s_1, \ldots, s_n) \in S$ is an *IESDS equilibrium* if each s_i survives IESDS.

IESDS – Algorithm

Note that in step 2 it is not sufficient to consider pure strategies. Consider the following zero sum game:

	X	Y
A	3	0
B	0	3
C	1	1

C is strictly dominated by $(\sigma_1(A), \sigma_1(B), \sigma_1(C)) = (\frac{1}{2}, \frac{1}{2}, 0)$ but no strategy is strictly dominated in pure strategies.

However, there are uncountably many mixed strategies that may dominate a given pure strategy ...

Recall $u_i(\sigma_i, \sigma_{-i})$ is linear in σ_i. So to decide strict dominance, we use linear programming ...

Intermezzo: Linear Programming

Linear programming is a technique for optimization of a linear objective function, subject to linear (non-strict) inequality constraints.

Formally, a linear program in so called *canonical form* looks like this:

$$\text{maximize} \sum_{j=1}^{m} c_j x_j \qquad \textit{(objective function)}$$

$$\text{subject to} \sum_{j=1}^{m} a_{ij} x_j \leq b_i \qquad 1 \leq i \leq n$$

$$x_j \geq 0 \qquad 1 \leq j \leq m \qquad \textit{(constraints)}$$

Here a_{ij}, b_k and c_j are real numbers and x_j's are real variables.

A *feasible solution* is an assignment of real numbers to the variables $x_j, 1 \leq j \leq m$, so that the *constraints* are satisfied.

An *optimal solution* is a feasible solution which maximizes the *objective function* $\sum_{j=1}^{m} c_j x_j$.

Intermezzo: Complexity of Linear Programming

We assume that coefficients a_{ij}, b_k and c_j are encoded in binary (more precisely, as fractions of two integers encoded in binary).

Theorem 32 (Khachiyan, Doklady Akademii Nauk SSSR, 1979)

There is an algorithm which for any linear program computes an optimal solution in polynomial time.

The algorithm uses so called ellipsoid method.

In practice, the Khachiyan's is not used. Usually **simplex algorithm** is used even though its theoretical complexity is exponential.

There is also a polynomial time algorithm (by Karmarkar) which has better complexity upper bounds than the Khachiyan's and sometimes works even better than the simplex.

There exist several advanced linear programming solvers (usually parts of larger optimization packages) implementing various heuristics for solving large scale problems, sensitivity analysis, etc.

For more info see
http://en.wikipedia.org/wiki/Linear_programming#Solvers_and_scripting_.28programming.29_languages

IESDS Algorithm – Strict Dominance Step

So how do we use linear programming to decide strict dominance in step 2 of IESDS procedure?
I.e. whether for a given s_i there exists σ_i such that for all σ_{-i} we have

$$u_i(\sigma_i, \sigma_{-i}) > u_i(s_i, \sigma_{-i})$$

Recall that by Lemma 29 we have that σ_i is strictly dominates σ_i' **iff** for all *pure strategy profiles* $s_{-i} \in S_{-i}$:

$$u_i(\sigma_i, s_{-i}) > u_i(\sigma_i', s_{-i})$$

In other words, it suffices to check the strict dominance only with respect to all *pure* profiles of opponents.

IESDS Algorithm – Strict Dominance Step

Recall that $u_i(\sigma_i, s_{-i}) = \sum_{s_i' \in S_i} \sigma_i(s_i') u_i(s_i', s_{-i})$.

So to decide whether $s_i \in S_i$ is strictly dominated by some mixed strategy σ_i, it suffices to solve the following system:

$$\sum_{s_i' \in S_i} x_{s_i'} \cdot u_i(s_i', s_{-i}) > u_i(s_i, s_{-i}) \qquad s_{-i} \in S_{-i}$$

$$x_{s_i'} \geq 0 \qquad s_i' \in S_i$$

$$\sum_{s_i' \in S_i} x_{s_i'} = 1$$

(Here each variable $x_{s_i'}$ corresponds to the probability $\sigma_i(s_i')$ assigned by the strictly dominant strategy σ_i to s_i')

Unfortunately, this is a "strict linear program" ... How to deal with the strict inequality?

IESDS Algorithm – Complexity

Introduce a new variable y to be **maximized** under the following constraints:

$$\sum_{s_i' \in S_i} x_{s_i'} \cdot u_i(s_i', s_{-i}) \geq u_i(s_i, s_{-i}) + y \qquad s_{-i} \in S_{-i}$$

$$x_{s_i'} \geq 0 \qquad s_i' \in S_i$$

$$\sum_{s_i' \in S_i} x_{s_i'} = 1$$

$$y \geq 0$$

Now s_i is strictly dominated **iff** a solution maximizing y satisfies $y > 0$

The size of the above program is polynomial in $|G|$.

So the step 2 of IESDS can be executed in polynomial time.

As every iteration of IESDS removes at least one pure strategy, IESDS runs in time polynomial in $|G|$.

IESDS in Mixed Strategie – Example

	X	Y
A	3	0
B	0	3
C	1	1

Let us have a look at the first iteration of IESDS.

Observe that A, B are not strictly dominated by any mixed strategy.

Let us construct the linear program for deciding whether C is strictly dominated: The program maximizes y under the following constraints:

$$\begin{aligned} 3x_A + 0x_B + x_C &\geq 1 + y \\ 0x_A + 3x_B + x_C &\geq 1 + y \\ x_A, x_B, x_C &\geq 0 \\ x_A + x_B + x_C &= 1 \\ y &\geq 0 \end{aligned}$$

The maximum $y = \frac{1}{2}$ is attained at $x_A = \frac{1}{2}$ and $x_B = \frac{1}{2}$.