

Computational Game Theory

Vincenzo Bonifaci

October 16, 2012

1 Games: examples and definitions

Game theory deals with situations in which multiple rational, self-interested entities (individuals, firms, nations, etc.) have to interact.

A normal-form game tries to model a situation in which the entities have to take their decisions simultaneously and independently.

An example is the following Rock-Paper-Scissors game. We can represent it by a table in which the rows correspond to decisions of Player 1, and the columns to decisions of Player 2.

P1, P2	rock	paper	scissors
rock	draw	P2 wins	P1 wins
paper	P1 wins	draw	P2 wins
scissors	P2 wins	P1 wins	draw

Definition 1.1. A *normal form game* is given by:

- a set N (set of *players*); often we use $N = \{1, 2, \dots, n\}$
- for each $i \in N$, a nonempty set S_i (*strategies* of player i)

The set $S := S_1 \times S_2 \times \dots \times S_n$ is called the set of *states* of the game.

- for each $i \in N$, a function $u_i : S \rightarrow \mathbb{R}$ (*utility* or *payoff function*)

Example 1.2 (Rock-Paper-Scissors).

u_1, u_2	rock	paper	scissors
rock	0, 0	-1, 1	1, -1
paper	1, -1	0, 0	-1, 1
scissors	-1, 1	1, -1	0, 0

Notice that Rock-Paper-Scissors is a *zero-sum* game: in any state of the game, the sum of the utilities of the players is constant. The Rock-Paper-Scissors game is also *finite*: the set N of players has finite cardinality, as do the strategy sets S_1, \dots, S_n .

Example 1.3 (Prisoner’s dilemma). Two suspects are interrogated in separate rooms. Each of them can confess or not confess their crime. If both confess, they get 4 years each in prison. If one confess and the other does not, the one that confessed gets 1 year and the other 5. If both are silent, they get 2 years each.

Like in this case, sometimes it is more natural to use *cost* functions $(c_i)_{i \in N}$ instead of utility functions $(u_i)_{i \in N}$; it is equivalent, since we can always define $u_i := -c_i$.

c_1, c_2	confess	silent
confess	4, 4	1, 5
silent	5, 1	2, 2

Notice that the Prisoner’s dilemma is *not* a zero-sum game; however it is a finite game.

So far we saw two-player games, but obviously there are games with more players.

Example 1.4 (Bandwidth sharing). A group of n users has to share a common Internet connection with finite bandwidth. Each user can decide what fraction of the bandwidth to use (any amount between none and all). The payoff of each user is higher if this fraction is higher, but is lower if the remaining available bandwidth is too small (packets get delayed too much).

We can model this by defining

- $N := \{1, \dots, n\}$;
- $S_i := [0, 1]$ for each $i \in N$;
- $u_i(s) := s_i \cdot (1 - \sum_{j \in N} s_j)$, where $s_i \in S_i$ is the strategy selected by player i and $s = (s_1, s_2, \dots, s_n)$.

Notice that this game is not finite: the set of players is finite, but the strategy sets have infinite cardinality.

Example 1.5 (“Chicken”). Two drivers are headed against each other on a single lane road. Each of them can continue straight ahead or deviate. If both deviate, they both get low payoff. If one deviates while the other continues, he is a “Chicken” and will get low payoff, while the payoff for the other player will be high. If both continue straight ahead, however, a disaster will occur which will cost a lot to the players, as both cars will be destroyed.

u_1, u_2	deviate	straight
deviate	0, 0	-1, 5
straight	5, -1	-100, -100

Notice that the type of games we discussed (normal-form games) are “one-shot” in the sense that players move simultaneously and interact only once. There are also model of games in which players move one after the other (*extensive games*) or in which the same game is played many times (*repeated games*). However, in the course we will focus on normal-form games.

Representing the game computationally

When we need to process a game computationally, we have to find some means of representing the game in a concise way. In a normal-form game with a constant number of players, we can represent the whole payoff table explicitly; its size will be polynomial in the total number of strategies. If the number of players is not constant (as in the bandwidth sharing game) we need to represent the functions that compute the payoffs, by encoding them in some formalism (e.g. as C programs or Turing machines).

We also notice that in general we cannot represent real values; in most cases we will need to assume that the codomain of the payoff functions is not \mathbb{R} , but rather \mathbb{Z} or \mathbb{Q} .

2 Solution concepts

After we have modeled a game, we would like to know which states of the game represent outcomes that are likely to occur, assuming that players are self-interested and rational. There are different ways to do this; each of them gives rise to a different *solution concept*. Different solution concepts have different interpretations, advantages and drawbacks.

2.1 Dominant strategy equilibrium

Consider a state of a game $\Gamma = (N, (S_i)_{i \in N}, (u_i)_{i \in N})$. The utility of a player i in state $s \in S$ will depend on both the action of player i himself (s_i) as well as on the actions of the other players, which we denote conventionally by s_{-i} . So we can rewrite $u_i(s)$ (utility of player i in state s) as $u_i(s_i, s_{-i})$. *Be careful* when reading (or using) this notation: we are not reordering the components of the vector s , we are just writing them differently. For example, with (z_i, s_{-i}) we simply mean the state vector that is obtained from s by replacing the i -th component of state s with z_i .

The idea of a dominant strategy equilibrium is that if a player has an action that is the best among his actions *independently of what the other players do*, then this is certainly a possible outcome of the game. This is formalized as follows.

Definition 2.1. State $s \in S$ is a *dominant strategy equilibrium* if for all $i \in N$ and for all $s' \in S$,

$$u_i(s_i, s'_{-i}) \geq u_i(s'_i, s'_{-i}).$$

(In terms of costs: $c_i(s_i, s'_{-i}) \leq c_i(s'_i, s'_{-i})$.)

Example 2.2 (Dominant strategy in the Prisoner's dilemma). Is (silent,silent) a dominant strategy in the Prisoner's dilemma game? The answer is no: if $s = (\text{silent}, \text{silent})$, there is a player ($i = 1$) and there is an alternative state $s' = (\text{confess}, \text{silent})$ for which $c_1(\text{silent}, \text{silent}) > c_1(\text{confess}, \text{silent})$. This contradicts the definition.

Is (confess,confess) a dominant strategy? We have to check 8 cases (2 players times 4 states) to be sure, but the answer is yes. The point is that no matter what the other player is doing, for each player it is cheaper to confess. So (confess,confess) is a dominant strategy.

A dominant strategy equilibrium represents a "strong" type of equilibrium: every player can rely on his strategy independently of what the others are doing. Unfortunately, it has a big drawback: it does not always exist!

Exercise 2.1. Show that the Chicken game has no dominant strategy equilibrium.

Since it does not always exist, we cannot use the dominant strategy equilibrium concept to predict what will happen in a game : the players will certainly do *something*, and this something will not in general be a dominant strategy equilibrium, simply because the game might not admit one.

2.1.1 Finding dominant strategy equilibriums

How do we find, given a game, its dominant strategy equilibriums? If the game is finite and there is a constant number of players this can be done efficiently. Since there is a polynomial number of states ($|S_1| \cdot |S_2| \cdot \dots \cdot |S_n|$, where n is constant) we simply check for every state whether it satisfies the condition in the definition of dominant strategy equilibrium.

2.2 Pure Nash equilibrium

The idea of a pure Nash equilibrium is of that of calling a state an equilibrium if for every player, *assuming that other players are not changing their action*, the player is selecting his “best” action. That is, no player has an incentive to deviate unilaterally from his action; no one has an interest to alter the “status quo”.

Definition 2.3. A state $s \in S$ is a *pure Nash equilibrium* (PNE) if for all $i \in N$ and for all $s'_i \in S_i$,

$$u_i(s_i, s_{-i}) \geq u_i(s'_i, s_{-i}).$$

The definition is superficially very similar to that of dominant strategy: take your time to appreciate the difference.

However, there is a similarity and in fact every dominant strategy equilibrium is also a pure Nash equilibrium (can you see why?).

The converse is not true: some games without dominant strategy equilibriums have pure Nash equilibria.

Example 2.4 (PNE in the Chicken game). Is the state $s = (\text{straight}, \text{straight})$ a PNE in the Chicken game? The answer is no: there is a player ($i = 1$) and an alternative strategy s'_i (deviate) such that $-1 = u_1(\text{straight}, \text{straight}) < u_1(\text{deviate}, \text{straight}) = -100$. This contradicts the definition.

Is the state $s = (\text{deviate}, \text{straight})$ a PNE in the Chicken game? Let's see. If player 1 knows that player 2 is going straight, deviating (-1) is better than going straight (-100). On the other hand, if player 2 knows that player 1 is deviating, going straight (5) is better than deviating (0). So $(\text{deviate}, \text{straight})$ is a PNE.

Notice that PNE need not be unique: in fact, in the Chicken game, there are two PNE (which is the other one?).

Let's look at a more complicated example.

Example 2.5 (PNE in the Bandwidth sharing game). Let's see what player i will do when the strategies of the other players are $s_1, \dots, s_{i-1}, s_{i+1}, \dots, s_n$. Let's define $t := \sum_{j \neq i} s_j$. From the point of view of player i , the quantity t is a constant. By definition of the payoffs we have $u_i(s) = s_i \cdot (1 - t - s_i)$.

Player i can control the one-dimensional variable $s_i \in [0, 1]$. If we take the derivative of $u_i(s)$ with respect to s_i we obtain

$$\frac{\partial}{\partial s_i} u_i(s) = 1 - t - 2s_i.$$

By standard analysis we know that the maximum of u_i is achieved when $\frac{\partial}{\partial s_i} u_i(s) = 0$ (or, possibly, when s_i is at an extreme point of $[0, 1]$, but this is not the case in our example because we get the worst possible payoff in that case). So the player will select $s_i = \frac{1}{2}(1 - t) = \frac{1}{2}(1 - \sum_{j \neq i} s_j)$. This will be true for all $i \in N$, so by symmetry we find out that $s_i = 1/(n + 1)$ for all i .

Unfortunately, although the PNE solution concept applies to a larger class of games, it has basically the same problem as that of a dominant strategy equilibrium: it does not always exist.

Exercise 2.2. Show that the Rock-Paper-Scissors game has no PNE.

2.2.1 Finding pure Nash equilibria

When the game is finite and the number of players is constant, we can find efficiently all pure Nash equilibria of the game via a simple enumeration of all states, as we did in the case of dominant strategy equilibria.

2.3 Mixed Nash equilibrium

So far there was no way for a player to interpolate between two actions: either he selects action s_i or he performs another action s_j . We now relax this constraint by allowing the player to choose actions with certain probabilities. For example he might choose action s_1 with probability $1/4$, action s_2 with probability $1/3$, and action s_3 with probability $5/12$. Such strategies are called *mixed*, in contrast with the usual deterministic *pure* strategies. Pure strategies are perhaps more natural, but often the strategies arising in a game are in fact mixed strategies.

Definition 2.6. A *mixed strategy* for player i is a probability distribution on the set of S_i of pure strategies. That is, it is a function $p_i : S_i \rightarrow [0, 1]$ such that $\sum_{s_i \in S_i} p_i(s_i) = 1$. A *mixed state* is a family $(p_i)_{i \in N}$ consisting of one mixed strategy for each player.

Notice that every pure state s has probability $p(s) := p_1(s_1) \cdot p_2(s_2) \cdot \dots \cdot p_n(s_n)$ of being realized.

Thus, a mixed state $(p_i)_{i \in N}$ induces an *expected payoff* for player i equal to $\sum_{s \in S} p(s) \cdot u_i(s)$. This is the expected payoff of a state selected probabilistically by the players according to their mixed strategies.

We can now define the notion of mixed Nash equilibrium (MNE).

Definition 2.7. A mixed state is a *mixed Nash equilibrium* if no player can unilaterally improve his expected payoff by switching to a different mixed strategy.

Since mixed strategies generalize pure strategies, it is not hard to see that every PNE is also a MNE. The opposite is not true. In fact, there are games without PNE that admit MNE. More than that: the surprising fact is that any finite game (game where N and S are finite) admits at least one mixed Nash equilibrium!

Theorem 2.1 (Nash 1950). *Every finite game admits at least one mixed Nash equilibrium.*

Example 2.8 (MNE for the Rock-Paper-Scissors game). We saw that the Rock-Paper-Scissors game has no pure Nash equilibria. According to Nash's Theorem it should have at least one equilibrium. In fact, we claim that if we define $p := (1/3, 1/3, 1/3)$, then (p, p) is a MNE.

Let's verify this. Consider for example player 1. We should check that when player 2 uses probability distribution p , player 1 has no incentive to play a mixed strategy different from p . (We should also do a similar check with the roles of the players reversed, but in this case everything will be symmetric.)

If player 2 uses mixed strategy p , and player 1 uses a generic mixed strategy $q = (a, b, c)$ where $a + b + c = 1$, then the expected payoff for player 1 becomes

$$\begin{aligned} & a \cdot 1/3 \cdot (0) + a \cdot 1/3 \cdot (-1) + a \cdot 1/3 \cdot (+1) + \\ & b \cdot 1/3 \cdot (+1) + b \cdot 1/3 \cdot (0) + b \cdot 1/3 \cdot (-1) + \\ & c \cdot 1/3 \cdot (-1) + c \cdot 1/3 \cdot (+1) + c \cdot 1/3 \cdot (0) = 0. \end{aligned}$$

So the expected payoff is a constant (0) no matter what a , b and c are! This means that there is no point for player 1 in changing them. Similarly, when player 1 plays $(1/3, 1/3, 1/3)$, player 2 has no incentive to change his strategy from $(1/3, 1/3, 1/3)$. The two players "lock" each other in the mixed Nash equilibrium.

At this point you might wonder why another mixed state, like $(1/2, 1/4, 1/4)$ for both players, is not a MNE. The reason is that if e.g. player 2 plays something different from $(1/3, 1/3, 1/3)$, then the player 1 is no longer indifferent between his possible responses. In this case, when player 2 plays $(1/2, 1/4, 1/4)$, it will be more convenient for player 1 to play $(0, 1, 0)$ than to play $(1/2, 1/4, 1/4)$: since player 2 is playing Rock more often than Paper or Scissors, it is best for player 1 to always play Paper (you can check this by computing the expected payoff for player 1). So $((1/2, 1/4, 1/4), (1/2, 1/4, 1/4))$ is not a MNE.

2.3.1 Finding mixed Nash equilibria

Finding MNE is considerably harder than finding dominant strategy equilibria or PNE, even when the game is finite and there are only a constant number of players, and even when there are only two players. Apparently we have to check for an infinite set of mixed states, so it is not even clear that we can do it in finite time!

Luckily, there are some notions that can help us.

Definition 2.9. A mixed strategy p_i is a *best response* to mixed strategies $p_1, \dots, p_{i-1}, p_{i+1}, \dots, p_n$ if for all mixed strategies p'_i of player i ,

$$\sum_{s \in S} p_1(s_1) \dots p_i(s_i) \dots p_n(s_n) \cdot u_i(s) \geq \sum_{s \in S} p_1(s_1) \dots p'_i(s_i) \dots p_n(s_n) \cdot u_i(s).$$

That is, a best response attains the maximum possible expected utility among all possible mixed strategies of this player. In fact, we can now say that *a mixed state is a MNE if and only if every player is playing a best response strategy.*

Definition 2.10. The *support* of a mixed strategy p_i is the set of all pure strategies that have nonzero probability in it: $\text{supp}(p_i) := \{j \in S_i : p_i(j) > 0\}$.

Example 2.11. If $p_i = (\frac{1}{3}, 0, 0, \frac{1}{2}, \frac{1}{6})$ then the support of p_i is $\{1, 4, 5\}$.

The following characterization will be very useful for computing mixed Nash equilibria.

Theorem 2.2. A mixed strategy p_i is a best response if and only if all pure strategies in $\text{supp}(p_i)$ are best responses.

Proof. If all strategies in $\text{supp}(p_i)$ are best responses, then since the mixed strategy is a convex combination of them, it will have the same expected payoff and also be a best response.

On the other hand, if mixed strategy p_i is a best response, all pure strategies in its support are best responses: suppose this was not the case, then by decreasing the probability of the pure strategy with *worst* expected payoff, and redistributing the remaining probability proportionally for the other pure strategies in the support, we could improve the expected payoff. But then p_i would not be a best response. \square

Thus the hard part in finding a MNE is finding the right supports.

Suppose that we are given a *finite* two-player game. This can be completely specified by two *payoff matrices* $A = (a_{ij})_{ij}, B = (b_{ij})_{ij} \in \mathbb{R}^{m_1 \times m_2}$ (which is why these games are also called *bimatrix games*), where $S_1 = \{1, \dots, m_1\}$ and $S_2 = \{1, \dots, m_2\}$ are the strategy sets. If we knew the supports $I \subseteq S_1$ and $J \subseteq S_2$, to check whether there is a MNE with these supports it would be enough to check whether the following system has a solution in the (vector) variables x, y :

$$\sum_{j \in J} y_j a_{kj} \leq \sum_{j \in J} y_j a_{ij} \quad \forall k \in S_1, \forall i \in I \quad (1)$$

$$\sum_{i \in I} x_i b_{ik} \leq \sum_{i \in I} x_i b_{ij} \quad \forall k \in S_2, \forall j \in J \quad (2)$$

$$\sum_{i \in I} x_i = 1 \quad (3)$$

$$\sum_{j \in J} y_j = 1 \quad (4)$$

$$x_i \geq 0 \quad \forall i \in I \quad (5)$$

$$y_j \geq 0 \quad \forall j \in J. \quad (6)$$

Intuitively, the equations (1) state that every pure strategy in the support of x (that is, I) is a best response to mixed strategy y : no other pure strategy k in S_1 can achieve better expected payoff. Similarly, equations (2) state that every pure strategy in the support of y is a best response to mixed strategy x . Equations (3)–(6) simply state that x and y are in fact mixed strategies (probability distributions on S_1 and S_2 , respectively).

Theorem 2.3. There is an algorithm that finds a MNE of a bimatrix game (A, B) where $A, B \in \mathbb{Q}^{m_1 \times m_2}$ in time $2^{m_1+m_2} \cdot \text{poly}(\text{bits}(A) + \text{bits}(B))$.

Proof. We simply enumerate all possible supports $I \subseteq S_1, J \subseteq S_2$ and for each of them check whether the above linear system is feasible. If it is, then (x, y) is a MNE. \square

Example 2.12. Let's go back to the Chicken game.

u_1, u_2	deviate	straight
deviate	0, 0	-1, 5
straight	5, -1	-100, -100

We saw that the game has two pure Nash equilibria: (deviate, straight) and (straight, deviate). Let's see if it has one mixed Nash equilibrium. It is easy to see that in this case, when the support of one of the players has size one, the other player best response is a single pure strategy. Thus, since we already investigated the PNE of this game, any other MNE (if there is one) will necessarily have a support of size at least two for *both* players. So there is no need to enumerate all the possible supports I and J : we can directly take $I = J = \{d, s\}$ where d and s are shorthand for "deviate" and "straight".

The linear program then gives us:

$$\begin{aligned} y_d \cdot 0 + y_s \cdot (-1) &= y_d \cdot 5 + y_s \cdot (-100) \\ x_d \cdot 0 + x_s \cdot (-1) &= x_d \cdot 5 + x_s \cdot (-100) \\ y_d + y_s &= 1 \\ x_d + x_s &= 1 \\ x_d, x_s, y_d, y_s &\geq 0. \end{aligned}$$

The solution is $x_d = y_d = 99/104$, $x_s = y_s = 5/104$. So we discovered another equilibrium. In this equilibrium both players will deviate from their route with high probability, but each of them has a small probability ($5/104$) of going straight.

Exercise 2.3. Find all MNE of the following bimatrix game.

u_1, u_2	Action 1	Action 2
Action 1	2, 1	0, 3
Action 2	1, 2	4, 1

Remark 2.1. It is not known whether finding a Nash equilibrium of a finite two-player game is a problem that can be solved in polynomial time, although there is some complexity-theoretic evidence that it is not. However the problem is definitely not NP-hard: in fact, the associated decision problem is trivial, as the answer is always "yes"!

2.3.2 The case of zero-sum games

In the special case where we have a *zero-sum* two-player game, it turns out that there is a polynomial time algorithm to find a Nash equilibrium. In fact, consider such a game; this can be specified by a *single* matrix $A = (a_{ij})_{ij}$, whose entries represent simultaneously the payoffs for the first (row) player as well as the costs for the second (column) player.

Assume that the column player knows that row player is playing mixed strategy x . Then the column player will look at the expected payoff vector xA , and since he wants to minimize his loss, he will only play strategies that correspond to minimum entries in this vector. So if we now consider things from the point of view of the row player, he can secure himself a payoff of v if he selects a mixed strategy x such that no matter what the second player plays, the payoff will be at least v . We thus have the following linear program for maximizing the “safety level” v :

$$\begin{aligned} v^* &= \max v \\ \sum_{i \in S_1} x_i a_{ij} &\geq v \quad \forall j \in S_2 \\ \sum_{i \in S_1} x_i &= 1 \\ x_i &\geq 0 \quad \forall i \in S_1. \end{aligned}$$

Similarly, for the column player we get the following program:

$$\begin{aligned} u^* &= \min u \\ \sum_{j \in S_2} y_j a_{ij} &\leq u \quad \forall i \in S_1 \\ \sum_{j \in S_2} y_j &= 1 \\ y_j &\geq 0 \quad \forall j \in S_2. \end{aligned}$$

We notice that these linear programs are duals of each other! We can now prove the following.

Theorem 2.4. *Optimum solutions for the above pair of linear programs give mixed strategies that form a Nash equilibrium of the two-person zero-sum game.*

Proof. By linear duality $v^* = u^*$. If the players play this pair of strategies, the row player cannot increase his win, as the column player is guaranteed by his strategy not to lose more than u^* . Similarly, the column player cannot decrease his loss under v^* . This means that the pair of strategies is at equilibrium. \square

Corollary 2.5. *There is a polynomial time algorithm for finding a mixed Nash equilibrium in a two-player zero-sum game.*

In fact, the quantity $v^* = u^*$ is called the *value* of the zero-sum game: it is the payoff that the first player can ensure for himself by playing the game at the best of his possibilities. Notice that the value might be negative: in that case it would be better for the first player not to play at all!

Example 2.13. Consider the following zero-sum game.

$u_1 = -u_2$	Action 1	Action 2
Action 1	2	-1
Action 2	1	3




Figure 1: Example 2.13

The first LP becomes:

$$\begin{aligned}
 & \max v \\
 & x_1 \cdot 2 + x_2 \cdot 1 \geq v \\
 & x_1 \cdot (-1) + x_2 \cdot 3 \geq v \\
 & x_1 + x_2 = 1 \\
 & x_1, x_2 \geq 0.
 \end{aligned}$$

If we substitute $x_2 = 1 - x_1$, we can rewrite the main inequalities as $x_1 + 1 \geq v$ and $-4x_1 + 3 \geq v$. From these we find out (Figure 1) that the best possible value of v is $7/5$, obtained when $x_1 = 2/5$, $x_2 = 3/5$. Similarly for the column player, we obtain that $y_1 = 4/5$, $y_2 = 1/5$.

2.3.3 Degenerate games

Sometimes it can happen that we have a “degenerate” kind of equilibrium. Consider a zero-sum game given by the following matrix:

$$\begin{pmatrix} 2 & 4 \\ 2 & 5 \end{pmatrix}$$

By solving for the second player’s equilibrium strategy we find $y_1 = 1$, $y_2 = 0$. However if we write the LP for the first player:

$$\begin{aligned}
 & \max v \\
 & 2x_1 + 2x_2 \geq v \\
 & 4x_1 + 5x_2 \geq v \\
 & x_1 + x_2 = 1 \\
 & x_1, x_2 \geq 0.
 \end{aligned}$$

we find out that any probability distribution (x_1, x_2) is feasible (and $v = 2$). What is happening? The point is that, as the second player will play the first column, it does not really matter which row the first player selects. So we have infinitely many MNE, of the form $((\epsilon, 1 - \epsilon), (1, 0))$ for any $\epsilon \in [0, 1]$.

2.3.4 Dominated strategies

Another useful concept to keep in mind in the study of equilibria is that of a *dominated strategy*.

Definition 2.14. A pure strategy s_i of a player $i \in N$ is *strictly dominated* by a strategy s'_i of the same player if, for each combination s_{-i} of strategies of the remaining players,

$$u_i(s_i, s_{-i}) < u_i(s'_i, s_{-i}).$$

Thus, a strictly dominated strategy is a strategy for which there is an alternative that is always strictly better for the player, independently of the actions of the others. As such, it is not rational to play a strictly dominated strategy and in fact it can be easily proven that they are never part of a Nash equilibrium.

A useful preprocessing step, when analyzing a game, is then to eliminate from it strategies that are strictly dominated. Since the strategy is strictly dominated, we are not “forgetting” any equilibrium in this way. This procedure can be iterated until no strategy is strictly dominated.

Example 2.15. Consider the following bimatrix game (A, B) :

$$A = \begin{pmatrix} 0 & 2 & 5 \\ 2 & 4 & -1 \end{pmatrix}, \quad B = \begin{pmatrix} 4 & 5 & -1 \\ 2 & -3 & 0 \end{pmatrix}$$

Looking at A , we see that no strategy of the row player is strictly dominated. Looking at B , we see that the third column is strictly dominated by the first one. We can thus eliminate the third column and obtain the simpler, but equivalent, game:

$$A = \begin{pmatrix} 0 & 2 \\ 2 & 4 \end{pmatrix}, \quad B = \begin{pmatrix} 4 & 5 \\ 2 & -3 \end{pmatrix}$$

In this new game, the first row in matrix A is strictly dominated by the second row. Thus player 1 will never play the first row. We obtain the even simpler game:

$$A = (2 \ 4), \quad B = (2 \ -3)$$

We can finally conclude that player 2 will select the column that gives payoff 2. We have thus reached a pure Nash equilibrium where both players have payoff 2. This is also a PNE of the original game, and from what we said it must be the *only* one (whether pure or mixed), because all the strategies we discarded cannot be part of an equilibrium.

The notion of strictly dominated strategy can be weakened to allow for equality of the payoffs, as follows.

Definition 2.16. A pure strategy s_i of a player $i \in N$ is *weakly dominated* by a strategy s'_i of the same player if, for each combination s_{-i} of strategies of the remaining players,

$$u_i(s_i, s_{-i}) \leq u_i(s'_i, s_{-i}).$$

Differently from strictly dominated strategies, we *cannot* eliminate weakly dominated strategies from a game while being sure that we do not remove a Nash equilibrium. In fact, if we did so in the example of Section 2.3.3, we would end up with a single pure equilibrium (where player 1 plays the second row and player 2 the first column), while we have already seen that the game had infinitely many Nash equilibria.

2.3.5 Other game examples

Example 2.17 (2/3 of the Mean). There are n players. Each player i selects a real number x_i between 0 and 100 (inclusive). Let $A := \sum_i x_i/n$ be the average of the numbers. The players whose number was closest to $2A/3$ get a payoff of 1, the others get 0.

Exercise 2.4. Determine the pure Nash equilibria of “2/3 of the Mean”.

Example 2.18 (Color the Nodes). Consider an undirected graph where each node represents a player. Each player selects a color: black or white. The payoff for player i is the number of adjacent nodes that have color different from i , minus the number of adjacent nodes having the same color as i .

Exercise 2.5. Prove that there is always at least one pure Nash equilibrium in “Color the Nodes”.

Example 2.19 (Catching the Votes). Each of n people chooses whether or not to become a political candidate, and if so which position to take. There is a continuum of citizens, each of whom has a favorite position; the distribution of favorite positions is given by a density function. A candidate attracts the votes of those citizens whose favorite positions are closer to his position than to the position of any other candidate; if k candidates choose the same position then each receives the fraction $1/k$ of the votes that the position attracts. The winner of the competition is the candidate who receives the most votes. Each person prefers to be the unique winning candidate than to tie for first place, prefers to tie for first place than to stay out of the competition, and prefers to stay out of the competition than to enter and lose.

Exercise 2.6. Formulate “Catching the Votes” as a strategic game and find the set of pure Nash equilibria when $n = 2$.