Introduction to Strategic Games

Dietmar Berwanger*

MPRI 2010/11

Contents

*LSV, ENS Cachan – dwb@lsv.ens-cachan.fr – revision: December 26, 2010

1 Games and solution concepts

1.1 The strategic form

For a finite set N of players, a *game in strategic form* is described by a tuple

$$\mathcal{G} = \big((S_i)_{i \in N}, (u_i)_{i \in N}\big)$$

with the following ingredients:

- a set S_i of *strategies* (or actions) for each player $i \in N$, and
- a *utility* function $u_i : \times_{i \in N} S_i \to \mathbb{R}$, for each player $i \in N$.

To play such a game, each player i chooses a strategy $s_i \in S_i$. This yields the *outcome* $s = (s_i)_{i \in N}$. The utility of the outcome for Player j is the value $u_j(s)$. Fundamentally, we assume that the players choose their actions independently and simultaneously, and that each player seeks to maximise his utility.

When speaking about situations that involve the players in N, we refer to a list of elements $x = (x_i)_{i \in N}$, one for each player, as a *profile*. We write x_{-i} to denote the list $(x_j)_{j \in N \setminus \{i\}}$ of elements in x for each player except i. Given an element x_i and a list x_{-i}, we denote by (x_i, x_{-i}) the profile $(x_i)_{i \in N}$. For clarity, we will always use subscripts to specify to which player an element belongs. If not quantified otherwise, we usually refer to Player i to mean *any* player. Often, but not always, we use the first natural numbers for naming the players and assume $N = \{1, \dots, n\}$.

With this notation, a game is a pair $\mathcal{G} = (S, u)$ where S is a profile of strategy sets and u is a profile of utility functions. We call S the *strategy space* of $\mathcal{G}$. A game is *finite* if its strategy space is finite.

A *best response* of Player i to a strategy profile s_{-i} of the other players is a strategy s_i that yields him the greatest utility, i.e.,

$$u_i(s_i, s_{-i}) \geq u_i(s'_i, s_{-i}) \quad \text{for all } s'_i \in S_i.$$

We specify that s_i is the *strongly* best response, if the inequality is strict for all $s'_i \in S_i$.

One main concern of game theory is to define an appropriate *solution concept*, that is a function that maps any game to a subset of its strategy space consisting of the outcomes that should be expected if the players would play rationally. Ideally, a solution concept should return a single strategy profile and thus predict the outcome of any (play of any) game. However, at present, there is no definitive answer to what it means to play rationally and thus, no unfailing solution concept.

We will discuss several proposals that are helpful for analysing particular game situations.

Bimatrix representation. A game for two players, $i \in \{1,2\}$, can conveniently be represented by a matrix where the rows are associated to the strategies of Player 1 and the columns to those of Player 2. The matrix entry at the intersection of the row for $s_1 \in S_1$ and the column for $s_2 \in S_2$ is the pair of utilities $(u_1(s), u_2(s))$, with $s = (s_1, s_2)$. For an example of such a representation, see Figure 1.

1.2 Dominance

Let $\mathcal{G}$ be a game with the usual notation and let $r_i, s_i \in S_i$ be two strategies of Player i. We say that r_i *dominates* s_i, if

$$u_i(r_i, s_{-i}) > u_i(s_i, s_{-i}) \quad \text{for all } s_{-i} \in S_{-i}.$$

Informally, r_i yields greater utility than s_i, no matter what the other players choose. We say that $s_i \in S_i$ is *dominated*, if there exists a strategy $r_i \in S_i$ that dominates it. A strategy that r_i dominates every other strategy $s_i \in S_i \setminus \{r_i\}$ is called *dominant* strategy. A *dominant strategy equilibrium* is a profile that consists of each player's dominating strategies.

In the Prisoner's Dilemma game represented in Figure 1, the profile (Confess, Confess) is a dominant strategy equilibrium.

Notice that in a dominant strategy equilibrium, each player plays his strongly best response to all strategy profiles of other players. Moreover, if a dominant strategy equilibrium exists, it is unique.

		Colin: Deny	Colin: Confess
Rose	Deny	$-1, -1$	-10, 0
Rose	Confess	$0, -10$	-8, -8

Figure 1: Prisoner's Dilemma

However, not all games admit a dominant strategy equilibrium. For instance, in the game “Battle of the Bismarck Sea” represented in Figure 2 no strategy is dominating, and no strategy is dominated. However, some strategies may be considered dominated in a weaker sense.

For a pair of strategies $r_i, s_i \in S_i$, we say that r_i *weakly dominates* s_i, if

$$\begin{aligned} u_i(r_i, s_{-i}) &\geq u_i(s_i, s_{-i}) \quad \text{for all } s_{-i} \in S_{-i}, \text{ and} \\ u_i(r_i, s_{-i}) &> u_i(s_i, s_{-i}) \quad \text{for some } s_{-i} \in S_{-i}. \end{aligned}$$

Informally, no matter what the other players choose, r_i is never worse than s_i, and sometimes better.

As we did for the strict case, we call a strategy $s_i \in S_i$ *weakly dominated* if there exists a strategy $r_i \in S_i$ that weakly dominates it. In turn, a strategy is *weakly dominating* if it weakly dominates all other strategies. A *weakly dominant equilibrium* is a profile that consists of each player's weakly dominating strategies.

In the game of Figure 2, the strategy North of Imamura weakly dominates his strategy South. However, none of the strategies of Kenney is a-priori undominated. Nevertheless, one may argue that Kenney can predict that Imamura will not use his dominated strategy South and therefore reason as if the game would only consist of the first column. In this smaller game Kenney would choose North as it dominates South. This type of reasoning is formalised through the notion of sequential elimination of (weakly) dominated strategies.

We say that a game $\mathcal{G} = (S, u)$ *reduces* to a game $\mathcal{G}' = (S', u)$ *by a one-stage elimination* if, for some player $i \in N$,

- $S'_i \subsetneq S_i$,
- every strategy $s_i \in S_i \setminus S'_i$ is dominated in $\mathcal{G}$,

and for every other player $j \neq i$, we have $S_j = S'_j$. (Strictly speaking we should restrict the utility function to the reduced strategy space, but we drop this detail here.)

An *elimination sequence* for $\mathcal{G}$ is a sequence of games $\mathcal{G}^0, \mathcal{G}^1, \ldots, \mathcal{G}^\ell$ starting from $\mathcal{G}^0 = \mathcal{G}$ and with the property that every $\mathcal{G}^k$ reduces to $\mathcal{G}^{k+1}$ by a one-stage elimination, for every $k < \ell$. An elimination sequence that cannot be prolonged is called *maximal*. If we set out with a finite game, every maximal elimination sequence is obviously finite.

Proposition 1. *For any finite game $\mathcal{G}$, all maximal elimination sequences of dominated strategies terminate at the same game.*

For a game $\mathcal{G} = (S, u)$, a strategy profile $s \in S$ is a *solution obtained by sequential elimination of dominated strategies*, if there exists a maximal elimination sequence $\mathcal{G}^0, \mathcal{G}^1, \ldots, \mathcal{G}^\ell$ with $\mathcal{G}^0 = \mathcal{G}$ and $\mathcal{G}^\ell = (\{s\}, u)$. In this case we say that the game is *solvable* (through iterated elimination of dominated strategies).

Proposition 1 implies that whenever a game is solvable, the solution is independent of the witnessing elimination sequence.

		Imamura: North	Imamura: South
Kenney	North	$2, -2$	$2, -2$
	South	$1, -1$	$3, -3$

Figure 2: Battle of the Bismarck Sea

For the weak notion of dominance, we can define the concept of a solution obtained by sequential elimination of weakly dominated strategies, analogously. For the game of Figure 2, the profile (South,South) is such a solution.

Notice, however, that the statement of Proposition 1 does not hold for weak dominance. Essentially, order dependence is due to the fact that the reason for eliminating a strategy may itself be eliminated in a later step, as illustrated in Figure 3. In this game, the strategy c_1 which justifies dominance of r_1 over r_2 would be eliminated in the second step; at this point, Player 1 may reconsider choosing r_2. This phenomenon makes the concept difficult to justify.

	c_1	c_2
r_1	$2, 2$	$3, 7$
r_2	$0, 6$	$3, -2$

Figure 3: Column c_1 which justifies dominance of row r_1 is eliminated

Due to order dependence a (single-profile) solution obtained through sequential elimination of weakly dominated strategies may not be unique. The game represented in Figure 4 admits as elimination sequences r_3, c_3, c_2, r_2 and r_2, c_2, c_1, r_3 which lead to two different solutions (r_1, c_1) and (r_1, c_3), respectively.

In summary, the notion of sequential elimination of weakly dominated strategies may appear to be of limited use for defining a general solution concept. Nevertheless it can be relevant in particular situations.

	c_1	c_2	c_3
r_1	$2, 12^*$	$1, 10$	$1, 12^*$
r_2	$0, 12$	$0, 10$	$0, 11$
r_3	$0, 12$	$1, 10$	$0, 13$

Figure 4: Different solutions from different elimination sequences

1.3 Nash equilibrium

The dominance-based concepts we have seen so far are helpless in many concrete cases. For instance, in the game of Figure 5, none of the strategies is even weakly dominated. An alternative concept is that of Nash equilibrium.

		Small: Press	Small: Wait
Big	Press	$5,1$	$4,4^*$
Big	Wait	$9,-1$	$0,0$

Figure 5: Big pig, small pig

Let $\mathcal{G}$ be a game with the usual notation. A profile $s \in S$ is a *Nash equilibrium* if s_i is a best response to s_{-i}, for every player $i \in N$. Likewise, we call a profile $s \in S$ is a *strong* Nash equilibrium, if s_i is the strongly best response to s_{-i}, for every player $i \in N$.

One argument in support of the Nash equilibrium concept is that it yields a consistent prescription: when an equilibrium outcome is proposed, none of the players has an incentive to deviate, provided he assumes that every other player will follow the proposal. Furthermore, Nash equilibrium may provide a solution to games when dominance-based concepts cannot, as it is the case in the game of figure 5.

Finally, it turns out that Nash equilibrium subsumes the dominance-based concepts, in a certain sense.

Proposition 2. (i) *If a game has a solution s obtained by sequential elimination of dominated strategies, then s is the unique Nash equilibrium of the game.*

(ii) *If a game $\mathcal{G}$ has a solution s obtained by sequential elimination of* weakly *dominated strategies, then s is a Nash equilibrium of $\mathcal{G}$.*

For (the strict version of) dominance, elimination sequences have the following property which holds even for games that are not solvable.

Proposition 3. *Let s be a Nash equilibrium of a game $\mathcal{G} = (S, u)$ and let $\mathcal{G}^0, \mathcal{G}^1, \ldots, \mathcal{G}^\ell$ be an elimination sequence of dominated strategies. Then $s \in S^k$, for the strategy space S^k of any game $\mathcal{G}^k$ in the sequence.*

However, the above property does not hold for elimination sequences of weakly dominated strategies, as one can see in the game of Figure 6, where the equilibrium (r_2, c_2) is eliminated in the first stage, because c_2 is weakly dominated by c_1.

	c_1	c_2
r_1	$2,6$	$3,2$
r_2	$0,2$	$3,2^*$

Figure 6: Nash equilibrium may be weakly dominated

Worse, the process of eliminating weakly dominated strategies may delete all Nash equilibria in a game, as shown in Figure 7.

	c_1	c_2	c_3
r_1	$1,3$	$1,3$	$2,1$
r_2	$0,2$	$0,2$	$2,2$
r_3	$7,1$	$7,1$	$0,3$

	c_1	c_2
r_1	$1,3$	$2,1$
r_2	$7,1$	$0,3$

Figure 7: Iterated weak dominance may eliminate all equilibria

In turn, strong equilibria do survive under elimination of weakly dominated strategies.

Nevertheless, Nash equilibrium is not the ultimate solution concept. On the one hand, there are games with no equilibria, such as the Matching Pennies game in Figure 8.

	head	tails
heads	$1,-1$	$-1,1$
tails	$-1,1$	$1,-1$

Figure 8: Matching Pennies

On the other hand, there are games with multiple equilibria. Typical examples are coordination games such as Battle of Sexes or the Chicken game of Figure 9. Here, the combination of equilibrium strategies can moreover result in the worst possible outcome.

	dove	hawk
dove	$1,1$	$0,2$
hawk	$2,0$	$-1,-1$

	meat	fish
red	$2,1$	$0,\ 0$
white	$0,0$	$1,\ 2$

Figure 9: Chicken and Battle of Sexes

One approach to dealing with the problem of multiple equilibria are *refinement* concepts. For instance, in the modified version of the Prisoner's Dilemma, given in Figure 10, one could discard the equilibrium (Confess, Confess) by arguing that the involved strategies are weakly dominated, or that the outcome is of utility to both players than the equilibrium (Deny, Deny).

The nonexistence problem can be overcome by introducing randomisation in the choice of strategies, as we will see in the next section.

		Colin	
		Deny	Confess
Rose	Deny	0, 0	-10, 0
	Confess	$0, -10$	-8, -8

Figure 10: Prisoner's Dilemma (modified)

1.4 Minimax

Minimax analysis is a decision-theoretic framework. Rather than relying on predictions about choices of other players, it is concerned with minimising the losses while maximising the gains for one particular player.

We refer in the following to a fixed game $\mathcal{G}$. For a player i, we say that a strategy $s_i \in S_i$ *guarantees* a utility value $v \in \mathbb{R}$, if $u_i(s_i, s_{-i}) \geq v$, for all strategies $s_{-i} \in S_{-i}$. The *safety level* of strategy s_i is

$$L_i(s_i) := \min_{s_{-i} \in S_{-i}} u_i(s_i, s_{-i}).$$

More generally, a player i can guarantee utility v in $\mathcal{G}$, if he has a strategy $s_i \in S_i$ such that $L_i(s_i) \geq v$. The safety level of Player i in the game $\mathcal{G}$ is

$$L_i := \max_{s_i \in S_i} L_i(s_i) = \max_{s_i \in S_i} \min_{s_{-i} \in S_{-i}} u_i(s_i, s_{-i}).$$

A *safety strategy* for Player i is a strategy that guarantees his safety level, i.e., a strategy $s_i \in S_i$ with $L_i(s_i) = L_i$.

Next, let $s_{-i} \in S_{-i}$ be a strategy profile of the opponents of Player i. The *punishment level* $H_i(s_{-i})$ is the greatest utility that Player i can achieve if his opponents play s_{-i}:

$$H_i(s_{-i}) := \max_{s_i \in S_i} u_i(s_i, s_{-i}).$$

The punishment level for Player i in the game $\mathcal{G}$ is

$$H_i := \min_{s_{-i} \in S_{-i}} H_i(s_{-i}) = \min_{s_{-i} \in S_{-i}} \max_{s_i \in S_i} u_i(s_i, s_{-i}).$$

A *punishment strategy* (of the opponents) against Player i is a strategy $s_{-i} \in S_{-i}$ which ensures that he cannot get more than his punishment level: $H_i(s_{-i}) = H_i$.

Proposition 4. *In any finite game, $L_i \leq H_i$, for all players i.*

Proposition 5. *Let $\mathcal{G}$ be a finite game and let $s \in S$ be a Nash equilibrium. Then, $H_i \leq u_i(s)$, for all players i.*

2 Mixed strategies

As long as we restrict our attention to choosing individual strategies, we may not be able to tell anything relevant about certain games. Consider for instance, the Matching Pennies game of Figure 8. There are no (weakly) dominated strategies, and no Nash equilibria; the safety level is -1 and the punishment level is $+1$, for both players. The situation changes, however, if we allow a player to randomise over his strategies. If, e.g., Player 1 tosses a fair coin and chooses his strategy according to the outcome, this would guarantee him utility 0 (in expectation).

2.1 The mixed extension of a game

Let $\mathcal{G} = (S, u)$ be a finite game. For every player $i \in N$, we denote by $\hat{S}_i$ the set $\Delta(S_i)$ of probability distributions over S_i, that is, the set of functions $\sigma_i : S_i \to [0, 1]$ with

$$\sum_{s_i \in S_i} \sigma(s_i) = 1.$$

Let $\hat{S}$ be the profile $(\hat{S}_i)_{i \in N}$. If we assume that each player i plays according to a distribution σ_i, independently of the others, the probability to obtain a fixed profile $s \in S$ as an outcome is:

$$\sigma(s) := \prod_{i \in N} \sigma_i(s_i).$$

Thus, the expected utility when playing according to σ is, for each player i,

$$\hat{u}_i(\sigma) := \sum_{s \in S} \sigma(s) \cdot u_i(s).$$

The *mixed extension* of a game $\mathcal{G}$ for a set N of players is the game $\hat{\mathcal{G}} := (\hat{S}, \hat{u})$, that is, the game with strategy set $\hat{S}_i$ and utility function $\hat{u}_i$, for each player $i \in N$. With an abuse of terminology, we call the strategies $\sigma_i \in \hat{S}_i$ mixed strategies of $\mathcal{G}$.

Notice that the basic concepts that we defined so far – best response, dominance, equilibrium, safety/punishment level – are not necessarily meaningful if we consider games with infinite strategy spaces. However the mixed extension of a finite game has a particular structure: the strategy space is a compact metric space and the utility functions are continuous – such games are called *regular*. Regular games behave in many ways like finite games.

Proposition 6. *Let $\mathcal{G}$ be a regular game.*

(i) *For any $s_{-i} \in S_{-i}$ there exists a best response $s_i \in S_i$.*

(ii) *All elimination sequences terminates in finitely many steps, and all maximal elimination sequences terminate at the same game.*

(iii) *Safety punishment levels are well defined by:*

$$\hat{L}_i = \max_{\sigma_i \in \hat{S}_i} \min_{\sigma_{-i} \in \hat{S}_{-i}} \hat{u}_i(\sigma_i, \sigma_{-i}), \quad \hat{H}_i = \min_{\sigma_{-i} \in \hat{S}_{-i}} \max_{\sigma_i \in \hat{S}_i} \hat{u}_i(\sigma_i, \sigma_{-i}),$$

and there exist safety and punishment strategies.

The *support* of a mixed strategy $\sigma_i \in \hat{S}_i$ is the set of all strategies s_i with $\sigma(s_i) > 0$. If the support of a strategy σ_i consists of a single element s_i, we say that σ_i is a *pure* strategy and often identify it with s_i.

The following remark shows that if a strategy σ_i can guarantee a value v against all pure opponent strategies, then it can also guarantee this value against all mixed opponent strategies. Similarly, if the opponent of i have a strategy σ_{-i} to prevent him from receiving more than v by using pure strategies, then this strategy will also prevent him from receiving more than v by using mixed strategies.

Proposition 7. *For a finite game $\mathcal{G}$, fix a player $i \in N$, and a value $v \in \mathbb{R}$.*

(i) *For every strategy $\sigma_i \in \hat{S}_i$,*

$$\begin{gathered} \textit{if } \hat{u}_i(\sigma_i, s_{-i}) \geq v \textit{ for all } s_{-i} \in S_{-i}, \textit{ then} \\ \hat{u}_i(\sigma_i, \sigma_{-i}) \geq v \textit{ for all } \sigma_{-i} \in \hat{S}_{-i}. \end{gathered}$$

(ii) *For every strategy profile $\sigma_{-i} \in \hat{S}_{-i}$,*

$$\begin{gathered} \textit{if } \hat{u}_i(s_i, \sigma_{-i}) \geq v, \textit{ for all } s_i \in S_i, \textit{ then} \\ \hat{u}_i(\sigma_i, \sigma_{-i}) \geq v, \textit{ for all } \sigma_i \in \hat{S}_i. \end{gathered}$$

Corollary 8.

$$\hat{L}_i = \max_{\sigma_i \in \hat{S}_i} \min_{s_{-i} \in S_{-i}} \hat{u}_i(\sigma_i, s_{-i}) \quad \textit{and} \quad \hat{H}_i = \min_{\sigma_{-i} \in \hat{S}_{-i}} \max_{s_i \in S_i} \hat{u}_i(\sigma_i, s_{-i}).$$

Proof. For all $\sigma_i \in \hat{S}_i$,

$$\min_{\sigma_{-i} \in \hat{S}_{-i}} u_i(\sigma_i, \sigma_{-i}) = \min_{s_{-i} \in S_{-i}} u_i(\sigma_i, s_{-i}),$$

□

Theorem 9. *For any finite game, $L_i \leq \hat{L}_i \leq \hat{H}_i \leq H_i$, for every player $i \in N$.*

If, for a player i in a game $\mathcal{G}$, we have $L_i = H_i = v$, then the utility v is called *value* of the game for player i. In this case, any safety strategy is called *optimal* strategy.

Theorem 10 (Minimax)**.** *Every game with two players has a value $\hat{L}_i = \hat{H}_i$, for each player $i \in 1, 2$.*

The original formulation of this theorem is due to von Neumann – we will see it in the next section; the proof is not straightforward.

Corollary 11. *Let $v \in \mathbb{R}$. If, for all $\sigma_2 \in \hat{S}_2$, there exists a pure strategy $s_1 \in S_1$ such that $\hat{u}_1(s_1, \sigma_2) \geq v$, then there exists a mixed strategy $\sigma_1 \in \hat{S}_1$ such that, for all $\sigma_2 \in \hat{S}_2$, we have $\hat{u}_1(\sigma_1, \sigma_2) \geq v$.*

2.2 Zero-sum games: von Neumann Theorem

Situations of strict competition are modeled by zero-sum games. These are games with two players, say $i \in \{1, 2\}$, with $u_1(s) = -u_2(s)$ for all outcomes s. That is, one player wins what the other loses. In zero-sum games, the safety and punishment levels of the two players are (additive) inverse to each other, respectively.

Proposition 12. *In every two player zero-sum game, we have:*

$$L_2 = -H_1 \quad \text{and} \quad H_2 = -L_1.$$

As a direct consequence of the Minimax Theorem 10, we obtain the original formulation of the Minimax Principle by von Neumann.

Theorem 13. *(von Neumann) For the mixed extension of any two-player zero-sum game, we have: $\hat{L}_1 = \hat{H}_1 = -\hat{L}_2 = -\hat{H}_2$*

In a two-player zero sum game, if optimal strategies exist, they are also Nash equilibria.

Proposition 14. *Let $\mathcal{G}$ be a finite or regular zero-sum two-player game.*

(i) *$\mathcal{G}$ has a value if, and only if, $\mathcal{G}$ has a Nash equilibrium.*

(ii) *Assume $\mathcal{G}$ has a value. Then, a profile $s \in S$ is a Nash equilibrium if, and only if, s_i is optimal, for every player i.*

(iii) *Assume $\mathcal{G}$ has a value. Then, a profile $s \in S$ is in Nash equilibrium if, and only if, $u_i(s) = L_i$, for every player i.*

For the particular case of mixed extensions of finite games, we know that values, and thus optimal strategies, always exist.

Corollary 15. *Let s be a mixed strategy profile of a finite two-player zero-sum game. Then,*

(i) *s is a Nash equilibrium if, and only if, s_i is optimal, for every player i.*

(ii) *s is in Nash equilibrium if, and only if $u_i(s) = L_i$, for every player i.*

Thus, in the case of two-player zero-sum games, the notion of Nash equilibrium can also be justified in terms of minimax optimality. This is an important robustness argument.

2.3 Computing values and optimal strategies

To compute the value v of a two-player game, together with optimal strategies σ_1, σ_2, let us assume $S_i = \{1, \ldots, n_i\}$ for each player $i \in \{1, 2\}$. We represent strategies $\sigma_1 \in \hat{S}_1$ by a vector $x = (x_1, \ldots, x_{n_1})$ with $x_i = \sigma_1(i)$, and likewise $\sigma_2 \in \hat{S}_2$ by a vector y. Finally, the utility function $u_1(i, j)$ is represented as a $n_1 x n_2$- matrix $(u_{ij})_{(i,j) \in S}$. The maxmin condition defining the safety level can be formulated as the following linear program.

$$\begin{array}{ll} \text{Maximise} & v \\ \text{subject to} & \begin{cases} \displaystyle\sum_{i=1}^{n_1} x_i u_{ij} \geq v & \text{for all } j \in \{1, \ldots, n_2\} \\ \displaystyle\sum_{i=1}^{n_1} x_i = 1 & \\ x_j \geq 0 & \text{for all } j \in \{1, \ldots, n_1\} \end{cases} \end{array}$$

The solution to this LP will yield the value v for player 1 together with a vector x representing an optimal strategy.

3 Nash Theorem

3.1 Equilibrium existence

A fundamental result at the basis of game theory is that Nash equilibria exist in every finite game, provided randomisation is allowed.

Theorem 16 (Nash)**.** *Every finite game has a Nash equilibrium in mixed strategies.*

Proofs of this theorem can be found in many places on the web. They typically rely on Brower's fixed-point theorem and are otherwise not difficult to follow. Interestingly, Brower's theorem can also be derived from a game-theoretic result – the determinacy of the HEX game. (See David Gale, The game of HEX and the Brower Fixed-Point Theorem, American Mathematical Monthly 86(10):818-827, 1979.)

An important argument in this proof, which will be useful for the computation of equilibria is the insight that a mixed strategy is a best response to a profile of (mixed) opponent strategies if, and only if, it puts nonzero weight on those pure strategies that are themselves the best responses to the profile.

Proposition 17. *Let $\sigma \in \hat{S}$ be a mixed strategy profile. The following statements are equivalent, for any $i \in N$:*

(i) σ_i is a best response to σ_{-i};

(ii) a strategy $s_i \in S_i$ is a best response to σ_{-i} if, and only if, $\sigma_i(s_i) \geq 0$.

3.2 Computing equilibria

Assume that the support of (a profile of) equilibrium strategies is given as a subset $\mathrm{support}_i \subseteq S_i$ for every player i. With an appropriate notation, we use Proposition 17 to describe a vector of Nash equilibrium outcomes v and strategies σ_i (described by vectors x_i of dimension n_i, respectively). This leads to a system consisting of the following constraints, for each player $i \in N$.

$$\begin{cases} u_i(s_i^j, x_{-i}) = v_i & \text{for each } s_i^j \in \mathrm{support}_i \\ u_i(s_i^j, x_{-i}) \leq v_i & \text{for each } s_i^j \notin \mathrm{support}_i \\ \displaystyle\sum_{j=1}^{n_i} x_i(j) = 1 & \\ x_i(j) \geq 0 & \text{for each } s_i^j \in \mathrm{support}_i \\ x_i(j) = 0 & \text{for each } s_i^j \notin \mathrm{support}_i \end{cases}$$

For the case $n = 2$, the resulting system is linear and can be solved with LP methods; the coefficient of equilibrium strategies are rational.

However, recall that the construction of the system relies on knowing support sets for the mixed strategy of each player. As these are not given beforehand, one approach is to try all possible support profiles, set up the system and verify whether the solution is correct. This yields an algorithm that is exponential in the strategy space. There exist better approaches, but nevertheless there are strong reasons to believe that the computational complexity of finding Nash equilibrium is intrinsically high (PPAD-complete), and that no tractable solutions may exist already for the two-player case.

For more than two players, precise solutions may not be computable at all. There are examples of three-player games where the only Nash equilibrium has irrational coefficients.

4 Potential games

The previous section suggests that finding equilibria of a game is computationally hard in the general case. In the following, we discuss games where pure-strategy equilibria exist and are moreover easy to calculate.

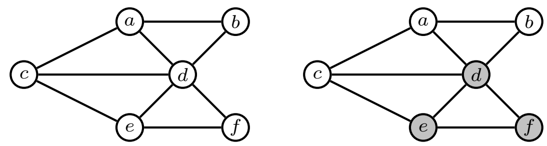

Figure 11: A MaxCut game instance and a play outcome

To start with, let us consider the following game called MaxCut. Such a game is described by a finite undirected graph $G = (V, E)$. Each player is associated with a vertex and has two strategies $\bullet$ and $\circ$. That is, $N = V$ and $S_v = \{\bullet, \circ\}$, for all $v \in V$. The utility of player v is the number of neighbours in G that choose a different strategy:

$$u_v = |\{w \ : \ (v, w) \in E \text{ and } s_v \neq s_w\}|.$$

The story tells that the players in V need to join a team, Black or White and the edges represent conflicts between players. Thus, everyone prefers to have more neighbours in the other team than in his own. The setting can be easily extended to graphs with weights on edges, reflecting how strong the conflict is. Figure 11 represents a game instance and a possible outcome of a play.

Notice that the outcome s of a play induces a cut in the graph via the partition $(V_\bullet, V_\circ)$ with $V_\bullet := \{v \in V \ : \ s_v = \bullet\}$ and $V_\circ := \{v \in V \ : \ s_v = \circ\}$. The *cutsize* of such a partition is the number of edges that link vertices from different sets: $\text{cutsize}(s) = \text{cutsize}(V_\bullet, V_\circ) := |V_\bullet \times V_\circ \cap E|$.

Proposition 18. *For every graph, the associate MaxCut game has a pure equilibrium.*

We show this with arguments of two kinds, a global and a local one.

A global argument. Let (X, Y) be a partition of V such that $\text{cutsize}(X, Y)$ is maximal, and consider the strategy profile s with

$$s_v := \begin{cases} \bullet & \text{if } v \in X; \\ \circ & \text{if } v \in Y. \end{cases}$$

Then, s is a Nash equilibrium: if a player v deviates, say from $\bullet$ to $\circ$, his utility increases by

$$\begin{aligned} u_v(\circ, s_{-v}) - u_v(\bullet, s_{-v}) &= |\{vE \cap V_\bullet\}| - |\{vE \cap V_\circ\}| \\ &= \text{cutsize}(V_\bullet \cup \{v\}, V_\circ \setminus \{v\}) - \text{cutsize}(V_\bullet, V_\circ) \end{aligned}$$

which cannot greater than 0, because $\text{cutsize}(V_\bullet, V_\circ)$ is maximal.

An optimal solution, and thus a pure equilibrium, for our example game is depicted in Figure 12.

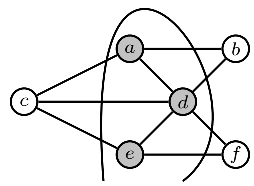

Figure 12: An optimal solution

A local argument. Let us consider Algorithm 1 that starts with an arbitrary partition and moves any node that has more neighbours in his own partition, to the other partition.

We claim that this algorithm terminates in polynomial time with an output that corresponds to a pure Nash equilibrium, and that the cutsize of the output is at most double of the optimal cutsize.

To see this, notice that, whenever a node v is moved, this improves the cutsize of $(V_\bullet, V_\circ)$ by a value corresponding to the improvement in the utility of Player v. Since the size of a cut is at most $|E|$, the algorithm terminates after at most $|E|$ many steps with a partition outcome where no player can improve his utility, i.e., a pure Nash equilibrium. At this point, every node has more neighbours in the other partition than in its own, i.e., more incident edges are in the cut than outside. Hence, the cutsize of the outcome is at least $\frac{|E|}{2}$. As the maximal cutsize is at most $|E|$, this means that the outcome yields guarantees a 2-approximation to a maximal cut.

Algorithm 1: A local improvement algorithm

Input: an undirected graph $G = (V, E)$
Output: a partition $(V_\bullet, V_\circ)$ that corresponds to a Nash equilibrium

Choose $V_\bullet \subseteq V$, arbitrary
$V_\circ = V \setminus V_\bullet$
repeat pick unhappy player and improve his choice
  if *exists* $v \in V_\bullet$ *with* $|vE \cap V_\bullet| > |vE \cap V_\circ|$ **then**
    $V_\bullet = V_\bullet \setminus \{v\}$; $V_\circ = V_\circ \cup \{v\}$
  if *exists* $v \in V_\circ$ *with* $|vE \cap V_\circ| > |vE \cap V_\bullet|$ **then**
    $V_\circ = V_\circ \setminus \{v\}$; $V_\bullet = V_\bullet \cup \{v\}$
until *nochange*

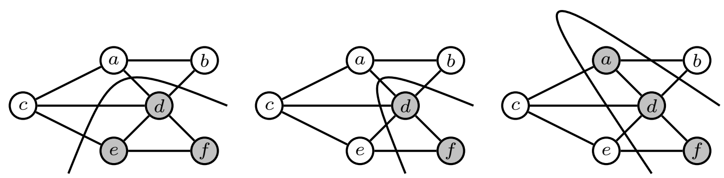

Figure 13: Improvement

4.1 Inefficiency of equilibria

The moral of the MaxCut example is twofold: it shows

(i) a family of games for which (pure) equilibria are easy to find, and

(ii) an optimisation problem that can be solved (up to a certain approximation) by viewing it as a game, that is by distributing it to a set of players and looking at what choices they make to maximise their utility.

When discussing optimisation problems in terms of games, we usually have a global objective in mind. For instance, maximising the *social welfare* – the sum of utilities off all players –, or the *global fairness* – the minimum utility of a player in the game. More generally, an *objective function* is a mapping $f : S \to \mathbb{R}$. The optimisation objective consists in either maximising or minimising this function.

When we model an optimisation problem as a game, we generally assume that the players will play a Nash equilibrium. Here, we restrict our attention to pure equilibria. One basic question is how far from the global optimum such an outcome may be, or, how much efficiency is lost when we let players follow their individual maximisation objective rather than imposing a behaviour that is globally optimal. To quantify this loss, we consider different *coordination ratios*.

The *price of stability* PoS of a game is the ratio between the best equilibrium and the global optimum. For a maximisation objective, e.g.:

$$\text{PoS} = \frac{\max\{f(s) \ : \ s \in S \text{ is a Nash equilibrium }\}}{\max\{f(s) \ : \ s \in S\}}.$$

Intuitively, the price of stability reflect the cost of restringing to solutions from which players have no incentive to deviate.

The *price of anarchy* PoA of a game is the ratio between the worst equilibrium and the global optimum. For a maximisation objective, e.g.:

$$\text{PoA} = \frac{\min\{f(s) \ : \ s \in S \text{ is a Nash equilibrium }\}}{\max\{f(s) \ : \ s \in S\}}.$$

Intuitively, the price of anarchy reflect the cost of letting the players choose a solution (of which we assume that it is an equilibrium) rather than imposing one.

For a successful game approach to optimisation problems, it will be helpful if the underlying games are simple, i.e., if finding equilibria is computationally easy. In the following we study some classes of such games.

4.2 Potential games

A finite game $\mathcal{G}$ is a *potential* game if there exists a function $\phi : S \to \mathbb{R}$ such that for all strategy profiles $s \in S$, all players $i \in N$, and all strategies $s_i' \in S_i$,

$$u_i(s_i, s_{-i}) > u_i(s_i', s_{-i}) \text{ implies } \phi(s_i, s_{-i}) > \phi(s_i', s_{-i}).$$

In this case, we say that ϕ is a potential function. The game is called *exact* potential game, if moreover

$$u_i(s_i, s_{-i}) - u_i(s_i', s_{-i}) = \phi(s_i, s_{-i}) - \phi(s_i', s_{-i}),$$

in which case, we say that ϕ is an exact potential function.

The MaxCut game from the previous section is an exact potential game with the cutsize as a potential function.

The following observation follows from the definition of potential functions.

Proposition 19. *Let $\mathcal{G}$ be a potential game and let ϕ be a potential function. Then, the set of pure Nash equilibria coincides with the set of local minima of ϕ:*

$$\begin{aligned} s \in S \textit{ is a Nash equilibrium} \quad & \textit{if, and only if,} \\ \phi(s_i, s_{-i}) \leq \phi(s_i', s_{-i}), \ & \textit{for all } i \in N, s_i' \in S_i. \end{aligned}$$

Another interesting property of potential games are that individual strategy improvement leads to equilibrium. For a game $\mathcal{G}$, an *improvement path* is a (finite or infinite) sequence $s^1, s^2, \ldots$ such that successive profiles s^k and s^{k+1} are the same except for one player i for which s_i^k is replaced by a better response to s_{-i}^k, that is, a strategy s_i^{k+1} with

$$u_i(s_i^{k+1}, s_{-i}^k) > u_i(s_i^k, s_{-i}^k).$$

Clearly, if an improvement path cannot be prolonged, it ends at a pure Nash equilibrium. It turns out that potential games are precisely those games where all improvement paths are finite.

Theorem 20. *A game $\mathcal{G}$ is a potential game if, and only if, all improvement paths for $\mathcal{G}$ are finite.*

Another useful property is that, for potential games, the price of stability can be bounded in terms of the ratio range between potential and social objective.

Proposition 21. *If $\alpha, \beta \in \mathbb{R}$ are such that*

$$\frac{1}{\alpha} f(s) \leq \phi(s) \leq \beta f(s), \ \textit{for all } s \in S,$$

then $\mathrm{PoS} \leq \alpha\beta$.

4.3 Examples

4.3.1 Load balancing

In a Load Balancing game, we have n players, each of which has a task of a certain duration: $t_1, \ldots, t_n$, respectively. There are m machines and each player can choose a machine to run his task on, i.e., $s_i \in \{1, \ldots, m\}$. The utility of player i is the total load of the machine that he has chosen:

$$u_i(s) = -\sum_{s_j = s_i} t_j.$$

This is a potential game. An example of a potential function is,

$$\phi(s) = \sum_{k=1}^{m} \big(\sum_{s_i = k} t_i\big)^2$$

that is, the sum of squares of machine loads. (Verify this.)

4.3.2 Connection games

A connection game is described by a directed graph $G = (V, E)$ with weighted edges $w : E \to \mathbb{R}_+$ representing non-negative connection costs. There are n players, $i \in \{1, \ldots, n\}$ each of which has a source and a target node: $(s_i, t_i) \in V \times V$. In a play, each player i chooses a path π_i from s_i to t_i. The cost of each edge is shared between the players who choose it; edges that are not taken are discarded. The utility of player i is:

$$u_i(\pi) = -\sum_{e \in \pi_i} \frac{w(e)}{n_e(\pi)}$$

with $n_e(\pi) := \{i \in N \ : \ e \in \pi_i\}$ being the number of players that choose the edge $e \in E$ in π.

As a social objective, we wish to maximise the social welfare $f(\pi) := \sigma_i u_i(\pi)$, which also corresponds to the negative sum of weights of all taken edges.

Figure 14 illustrates a simple routing game for n players with common source and target node . Observe that we have two pure equilibria in this game: all n players choose the upper edge or all choose the lower edge. Thus, the price of anarchy is – PoA $= n$ – this is also an upper bound.

Routing games are potential games. A potential function is

$$\phi(\pi) = \sum_{e \in \pi} w(e)\left(1 + \frac{1}{2} + \cdots + \frac{1}{n_e(\pi)}\right).$$

Exercise: Verify that ϕ is an exact potential function.

The sum $H_n := \sum_1^n \frac{1}{n}$ is the n-th harmonic number. We can verify that, for any connection game, we have

$$f(\pi) \leq \phi(\pi) \leq H_n f(s).$$

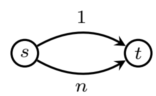

Figure 14: Tight upper bound for PoA

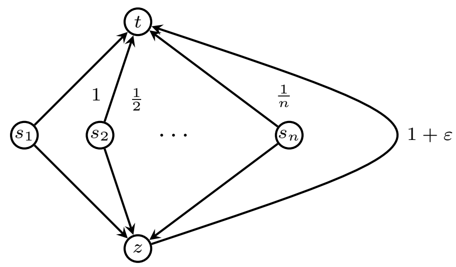

Figure 15: Tight upper bound for PoS

From Proposition 21 it then follows that the price of stability in a connection game for n-players is at most H_n. As H_n grows as fast as $\log n$, this gives us an asymptotic bound for the price of stability of $\log n$. The game in Figure 15 shows that this bound is tight (with ε being a vanishing positive number).

4.3.3 Congestion games

In a congestion game, we have n players $i = \{1, \dots, n\}$ and a set E of resources. For each resource $e \in E$, we have a *latency* function $\ell_e : \{1 \dots, n\} \to \mathbb{R}$. Each player i can choose among certain subsets of resources: $S_i \subseteq 2^E$. For an outcome s, the *congestion* of a resource is the number of players that use it: $n_e = |\{i \ : \ e \in s_i\}|$, and it determines the latency of the resource. Finally, the utility of player i is the (inverse of the) aggregated latency of the resources he chose:

$$u_i(s) = -\sum_{e \in S_i} \ell_e(n_e).$$

Proposition 22. *Every congestion game is a potential game.*

To see this, we take

$$\phi(s) = \sum_{e \in E} \sum_{k=1}^{n_e} \ell_e(k).$$

Surprisingly, the converse holds as well. We call two games equivalent, if they are isomorphic up to removal of duplicate strategies for each player, i.e., strategies that yield the same utility regardless of the choice of other players.

Theorem 23. *For every potential game there exists an equivalent congestion game.*

The original proof was published in (Rosenthal, R. W., A Class of Games Possessing Pure-Strategy Nash Equilibria, Int. Journal of Game Theory (2), 6567, 1973). It is not straightforward, but easy to follow.