Special Topics in Algorithmic Game Theory
Lecture 1 – Supplementary Notes

Yiannis Giannakopoulos*

April 11, 2018

1 Normal Form Games

A (finite) *game* $\mathcal{G} = (N, \{S_i\}_{i\in N}, \{u_i\}_{i\in N})$ consists of

- a finite set of players $N = \{1, 2, \dots, n\}$
- for each player $i \in N$, a finite set S_i of strategies (or actions)
- for each player $i \in N$, a *utility* function $u_i : S_1 \times \dots \times S_n \longrightarrow \mathbb{R}$.

An n-tuple $\mathbf{s} = (s_1, \dots, s_n) \in \mathbf{S} = S_1 \times \dots S_n$ of strategies, one for each player $i = 1, \dots, n$, is called a *strategy profile* or *outcome* of $\mathcal{G}$.

A game $\mathcal{G}$ is called *zero-sum* if, at any outcome $\mathbf{s} \in \mathbf{S}$,

$$\sum_{i\in N} u_i(\mathbf{s}) = 0.$$

A *mixed* strategy σ_i of player i is a probability distribution over her (pure) strategies S_i. Formally, the set Δ_i of mixed strategies of player $i \in N$ is thus defined as

$$\Delta_i = \left\{ p \in [0,1]^{S_i} \;\middle|\; \sum_{s_i \in S_i} p(s_i) = 1 \right\}.$$

We denote by $\Delta = \Delta_1 \times \dots \times \Delta_n$ the set of *mixed* strategy profiles of $\mathcal{G}$.

Notation For any vector $\mathbf{x} = (x_1, \dots, x_n) \in \mathbb{R}^n$ and positive integer $i \leq n$, we use the standard game-theoretic notation

$$\mathbf{x}_{-i} = (x_1, \dots, x_{i-1}, x_{i+1}, \dots, x_n)$$

for the $(n-1)$–dimensional vector that we get from $\mathbf{x}$ if we remove its i-th coordinate. Then, for any $y \in \mathbb{R}$, we can use the following shorthand notation

$$(y, \mathbf{x}_{-i}) = (x_1, \dots, x_{i-1}, y, x_{i+1}, \dots, x_n),$$

so that, for example, $\mathbf{x} = (x_i, \mathbf{x}_{-i})$.

*Chair of Operations Research, Department of Mathematics, TU Munich. Email: yiannis.giannakopoulos@tum.de

1.1 Bimatrix Games

For the special case of $n = 2$, it is common to represent two-player games through their *payoff-matrix* (see, for example, the game examples discussed in the following section). The number of rows of such a matrix is equal to the cardinality of the strategy set S_1 of player 1 (called the *row player*), and the number of its columns to the cardinality of the strategy set S_2 of player 2 (called the *column player*). Each entry of this matrix, is a tuple: its first component gives the corresponding utility of player 1 and its second component that of player 2. Formally, if we enumerate the strategy sets $S_1 = \{s_{1,1}, \dots, s_{1,m_1}\}$, $S_2 = \{s_{2,1}, \dots, s_{2,m_2}\}$, then the (i, j)-th entry of the payoff matrix is $(u_1(s_{1,i}, s_{2,j}), u_2(s_{1,i}, s_{2,j}))$.

2 Solution Concepts

Fix some game $\mathcal{G} = (N, \{S_i\}_{i\in N}, \{u_i\}_{i\in N})$. A (pure) strategy $s_i \in S_i$ will be called (weakly) *dominant* for player i if it maximizes her utility, no matter the actions of the other players. Formally, for any $s'_{-i} \in S_{-i}$ and all $s'_i \in S_i$:

$$u_i(s_i, \mathbf{s}'_{-i}) \geq u_i(s'_i, \mathbf{s}'_{-i}). \tag{1}$$

A profile $\mathbf{s} = (s_1, \dots, s_n)$ is a *dominant strategy equilibrium* of $\mathcal{G}$, if s_i is a dominant strategy of player i, for all players $i \in N$.

A profile $\mathbf{s} = (s_1, \dots, s_n)$ is a *pure Nash equilibrium* of $\mathcal{G}$ if no player has an incentive to unilaterally deviate from it. Formally, for any player $i \in N$ and strategies $s'_i \in S_i$,

$$u_i(\mathbf{s}) \geq u_i(s'_i, \mathbf{s}_{-i}). \tag{2}$$

Similarly, considering the *expected* utilities of the players, a *mixed* strategy profile $\sigma = (\sigma_1, \dots, \sigma_n) \in \varDelta$ will be called *(mixed) Nash equilibrium* if, for any player $i \in N$ and all (pure) strategies $s'_i \in S_i$,

$$\mathbf{E}_{s\sim\sigma}\left[u_i(\mathbf{s})\right] \geq \mathbf{E}_{\mathbf{s}_{-i}\sim\sigma_{-i}}\left[u_i(s'_i, \mathbf{s}_{-i})\right]. \tag{3}$$

It is not difficult to see (verify this on your own!) that the following hierarchy of the sets of equilibria holds for any game:

$$\{\text{dominant strategy}\} \subseteq \{\text{pure Nash}\} \subseteq \{\text{mixed Nash}\}.$$

Example 1 (Prisoners' Dilemma)**.** The following game has a unique (verify!) dominant strategy equilibrium, namely (*confess*, *confess*).

		Prisoner 2: *confess*	Prisoner 2: *silent*
Prisoner 1	*confess*	−5, −5	−1, −10
	silent	−10, −1	−2, −2

Table 1: Prisoner's Dilemma

Example 2 (Battle of the Sexes)**.** The following game has *no* dominant strategy equilibria (check this!). However, it has two pure Nash equilibria, namely (*film*, *film*) and (*football*, *football*). (It also has one additional *mixed* Nash equilibrium; can you find it?)

		Boy	
		film	*football*
Girl	*film*	10, 7	5, 5
	football	1, 1	7, 10

Table 2: Battle of Sexes

2.1 Cost-Minimization Games

So far we have modelled the players as utility-maximizers: their own, selfish goal is to choose a strategy that gives them the *highest* possible "happiness", quantified by their utility function u_i. However, many applications naturally call for a different modelling, where at every outcome of the game there is an actual *cost* induced to each player; now, each player is selfishly trying to *minimize* her own cost, instead of maximizing her utility. For example, this is the case with the Prisoners' Dilemma game described above, and that's why we used negative entries when we modelled the payoff matrix. In such games, we usually use cost functions $C_i : \mathbf{S} \longrightarrow \mathbb{R}$, which can be simply viewed as negative utilities, i.e., $C_i(\mathbf{s}) = -u_i(\mathbf{s})$ for any strategy profile $\mathbf{s}$. Thus, formally we can define a *cost-minimization game* $\mathcal{G} = (N, \{S_i\}_{i \in N}, \{C_i\}_{i \in N})$ as simply being a standard, *utility-maximization* game $\mathcal{G} = (N, \{S_i\}_{i \in N}, \{-C_i\}_{i \in N})$.

All solution concepts described above, can be adapted in the natural way to cost-minimization games, by simply switching the direction of the inequality in their defining conditions (that is, (1), (2) and (3)). For example, condition (2) for the pure Nash equilibrium now becomes

$$C_i(\mathbf{s}) \leq C_i(s_i', \mathbf{s}_{-i}).$$