Game-theoretic models in energy systems and control
DTU Summer School 2018
Modern Optimization in Energy Systems

Maryam Kamgarpour

Contents

1 Introduction

1.1 Motivation

Previously, you saw the electricity procurement from the central operator perspective. The supply and demand offers are given to her and she optimizes the procurement cost while ensuring constraint satisfaction. The other side of this story is that market participants are strategic decision-makers who aim to optimize their profit by choosing their bids. Each participant's profit is a function of all participants submitted bids. To account for this coupling in participants' strategic decision-making, we can use game theory. Strategic decision-making in most power markets is very difficult to analyze, not only because in general, analyzing equilibria in games is a very challenging mathematical problem but also because the nature of power markets make them one of the most difficult classes of games. In particular, due to stochasticity and time-coupling, one has to consider optimal strategic behaviors for a dynamic repeated game with incomplete information. While we do not attempt to address these challenges in a half-day course, our goal is to learn the basic tools to understand the challenges so you can go into more depth based on your specific interests. Let's look at two concrete simple examples, where one gets suboptimal results assuming market participants are price-takers.

Example 1 (Strategic bidding)**.** Consider the LMP market shown in the slide. In this simple example, we see that by deviating from truthful strategies, generator 3 can increase her profit. At the same time, the ISO has to pay more for procuring the required 20 MWh. Here, we need to study auctions as games between strategic generators and examine the *incentive compatibility* property.

Example 2 (Real-time pricing)**.** In a real-time pricing scheme, the price of electricity is a monotone increasing function of the demand. As a very simple model, let $p_t = \sum_{i=1}^{N} x_t^j - r_t$ where r_t denotes the production of renewables at time t and x^j is the power consumption of consumer j. The idea is that price should increase if there is less renewable (we use expensive generators). This should then incentivize consumers to shift consumption to periods with high renewables. For a horizon of length 2, let $r_1 = 0$, $r_2 = 4$ and $N = 10$ consumers with constraint $x_1^j + x_2^j = 1$. If each consumer optimize her production assuming she is a price-taker, then she will shift her entire demand to time step 2. Consequently, each consumer will incur higher price than the case in which she uniformly divides her demand between the two periods. The demand exhibits an undesirable peak.

1.2 Objectives

By the end of the course, you will learn which tools to use to analyze strategic decision-making in power markets. In particular, you will be able to formulate a basic game to model various power market auctions from the participating agents' perspective. You would be able to characterize Nash equilibria under complete information and their Bayesian variants under partial information. Finally, you will be able to compute the Nash equilibrium for a few simple electricity markets.

1.3 Organization

Motivated by the above discussion, in Section 2 we consider continuous action games, characterize their equilibria and present one approach for computing equilibria. In Section 3 we consider auctions and analyze pay-as-bid and Vickrey-Clarke-Groves auction. In Section 4 we consider games under partial information. We will summarize and outline few open problems in Section 5. Homework problems are provided in Section 6. Finally, the basic mathematical tools are given in the Appendix.

2 Games with continuous action domains

We consider games in which each player's decision space is a subset of $\mathbb{R}^n$. We will address existence and uniqueness and computation of Nash equilibria in these games. Let us begin with a warm-up.

2.1 Existence and uniqueness of equilibria

Example 3 (Cournot competition)**.** Consider two electricity producers competing in a market. Player (producer) j decides on the quantity to produce denoted by $x^j \in \mathbb{R}_+$ for $i = 1, 2$, and has a production marginal cost of $c \in \mathbb{R}_+$. The market price $p : \mathbb{R}^2 \to \mathbb{R}$, is a linearly decreasing function of the total production $x^1 + x^2$, that is, $p(x^1, x^2) = a - b(x^1 + x^2)$.

1. Write each producer's loss as a function of the quantities produced.
2. How would each player minimize her loss? Find the production quantity (x^1, x^2) that simultaneously optimizes both losses.
3. How does your answer above change if each producer has a maximum capacity of k units.

Solution. Denote each player's loss by $J^j : \mathbb{R}^2 \to \mathbb{R}$ for $i = 1, 2$.

1. For $i = 1, 2$, $J^j(x) = (c - p)x^j$, where $p = a - b(x^1 + x^2)$.
2. Each player computes $\min_{x^j \in \mathbb{R}} J^j(x)$. Notice that J^j is strongly convex in x^j for each player. Hence, we use the first order optimality condition for this unconstrained optimization problem, that is, $\nabla_{x^j} J^j(x) = c - (a - b(x^1 + x^2)) - bx^j = 0$, for $i = 1, 2$. Hence, $x^j = \frac{a-c}{3b}$.
3. From each player's perspective, we need to consider a constrained optimization problem. Each player computes $\min_{x^j \in [0,k]} J^j(x)$. The first order optimality condition is given by $\nabla_{x^j} J^j(x)(\tilde{x}^j - x^j) \geq 0$ for all $\tilde{x}^j \in [0, k]$ for $i = 1, 2$.

Motivated by the above example, let us formally define our game. We consider an N player game with player set denoted by $\mathcal{N} = \{1, 2, \ldots, N\}$. For $i \in \mathcal{N}$, let $x^j \in K^i \subset \mathbb{R}^n$, where K^i is non-empty, closed and convex. Let $J^j : \mathbb{R}^{Nn} \to \mathbb{R}$ denote player j's objective function. Let $K = K^1 \times K^2 \times \cdots \times K^N \subset \mathbb{R}^{Nn}$. We equivalently represent this game by $\Gamma(\mathcal{N}, K, \{J^j\}_{i\in\mathcal{N}})$.

Definition 1. A point $x \in K$ is a Nash equilibrium for $\Gamma(\mathcal{N}, K, \{J^j\}_{i\in\mathcal{N}})$ if and only if

$$J^j(x^j, x^{-j}) \leq J^j(\tilde{x}^j, x^{-j}), \quad \forall \tilde{x}^j \in K^i, \quad \forall i \in \mathcal{N}. \tag{1}$$

Hence, at a Nash equilibrium, each player chooses a *best-response* to other players' strategies. The Nash equilibrium is stable in the sense that no player has an incentive to unilaterally deviate from her choice. Do Nash equilibria always exist?

Example 4 (Bertrand competition)**.** Consider two producers as before. This time, each producer decides on the price she will charge her consumers for the electricity provided. Hence, $x^j \in \mathbb{R}_+$ for $i = 1, 2$ denotes the price announced by each producer. As before, the production marginal costs are identically $c \in \mathbb{R}_+$. The total demand is one unit and the consumers choose to buy from the producer with the lowest price. Furthermore, if both firms declare the same price, then half of the demand chooses firm 1 and the other half chooses firm 2.

1. Write each producer's profit as a function of the price they charge.

2. Derive the Nash equilibrium.
3. What is the Nash equilibrium if each producer has the capacity to serve maximum 2/3 of the unit demand?

Solution. We adopt the same notation as in the previous example.

1. The losses are given by
$$J^1(x) = \begin{cases} c - x^1, & \text{if } x^1 < x^2 \\ \frac{c-x^1}{2}, & \text{if } x^1 = x^2 \\ 0, & \text{if } x^1 > x^2. \end{cases}$$
The cost $J^2(x)$ is defined in exactly the same manner.

2. We can verify that the only Nash equilibrium is $(x^1, x^2) = (c, c)$. To see this, assume without loss of generality, that player 1 announces a price $x^1 > x^2$. Then, all consumers will go to player 2 and thus, player 1 will make no profit. She has incentive to decrease her price until it reaches x^2. If $x^2 > c$, then she can still decrease x^1 and attract all consumers, hence making a profit. On the other hand, for any $x^1 < c$, she will make a loss rather than a profit. By symmetry, the same argument holds for player 2. Hence, at (c, c), no player can increase her profit by unilaterally deviating from her strategy.

3. Let's write the loss functions again
$$J^1(x) = \begin{cases} \frac{2(c-x^1)}{3}, & \text{if } x^1 < x^2 \\ \frac{c-x^1}{2}, & \text{if } x^1 = x^2 \\ \frac{c-x^1}{3}, & \text{if } x^1 > x^2. \end{cases}$$
Here, we can see that there is no Nash equilibrium. In particular, if $x^1 = x^2$ then player 1 has an incentive to decrease its cost to get more of the demand. Now if $x^1 < x^2$, then player 2 has incentive to decrease its cost. Continuing this way, we arrive at both players choosing $x^1 = x^2 = c$. When $(x^1, x^2) = (c, c)$, each player gets half of the demand and makes zero loss. However, they still have incentive to increase their cost because with a slight increase of the cost they can make negative loss. Note that a mixed strategy Nash equilibrium exists. In this case, player can choose a price based on a probability distribution.

The above motivates us to ask under which conditions on the game $\Gamma(\mathcal{N}, K, \{J^j\}_{i\in\mathcal{N}})$ we can ensure existence of Nash equilibria. The following theorem from [2] provides one such sufficient condition. Note that the theorem is valid for a more general case in which the action spaces of the agents may be coupled, that is, $(x^1, \dots, x^N) \in K \subset \mathbb{R}^{Nn}$.

Theorem 1. *Suppose $K \subset \mathbb{R}^{Nn}$ is compact convex and J^j are continuous in $x \in K$ and convex in x^j for fixed x^{-j}, $\forall j = 1, \dots, N$. Then, $\Gamma(\mathcal{N}, K, \{J^j\}_{i\in\mathcal{N}})$ has a Nash equilibrium.*

Proof. Provided in [2] and we recall it for completeness. Define $\rho(x, \tilde{x}) = \sum_{i=1}^{N} J^j(x^1, \dots, \tilde{x}^j, \dots, x^N)$. Hence, $\rho : K \times K \to \mathbb{R}$. By continuity of J^j, ρ is continuous on $K \times K$ and by convexity of J^j in x^j for fixed x^{-j}, ρ is convex in $\tilde{x}$ for fixed x. Consider $\Omega(x) = \{\tilde{x} \mid \rho(x, \tilde{x}) = \min_{z\in K} \rho(x, z)\}$. In general Ω is a set-valued map, that is, it maps each point in K to a subset of K. Given that K is convex, ρ is convex in z and ρ is continuous, it follows that the map $x \mapsto \Omega(x)$ is lower semi-continuous (see The Minimum Theorem in the appendix) and $\Omega(x) \subset K$ is compact for each $x \in K$. By Kakutani's fixed point theorem (see appendix), there exists $x_0 \in K$ such that $x_0 \in \Omega(x_0)$. We can verify that this x_0 is a Nash equilibrium as follows. Notice that $\rho(x_0, x_0) = \min_{z \in K} \rho(x_0, z)$. We show that $J^j(x_0^j, x_0^{-j}) \leq J^j(\tilde{x}^j, x_0^{-j})$ for each j and for all $\tilde{x}^j \in K$. To see this, suppose that there exists j and $\tilde{x}_0 = (\tilde{x}^j, x_0^{-j}) \in K$ such that $J^j(\tilde{x}_0) = J^j(\tilde{x}^j, x_0^{-j}) < J^j(x^i, x_0^{-j})$. Then $\rho(x_0, \tilde{x}_0) < \rho(x_0, x_0)$. This however contradicts $x_0 \in \Omega(x_0)$. Hence, x_0 is a Nash equilibrium by definition. □

Exercise 1. Do the Cournot and Bertrand models previously introduced meet the assumptions above for existence of Nash equilibria?

Solution. In the Cournot competition, $J^j(x^1, x^2) = -px^j + c^i x^j$ and J^j is continuous in $x = (x^1, x^2)$ and convex in x^j. In the first part of the exercise, without capacity constraints, the decision space $K = \mathbb{R} \times \mathbb{R}$ is not compact. Nevertheless, we were able to uniquely characterize the Nash equilibrium. This highlights that the above conditions are only sufficient and not necessary. In the second part, the action space is $K = [0, k] \times [0, k]$, which is also compact. Hence, there exists a Nash equilibrium. In the Bertrand competition, the cost functions are not continuous. We are not able to guarantee existence of Nash equilibrium from the theorem above.

For the rest of this section, let us assume $K = K^1 \times \cdots \times K^N \subset \mathbb{R}^{Nn}$, with each K^j compact and convex and that J^j is continuous in $x \in K$ and convex in x^j for fixed x^{-j}. These assumptions ensure existence of Nash equilibria. We now ask the question of how can we compute the Nash equilibria? In general, this is a difficult task. However, if J^j are differentiable in x^j, we can establish a link between convex games and convex optimization to answer this question.

2.2 Connection to optimization

Our approach is to connect Nash equilibria to the solution of a variational inequality problem.

Definition 2. Let $K \subset \mathbb{R}^n$, $F : K \to \mathbb{R}^n$. The *variational inequality* problem $\text{VI}(K, F)$ consists of finding an $x \in K$ such that

$$F^T(x)(y - x) \geq 0, \quad \forall y \in K. \tag{2}$$

The set of solutions of $\text{VI}(K, F)$ is denoted by $\text{SOL}(K, F)$.

Assume J^j is differentiable as a function of x^j, with a continuous derivative denoted by $\nabla_{x^j} J^j : \mathbb{R}^{Nn} \to \mathbb{R}^n$.

Exercise 2. Given $\Gamma(\mathcal{N}, K, \{J^j\}_{i \in \mathcal{N}})$, define the game map $F_\Gamma : \mathbb{R}^{Nn} \to \mathbb{R}^{Nn}$ as

$$F_\Gamma = (\nabla_{x^1} J^1; \nabla_{x^2} J^2; \ldots; \nabla_{x^N} J^N). \tag{3}$$

Show that x is a Nash equilibrium for $\Gamma(\mathcal{N}, K, \{J^j\}_{i \in \mathcal{N}})$ if and only if $F_\Gamma(x)^T(y - x) \geq 0$ for all $y \in K$. In other words, x is a Nash equilibrium if and only if x is a solution of a variational inequality problem (see Definition 2 in the appendix).

Solution. At a Nash equilibrium $x \in \mathbb{R}^{Nn}$, $J^j(x^j, x^{-j}) \leq J^j(\tilde{x}^j, x^{-j})$, $\forall \tilde{x}^j \in K^i$ and $\forall i \in \mathcal{N}$. Hence, $x^j = \arg\min_{\tilde{x}^j \in K^i} J^j(\tilde{x}^j, x^{-j})$, $\forall i \in \mathcal{N}$. From convexity of J^j in x^j and Exercise 19 it follows that $\nabla_{x^j} J^j(x^j, \tilde{x}^j)^T(\tilde{x}^j - x^j) \geq 0$, $\forall \tilde{x}^j \in K^i$. Summing together the inequalities corresponding to each j, we obtain $\sum_{i=1}^N \nabla_{x^{-j}} J^j(x^j, x^{-j})^T(\tilde{x}^j - x^j) \geq 0$. This is equivalent to $(\nabla_{x^1} J^1; \nabla_{x^2} J^N; \ldots; \nabla_{x^N} J^N)^T(y - x) \geq 0$ for all $y \in K$ (verify this). Hence, $F_\Gamma(x)^T(y - x) \geq 0$ as desired.

As the above result shows, to study Nash equilibria in a convex game we can equivalently study the solutions to a variational inequality problem corresponding to the game map. Since variational inequalities are well-studied, we can then address existence and uniqueness of Nash equilibria for $\Gamma(\mathcal{N}, K, \{J^j\}_{i\in\mathcal{N}})$ by studying the same questions for variational inequalities. There are several sufficient conditions for existence and uniqueness of solutions to a variational inequality problem, as discussed in Chapter 2 of [4]. The one we will be using is based on *monotone operators*.

Definition 3. The map $F : D \to \mathbb{R}^n$ is *monotone* if $\langle F(x) - F(y), x - y\rangle \geq 0$ for all $x, y \in D \subseteq \mathbb{R}^n$. If the above inequality is strict for all $x \neq y$, then F is *strictly monotone*.

First, we look at a few characterization of monotone maps. Then, we will see how using monotonicity, we can conclude results about existence and uniqueness of solutions to $\mathrm{VI}(K, F)$.

Exercise 3. Consider the linear map defined by the matrix multiplication $A : x \mapsto Ax$, $A \in \mathcal{S}^n$. Under which conditions on the matrix A is this operator monotone? strictly monotone?

Solution. We need $\langle Ax - Ay, x - y\rangle \geq 0$ for all $x, y \in \mathbb{R}^n$. This is equivalent to $z^T Az \geq 0$ for all $z \in \mathbb{R}^n$, which holds if $A \in \mathcal{S}^n_+$. Similarly, for strict monotonicity, we need $A \in \mathcal{S}^n_{++}$. If A is not symmetric, then this condition is equivalent to the symmetric part of A being positive semidefinite.

Exercise 4. Let $f : \mathbb{R}^n \to \mathbb{R}$ be continuously differentiable and convex. Show that $\nabla f : \mathbb{R}^n \to \mathbb{R}^n$ is monotone. If in addition f is strictly convex, show that ∇f is strictly monotone.

Solution. We use the the equivalent characterization of convex differentiable functions in Equation (13). Note that from convexity of f we have, for every $x, y \in \mathbb{R}^n$

$$f(y) \geq f(x) + \nabla f(x)^T (y - x),$$
$$f(x) \geq f(y) + \nabla f(y)^T (x - y).$$

If we add the above two inequalities, we obtain $(\nabla f(y) - \nabla f(x))^T (y - x) \geq 0$ as desired.

For a map $F : \mathbb{R}^n \to \mathbb{R}^n$, let $JF := \left(\frac{\partial F_i}{\partial x_j}\right) \in \mathbb{R}^{n\times n}$ denote the Jacobian of F.

Proposition 1. *Let $F : D \to \mathbb{R}^n$ be continuously differentiable on the open convex set $D \subset \mathbb{R}^n$. Then, F is monotone on D if and only if $JF(x)$ is positive semi-definite for all $x \in D$. F is strictly monotone on D if and only if $JF(x)$ is positive definite for all $x \in D$.*

Finally, we state our existence and uniqueness result.

Proposition 2. *Let $F : \mathbb{R}^n \to \mathbb{R}^n$ be a continuous operator. Let $K \subset \mathbb{R}^n$ be compact and convex. Then, $SOL(K, F)$ is non-empty and compact. Furthermore, if F is strictly monotone then $SOL(K, F)$ is a singleton (the solution to the $VI(K, F)$ is unique).*

For the proof, please see Chapter 2 of [4]. We can readily link the above result to the existence and uniqueness of Nash equilibria for our convex game.

Corollary 1. $\Gamma(\mathcal{N}, K, \{J^j\}_{j\in\mathcal{N}})$ has a unique Nash equilibrium if K^j are convex and compact, J^j is continuously differentiable in x^j and the game map F_Γ is strictly monotone.

Exercise 5 (Cournot competition continued)**.** Let us return to Example 4. Consider the case in which each player's decision space is $[0, k^i] \subset \mathbb{R}$, $i = 1, 2$. Characterize the Nash equilibrium using the variational inequality framework. Is there a unique Nash equilibrium in this setting?

Solution. Verify that $F_\Gamma(x) = [\nabla_{x^1} J^1(x^1, x^2); \nabla_{x^2} J^2(x^1, x^2)]$ is strictly monotone. Hence, this game has a unique Nash equilibrium, which is the solution of $\mathrm{VI}(K, F_\Gamma)$, with $K = [0, k^1] \times [0, k^2]$.

2.3 Computing Nash equilibria

We now present an approach to find the Nash equilibrium for a game satisfying assumptions of Corollary 1. Before presenting the approach, consider the following questions. It might be helpful to keep in mind the Cournot and the Bertrand competitions.

1. Given a convex game $\Gamma(\mathcal{N}, K, \{J^j\}_{i\in\mathcal{N}})$ how can each player compute its Nash equilibrium strategy?
2. Suppose each player does not know the other players' constraint sets or cost functions. How does your answer above change?

Our approach for computing the Nash equilibrium will be to design an iterative algorithm $x_{t+1} = f(x_t)$, whose fixed point $x = f(x)$, would be the Nash equilibrium. Here, x would be the stacked vector of decisions of the players. If we can ensure f is a contraction, then we are sure that the above iteration converges to the fixed point of f and hence, the Nash equilibrium. This is due to the well-known *Banach fixed point Theorem*, see Theorem 2 in the appendix.

Consider a game $\Gamma(\mathcal{N}, K, \{J^j\}_{i\in\mathcal{N}})$, which satisfies the conditions for existence and uniqueness of Nash equilibrium. In particular, assume $K = K^1 \times \dots K^N$ is compact convex, J^j are continuous in all variables, convex and continuously differentiable in each x^j for fixed x^{-j} and F_Γ, the game mapping, is strictly monotone. We consider a standard projected gradient descent approach for our iterative functions f. Given a starting point x_0 at time $t = 0$, the next iterates are obtained by $x_{t+1} = \Pi_K(x_t - \gamma F_\Gamma(x_t))$. The function Π_K is a projection operator, that is, it takes an element x and returns its projection to the compact convex constraint set K. The parameter $\gamma > 0$ is a step size that we choose. Recall the definition of the projection operator.

Definition 4. Given a convex set $K \subset \mathbb{R}^n$, the *projection operator* $\Pi_K : \mathbb{R}^n \to K$ is defined as

$$\begin{aligned} \Pi_K(x) := \arg\min_y \quad & \frac{1}{2}(y-x)^T(y-x) \\ \text{s.t.} \quad & y \in K. \end{aligned} \tag{4}$$

The interpretation of the above optimization problem in the decision variable y is that the projection operator Π_K maps any $x \in \mathbb{R}^n$ to an $\bar{x} \in K$, which has the minimum Euclidean distance to x.

Exercise 6. Let $\bar{x} = \Pi_K(x)$. Show that

$$(y - \bar{x})^T(\bar{x} - x) \geq 0, \quad \forall y \in K.$$

Solution. Let $f(y) = \frac{1}{2}(y-x)^T(y-x)$ denote the objective function above. Notice that this is a convex function in the decision variable y with $\nabla_y f(y) = y - x$. If $\bar{x}$ is an optimum for the optimization problem, from convexity of $f(y)$ it follows that

$$\nabla f(\bar{x})^T(y - \bar{x}) = (\bar{x} - x)^T(y - \bar{x}) \geq 0, \quad \forall y \in K.$$

Hence, we have the desired result.

Exercise 7. Show that for any $\gamma > 0$

$$x \in \text{SOL}(K, F) \iff x = \Pi_K(x - \gamma F(x)). \tag{5}$$

Solution. Recall that $x \in \mathrm{SOL}(K,F) \iff (y-x)^T F(x) \geq 0, \quad \forall y \in K$. From Exercise 6 we have $x = \Pi_K(x - \gamma F(x))$ if and only if

$$(y-x)^T (x - (x - \gamma F(x))) \geq 0, \quad \forall y \in K \iff (y-x)^T F(x) \geq 0, \quad \forall y \in K.$$

From the above, we conclude that the Nash equilibria in $\Gamma(\mathcal{N}, K, \{J^j\}_{i\in\mathcal{N}})$ can be characterized as fixed points of the operator $x \mapsto \Pi_K(x - \gamma F_\Gamma(x))$. The question of computing Nash equilibria is reduced to the computation of fixed points of the operator $x \mapsto \Pi_K(x - \gamma F_\Gamma(x))$. We will design γ so that this map is a contraction. To do so, first we show that $\|\Pi_K(g(x)) - \Pi_K(g(y))\|_2 \leq \|x-y\|_2$. It then follows that if $g : x \mapsto x - \gamma F_\Gamma(x)$ is a contraction with respect to the 2−norm, then so is the composition map $\Pi_K \circ g : x \mapsto \Pi_K(x - \gamma F_\Gamma(x))$ (verify this).

Exercise 8. Show that $\Pi_K : \mathbb{R}^n \to K$ is a non-expansive operator, that is,

$$\|\Pi_K(y) - \Pi_K(x)\|_2 \leq \|y - x\|_2, \quad \forall x, y \in \mathbb{R}^n.$$

Solution. Let $\bar{x} = \Pi_K(x)$. Then, from Exercise (6) we have that

$$(y - \bar{x})^T (\bar{x} - x) \geq 0, \quad \forall y \in K.$$

Since the above holds for any $x \in R^n$ and $y \in K$, by letting $u, v \in \mathbb{R}^n$ arbitrary and choosing $x = v$ and $y = \Pi_K(u)$, we have that

$$(\Pi_K(u) - \Pi_K(v))^T (\Pi_K(v) - v) \geq 0.$$

Repeating the same steps but with u and v interchanged, we obtain

$$(\Pi_K(v) - \Pi_K(u))^T (\Pi_K(u) - u) \geq 0.$$

Now, if we add the above two inequalities, we obtain

$$\begin{aligned} &\Pi_K(u)^T \Pi_K(v) - \Pi_K(v)^T \Pi_K(v) - \Pi_K(u)^T v + \Pi_K(v)^T v \\ &+ \Pi_K(v)^T \Pi_K(u) - \Pi_K(v)^T u - \Pi_K(u)^T \Pi_K(u) + \Pi_K(u)^T u \geq 0. \end{aligned}$$

Grouping terms, we obtain

$$\|\Pi_K(u) - \Pi_K(v)\|_2^2 \leq (\Pi_K(u) - \Pi_K(v))^T (u - v). \tag{6}$$

Now, from Cauchy-Schwarz inequality, we have

$$|(\Pi_K(u) - \Pi_K(v))^T (u - v)| \leq \|\Pi_K(u) - \Pi_K(v)\| \|u - v\|.$$

Finally, dividing the above by $\|\Pi_K(u) - \Pi_K(v)\|_2$ and combining with Inequality (6), we obtain the desired result.

Now, to ensure $x \mapsto x - \gamma F_\Gamma(x)$ can be a contraction we make the following assumptions on F_Γ.

Assumption 1. Let $K \subset \mathbb{R}^n$, $F : K \to \mathbb{R}^n$, and let $L, \mu \in \mathbb{R}_{>0}$ be such that $\forall x, y \in K$

$$(F(x) - F(y))^T (x - y) \geq \mu \|x - y\|_2^2 \tag{7}$$

$$\|F(x) - F(y)\|_2 \leq L \|x - y\|_2. \tag{8}$$

Note that the first assumption is satisfied when F_Γ is strictly monotone. The second assumption is requiring that the game mapping is Lipschitz continuous.

Proposition 3. *Under Assumption 1, if* $0 < \gamma < \frac{2\mu}{L^2}$*, then the iteration*

$$x_{t+1} = \Pi_K(x_t - \gamma F_\Gamma(x_t)), \tag{9}$$

converges to the Nash equilibrium of the game $\Gamma(\mathcal{N}, K, \{J^j\}_{i\in\mathcal{N}})$.

Solution. We already showed that the fixed point of $x \mapsto \Pi_K(x - \gamma F_\Gamma(x))$ is the Nash equilibrium of the game. Hence, we now simply verify the conditions for the map above to be a contraction. In the derivation below, all norms are the 2−norm.

$$\begin{aligned}
&\|\Pi_K(x - \gamma F_\Gamma(x)) - \Pi_K(y - \gamma F_\Gamma(y))\|^2 \le \|x - \gamma F_\Gamma(x) - (y - \gamma F_\Gamma(y))\|^2 \\
&= \|x - y\| + \gamma^2\|F_\Gamma(x) - F_\Gamma(y)\|^2 - 2\gamma(F(x) - F(y))^T(x - y) \\
&\le (1 + \gamma^2 L^2 - 2\mu\gamma)\|x - y\|^2,
\end{aligned}$$

where the first inequality above is due to the non-expansion of the projection operator, the second equality is simply writing out the 2−norm squared, the second inequality is due to Assumption 1. Finally, we see that $0 < (1 + \gamma^2 L^2 - 2\mu\gamma) < 1$ if and only if $0 < \gamma < \frac{2\mu}{L^2}$ as desired.

Notice that to implement the above algorithm, the j-th player needs to know only K^j (its own constraint set) as well as $\nabla_{x^j} J^j(x)$, which is the gradient of its own cost function with respect to its decision variables. To evaluate the gradient term $\nabla_{x^j} J^j(x_t)$, the player may need some information about strategies x_t^{-j}. In several games such information can be inferred from the game data. For example, in the Cournot competition, each producer only needs to know the total production of all firms rather than individual firms production to compute this term. Another potential source for coordination in the above algorithm is that all players are using the same step size γ. The latter requirement is not so stringent and one can design a convergent algorithm as long as players' individual step sizes satisfy certain conditions to ensure the operator is contractive.

Exercise 9. Consider the Cournot game introduced at the beginning of the previous lecture.

1. Propose a step size γ to ensure convergence of the iterations (9) to the Nash equilibrium.
2. What does each player need to know about the game to implement this algorithm?
3. How would your answer change if the players cannot coordinate choosing the step sizes γ?
4. What can you say about the convergence speed of the algorithm?

We briefly note that there are other approaches for solving the Nash equilibrium problem. For example, as you already learned, one can use the Karush-Kuhn-Tucker first-order optimality condition in case of convex games [24]. Furthermore, we can consider coupling constraints, that is, the case where K is not Cartesian product of individual constraints K^j, $j = 1, \ldots, N$. In this case, the problem is called a *generalized Nash equilibrium* problem. It is very difficult to ensure uniqueness of Nash equilibria in the presence of coupling constraints. However, one can find a subset of Nash equilibria through formulating the variational inequality problem [25]. Note that several electricity markets might exhibit both individual constraints K^j and coupling constraints. Nevertheless, it is possible to design decentralized algorithms, where the participants can keep their strategy sets and objectives private, while a coordinator broadcasts sufficient aggregate information and dual variables to ensure the coupling constraints are satisfied upon convergence [5, 7].

3 Auctions, mechanism design

This is a very brief introduction to auction theory and practice. We discuss the first-price (also known as pay-as-bid) and second-price auctions using game theory. We then consider auctions of multiple possibly divisible items. We discuss desired properties of a mechanism and discuss a generalization of the second-price auction known as the Vickrey-Clarke-Groves (VCG) mechanism.

3.1 First-price and second-price sealed-bid auctions

Consider a reverse auction motivated by auctions arising in electricity markets. Note that in a reverse auction, the auctioneer wants to purchase an item. Hence, it picks the bid with the lowest value (in a forward auction on the other hand, the auctioneer is selling an item and hence, she will give the item to the player with the highest bid.) In a first-price reverse auction, the player with the lowest bid wins and gets paid the price she bids. In a second-price auction, the bidder with the lowest bid wins (as before) but gets paid the price of the second lowest bid.

Exercise 10. Form groups of three. Two of you are electricity providers and one is the system operator. The operator wants to purchase 1 MW capacity. As the operator, you decide which auction you will run and you will announce it to the participants. As the providers, decide how much you will offer for the 1 MW capacity (do not discuss with the other provider). Write your offer price on a piece of paper and pass it on to the auctioneer. The auctioneer decides the winner and the payment. Discuss your choices afterwards.

As you notice in the exercise, from the players' perspective it is difficult to determine a bidding strategy to maximize the profit. From the auctioneer's perspective, it is difficult to know how to set the auction rule to maximize the auction's revenue. The latter is an instance of *mechanism design*. We can use game theory to analyze both problems. In particular, if we could analyze Nash equilibria of different mechanisms, then we could figure out optimal bidding strategies, profits of each participant and the total auctioneer's revenue.

3.2 Equilibria in auctions

Let us fix the number of auction participants (players) to N. Player j's true (private) value of the item is denoted by $t^j \in \mathbb{R}_+$ and her bid is denoted by $x^j \in \mathbb{R}_+$. The vector $x = [x^1, x^2, \dots, x^N]$ denotes the bids of all players and x^{-j} denotes the vector x with the player j's bid removed. The set of all players is denoted by $\mathcal{N}$. The payment made by each player to the auctioneer (or by the auctioneer to each player in case of a reverse auction) is denoted by $p^j \in \mathbb{R}$. The utility (also referred to as payoff or profit) of the winner(s) is given by $J^j = t^j - p^j$ in a forward auction and $J^j = p^j - t^j$ in a reverse auction. For non-winners, $J^j = 0$. If participant j bids $x^j = t^j$ we say she is bidding truthfully. If however, $x^j < t^j$ she is underbidding, whereas if $x^j > t^j$ she is overbidding.

Exercise 11. Formulate the first-price and second-price reverse auctions discussed in the above exercise each as a game. What are the strategy sets? What are the payoffs?

Solution. Simply define the winner selection and the payment rules as a function of the bids. *First-price auction*: the winner $j^\star \in \mathcal{N}$ is found as $j^\star = \arg\min_{j \in \mathcal{N}} x^j$. The winner is paid its bid price $p^{j^\star} = x^{j^\star}$, while other participants are paid zero. So, the utility of the winner $j^\star$ is $J^{j^\star} = t^{j^\star} - x^{j^\star}$, while other players have zero utility.

Second-price auction: the winner $j^\star$ is found as before, that is, $j^\star = \arg\min_{j \in \mathcal{N}} x^j$. The winner in this case is paid the second highest price: $p^{j^\star} = \min_{j \neq j^\star} x^j$, while other players are paid zero. So, the utility of the winner is $J^{j^\star} = t^{j^\star} - \min_{j \neq j^\star} x^j$.

Definition 5. A bidding profile x is a Nash equilibrium if $\forall j$

$$J^j(x^j, x^{-j}) \geq J^j(\tilde{x}^j, x^{-j}), \quad \forall \tilde{x}^j \in \mathbb{R}_+.$$

A stronger notion of Nash equilibrium is a dominant strategy Nash equilibrium. It says that no matter what the other players choose, it is best to play x^j.

Definition 6. A bidding strategy is a dominant strategy Nash equilibrium if $\forall j$

$$J^j(x^j, x^{-j}) \geq J^j(\tilde{x}^j, x^{-j}), \quad \forall \tilde{x}^j,\ \forall x^{-j}.$$

Clearly, for a participant to make profit in a first-price reverse auction she must overbid (in forward auction she must underbid). Indeed, it is not hard to verify that if she overbids by an amount only slightly lower than the second lowest price, she will win the auction while making more profit. We are interested in auctions which incentivize participants to bid truthfully.

Proposition 4. *Truthful bidding is a dominant strategy Nash equilibrium in a second-price auction.*

Proof. Please note that the proof is written for the normal auction (not the reverse auction). We show that both overbidding and underbidding strategies are weakly dominated by truthful bidding.

- Case 1 - overbidding, $x^j > t^j$. We have the following three possibilities:
  - $\max_{i \neq j} x^i < t^j$: overbidding and truthful bidding both result in winning the auction with the same payoff.
  - $\max_{i \neq j} x^i > x^j$: overbidding and truthful bidding both result in losing the auction with zero payoff.
  - $t^j < \max_{i \neq j} x^i < x^j$: overbidding results in winning the auction but a negative payoff $J^j = t^j - \max_{i \neq j} x^i < 0$; truthful bidding results in losing the auction and having zero loss.
- Case 2 - underbidding, $x^j < t^j$. We have the following three possibilities:
  - $\max_{i \neq j} x^i < x^j$: underbidding and truthful bidding both result in winning the auction with the same payoff.
  - $\max_{i \neq j} x^i > t^j$: underbidding and truthful bidding both result in losing the auction with zero payoff.
  - $x^j < \max_{i \neq j} x^i < t^j$: underbidding results in losing the auction; truthful bidding results in winning the auction with $J^j = t^j - \max_{i \neq j} x^i > 0$ payoff.

We just showed that the utility of player j does not improve by deviating from truthful bidding. Since j and x^{-j} were arbitrary, this implies that truthful bidding is a dominant strategy Nash equilibrium. Intuitively, this property holds because $\forall j \in \mathcal{N}$, the payment to the winner, namely $\max_{i \neq j} x^i$, does not depend on the winner's bid x^j.

□

As we can see, the second-price auction has a beautiful property that each player's profit is maximized by bidding truthfully. This property is referred to as *incentive compatibility*. It seems it could be a good idea to use this mechanism for electricity markets. Let us consider a slightly generalized example for a reverse auction motivated by electricity markets.

Exercise 12. A user wants to purchase 800 KW of electrical power. There are three providers. Player 1 offers 400 KW at 25 CHF. Player 2 offers 600 KW at 30 CHF and 400 KW at 24 CHF conditional bids, meaning that only one of the two bids can be accepted. Player 3 offers 400 KW for 26 CHF and 200 KW for 18 CHF conditional bids. We assume that we can only buy the whole amounts offered from each user (the bids are not divisible). How do we generalize the second-price auction to this setting?

To be able to address the above example and more complex scenarios with players submitting multiple bids or bid curves, and in the presence of coupling constraints on the amount of electricity to purchase, we consider a more general auction setting.

3.3 VCG mechanism

We consider auctions where each participant can submit multiple bids. A bid in this general setting consists of a price vector c^j and the item (or amount in case of divisible bid) the bidder is buying (or selling in a reverse auction) denoted by m^j. Let N^j denote the number of bids of player j. Hence, $x^j = (c^j, m^j)$, where $c^j \in \mathbb{R}^{N^j}$ and $|m^j| = N^j$. Let $M = \sum_{j=1}^{N} N^j$. Let $x = [x^1, x^2, \ldots, x^N]$ as before denote the decision of all players. Note that a generalization to bid curves $c^j : \mathbb{R}_+ \to \mathbb{R}_+$ is straightforward. We do not discuss it here for simplicity in notation, but you can find it in [16].

A mechanism is defined by a *choice function* $\delta : x \mapsto \{0,1\}^M$ and a *payment rule* $p : x \mapsto \mathbb{R}^M$. The choice function indicates which bids are accepted, in particular, those that get mapped to 1. The payment rule determines the payment made by (or given to, in a reverse auction) each player. Given $\delta \in \{0,1\}^M$, the total payment (or cost in a reverse auction) is $c^T\delta$ (the super-index T is denoting transpose of a vector, whereas the super-indices j, i denote a particular player (j, i respectively) bids. There may exist some constraints on the choice function. In Exercise 12 for example, the constraint was to obtain 800 KW of electrical power. Hence, $m^T\delta = 800$. Furthermore, the bids from each player were mutually exclusive, $\mathbf{1}^T\delta^j \leq 1$, where $\mathbf{1}$ is a $\mathbb{R}^{N^j}$ dimensional vector of unit entries. We write such constraints in general as $(\delta, m) \in \mathcal{C}$, where $\mathcal{C}$ denotes a constraint set. Note that we consider constraints only on the amount of bids m_j's and not on the prices c_j's.

Example 5. Formulate the setup in Exercise 12 as a reverse auction using the notation above.

There are certain properties we may require from a mechanism. For example, the bidders should be incentivized to submit truthful bids. This ensures a *socially efficient* choice, when the operator optimizes the total procurement cost (the true costs are being minimized). The Vickrey-Clarke-Groves mechanism can achieve this goal in the above general auction setting.

The Vickrey-Clarke-Groves auction mechanism is a generalization of the second-price auction to the cases of multiple possibly non-divisible items. It is named after three economists who successively (1961, 1971, 1973) generalized the idea of the second-price auction. To define the VCG mechanism in our current setting, we consider the objective function which chooses the bids that satisfy the constraint while procuring items from bidders who offer the best price $J^\star(x) = \max_{(\delta,m)\in\mathcal{C}} c^T\delta$ and the optimizer is denoted by $\delta^\star(x)$. If bidders are being truthful, then this social choice function enjoys the *social welfare maximization* property. Let $J^\star(x^{-j})$ denote the optimal cost when the bid of player j is removed from the optimization problem. It follows that $J^\star(x^{-j}) \leq J^\star(x)$ in a forward auction ($J^\star(x^{-j}) \geq J^\star(x)$ in a reverse auction). Note that any result below that is defined with respect to a forward auction can be trivially extended to reverse auctions by replacing max with min in the optimization and by appropriately defining the utilities of players with respect to payments received (instead of payments made).

Definition 7. The VCG mechanism choice function and payment rule are defined as

$$\delta^\star(x) = \arg\max_{(\delta,m)\in\mathcal{C}} c^T\delta \tag{10}$$

$$p^j(x) = J^\star(x^{-j}) - (J^\star(x) - c^{jT}\delta^{\star j}), \quad \forall j \in \mathcal{N}. \tag{11}$$

Exercise 13. Show that for a single non-divisible item forward auction, similar to the one considered in the first example, the VCG mechanism is equivalent to the second-price auction.

Solution. Notice that in a single item auction with each player submitting a single bid, we can consider x^j to be the price the player offers for the item (no need to introduce m^j).
Let $c^j = x^j$ based on the above observation. The VCG choice function (10) is equivalent to

$$\delta^\star(x) = \arg\max_{\delta\in\{0,1\}^N} c^T\delta = \arg\max_j c^j,$$

which is the choice function corresponding to the second-price auction. For the payment rule, consider two cases of player j being the winner or not. In the first case, $\delta^{\star j}(x) = 1$ for player j and $\delta^{\star i}(x) = 0$ for all other players $i \neq j$. Consequently, $J^\star(x) = c^{jT}\delta^{\star j}$. Hence, the payment rule (11) reduces to $p^j(x) = J^\star(x^{-j})$. Note that $J^\star(x^{-j}) = \max_{i\neq j} x^j$. Hence, we have the same payment as the second-price auction. In the case in which j is a losing bidder, $x^{\star j}(x) = 0$. Furthermore, $J^\star(x^{-j}) = J^\star(x)$. Consequently, we have that $p^j(x) = 0$, consistent with the second-price auction.

Exercise 14. Determine the VCG payments in Exercise 12.

Solution. The winners are player 2 for 600 CHF and player 3 for 200 CHF. To determine the payment made to each winner, we use formula (11):

$$\begin{aligned} p^2(x) &= J^\star(x^{-2}) - (J^\star(x) - c^{2T}x^{\star 2}) \\ &= 51 - (48 - 30) = 33 \\ p^3(x) &= J^\star(x^{-3}) - (J^\star(x) - c^{3T}x^{\star 3}) \\ &= 49 - (48 - 18) = 19. \end{aligned}$$

The following result is a generalization of the incentive-compatibility property of the second-price auction shown in Proposition 4.

Proposition 5. *Truthful bidding is the dominant strategy Nash equilibrium in the VCG mechanism.*

Proof. For an arbitrary player j consider the truthful bid $t^j = (c^j, m^j)$ and any other bid $\tilde{x}^j = (\tilde{c}^j, \tilde{m}^j)$. Let t denote the bids where player j places her truthful bid and other players place any arbitrary bid x^{-j}. On the other hand, $\tilde{x}$ denotes the bid profile where player j places $\tilde{x}^j$ and other players place x^{-j}. Let $J^\star(t)$ and $\delta^\star(t)$ denote the optimal value and optimizer corresponding to t. Similarly, define $J^\star(\tilde{x})$ and $\delta^\star(\tilde{x})$ as the optimal value and the optimizer corresponding $\tilde{x}$. Recall the VCG payment rule (11). So, player j's payoff under truthful bidding is

$$\begin{aligned} J^j(t) &= c^{jT}\delta^{\star j}(t) - p^j(t) \\ &= c^{jT}\delta^{\star j}(t) - \big(J^\star(x^{-j}) - (J^\star(t) - c^{jT}\delta^{\star j}(t)\big) \\ &= J^\star(t) - J^\star(x^{-j}). \end{aligned}$$

The payoff under the bid $\tilde{x}$, on the other hand is

$$\begin{aligned}
J^j(\tilde{x}) &= c^{jT}\delta^{\star j}(\tilde{x}) - p^j(\tilde{x}) \\
&= c^{jT}\delta^{\star j}(\tilde{x}) - \big(J^\star(x^{-j}) - (J^\star(\tilde{x}) - \tilde{c}^{jT}\delta^{\star j}(\tilde{x})\big) \\
&= c^{jT}\delta^{\star j}(\tilde{x}) - \big(J^\star(x^{-j}) - \sum_{i\neq j} c^{iT}\delta^{\star i}(\tilde{x})\big) \\
&= c^{jT}\delta^{\star j}(\tilde{x}) + \sum_{i\neq j} c^{iT}\delta^{\star i}(\tilde{x}) - J^\star(x^{-j}).
\end{aligned}$$

Now, notice that $c^{jT}\delta^{\star j}(\tilde{x})+\sum_{i\neq j} c^{iT}\delta^{\star i}(\tilde{x}) \leq J^\star(t)$ because $\delta^\star(\tilde{x})$ is a feasible suboptimal allocation for the bid profile t (it is optimal for the bid profile $\tilde{x}$). Hence, we have $J^j(\tilde{x}) \leq J^j(t)$ as desired. Since the proof was for an arbitrary player j and for any x^{-j}, we conclude that truthful bidding is the dominant strategy Nash equilibrium. □

Exercise 15. Verify that the above proposition is true if $J^\star(x^{-j})$ in (11) is replaced by any other positive function $h(x^{-j})$. The main insight here is that the payment should not depend on player j's bid to ensure incentive compatibility. The choice of $h(x^{-j}) = J^\star(x^{-j})$ is referred to as the Clarke Pivot rule. This choice ensures an additional desirable property of the mechanism, which an arbitrary h may not do. In particular, using $J^\star(x^{-j})$ ensures the minimum total payment from (to in a reverse auction) each player that ensures players face a non-negative utility (price recovery). This is referred to as individual rationality because it ensures players have incentive to participate in the auction (see Proposition 6).

Proposition 6. *At the dominant strategy incentive compatible Nash equilibrium of the VCG mechanism, each player obtains a non-negative utility.*

Proof. Note that $J^\star(x^{-j}) \leq J^\star(x)$ because the former maximization has one fewer player. The utility under truthful bidding is $J^j(x) = J^\star(x) - J^\star(x^{-j})$. It follows that $J^j(x) \geq 0$. □

The VCG mechanism also ensures the utility of all participants including the auctioneer is maximized. This property is referred to as social welfare optimization and is shown here.

Proposition 7. *At the dominant strategy incentive compatible Nash equilibrium of the VCG market, the sum of utilities of all players (including the auctioneer) is maximized.*

Proof. The utility of player j under truthful bidding is $J^j(t) = c^{jT}\delta^{\star j}(t) - \big(J^\star(t^{-j}) - (J^\star(t) - c^{jT}\delta^{\star j}(t)\big)$. The utility of the auctioneer is $J^0(t) = \sum_{j=1}^N \big(J^\star(t^{-j}) - (J^\star(t) - c^{jT}\delta^{\star j}(t)\big)$. Adding these utilities, we obtain $J^0(t) + \sum_{j=1}^N J^j(t) = J(t) = \max_{(\delta,m)\in\mathcal{C}} c^T\delta$.

□

Unfortunately, the VCG mechanism suffers from several drawbacks.

Example 6. Consider purchase of 800 KW of electrical power as before. There are five providers. Player 1 offers 800 KW at 50 CHF. The next four players each offer 200 KW for 0 CHF. What is the VCG outcome? What is the total payment made by the auctioneer?

The above highlights some of the pathologies of the VCG mechanism. In particular, note that the auctioneer ends up paying a very high price despite the fact that it is possible to procure electricity at a much lower price. Furthermore, the example shows shill bidding and collusion can occur in a VCG mechanism and that the auctioneer's total payment is non-monotone as a function of the number of bidders. Shill bidding means a bidder can submit bids with multiple identities. Collusion means bidders can join forces to increase their utilities. Finally, payoff monotonicity means that as the number of players increases (more competition in the market), each players' utilities should not increase. For a detailed example of these pathologies please see [12].

Some of the problems of the VCG mechanism can be avoided by ensuring the VCG utilities are in a set referred to as the *core*. The core is a concept from coalitional game theory. Unfortunately, one needs to place additional assumptions on the bids or constraint function $\mathcal{C}$ in a general auction to ensure the outcome of the VCG mechanism are in the core [16]. Alternatively, one can define payment rules that are immune to collusion and shill bidding or ensure monotonicity of the payoffs. However, these mechanisms in general may not ensure incentive compatibility. Furthermore, since the theoretical analysis becomes very challenging for such general mechanisms, one has to resort to heuristic approaches and simulations to understand equilibria and revenues.

4 Games with incomplete information

We saw that in a second-price auction, the dominant strategy Nash equilibrium is to bid truthfully. What about in a first-price auction? What is a Nash equilibrium bidding strategy?

Exercise 16. Let $t^1 > t^2 > \cdots > t^N$ denote the true valuations of bidders $1, 2, \dots, N$ for an item being auctioned. Verify that in a first-price auction, a Nash equilibrium strategy is given by $x^1 = t^2$ and $x^j = t^i$ for $i \neq 1$. In words, the bidder with the highest valuation of the item should bid the second highest price and other bidders should bid truthfully.

Solution. We can verify that no player has incentive to unilaterally deviate from the above strategies. Clearly, for any $x^1 > t^2$ player 1 still wins the auction but has to pay higher amount so she has a lower utility. For $x^1 < t^2$ player 1 no longer wins the auction. For all other players, for any $x^j > t^i$ the players could potentially win the auction but will have negative utility. However, for $x^j \leq t^i$, they will have zero utility.

Obviously, the players do not know others' valuation of the items. They have *incomplete information* about the game. Hence, they cannot compute the Nash equilibrium strategy above. One may then only consider auctions such as second-price or VCG whose dominant strategy Nash equilibrium is truthful bidding. However, such mechanisms in general suffer from other shortcomings. Alternatively, we can try to analyze games with incomplete information.

We will consider the Bayesian approach to analyze games with incomplete information. In this approach, all the unknowns in the model are captured by an uncertainty, and a prior distribution on the uncertainty set is assumed. The original formulation of Bayesian games is due to [19]. Ever since, the topic has become standard in most game theory courses and books [20, 11, 15].

Definition 8 (Bayesian game)**.** A Bayesian game consists of N players.

- A set $K = K^1 \times \cdots \times K^N$ of action spaces, with K^j being the action space for player j.
- A set of types $T = T^1 \times \cdots \times T^N$. Player j's type is $t^j \in T^j$ and is known only to her.
- A prior probability distribution D on the set of types T.
- A set of utilities $(J^1, \dots, J^N)$, where $J^j : T \times K \to \mathbb{R}$ is player j's utility.

		player 2	
		cooperate	defect
player 1	cooperate	2, 2	0, 3
	defect	3, 0	1, 1

Figure 1: Payoff matrix if player 2 is selfish

In the Bayesian formulation above, we assume the unknowns in the game are captured by the "type" of players. And all players have the same common prior distribution D about the types (this can be generalized to consider players having different priors on the types).

In a Bayesian game, each player must choose its action $x^j \in K^j$ not knowing others' types t^{-j}, but only the distribution D over all types. How should she select x^j? The idea is to associate an action for each type of the player. We define a pure strategy $s^j : T^j \to K^j$ as a map from player j's type to its action space. Given that player j knows her type t^j, she can use Bayes' rule[1] to compute the probability of others' types conditioned on her type, namely, $D^j := D(t^{-j}|t^j)$, so D^j is the distribution of types of other players, conditioned that type of player j is t_j. Her expected utility then is given by

$$E_{D^j} J^j(t^j, t^{-j}, s^j(t^j), s^{-j}(t^{-j})) := \sum_{t^{-j}} D^j(t^{-j}|t^j) J^j\big(t^j, t^{-j}, s^j(t^j), s^{-j}(t^{-j})\big).$$

Note that in the above, we assume the set of types T has finite cardinality. The generalization to uncountable types would simply replace the sum above with an integral.

We will use the expected utility defined above to define a Bayesian Nash equilibrium strategy:

Definition 9 (Bayesian Nash equilibrium)**.** A strategy profile $\{s^j\}_{j=1}^N$ with $s^j : T^j \to K^j$ is Bayesian Nash equilibrium if each s^j is a best-response strategy to s^{-j} for all possible types t^j, that is,

$$E_{D^j} J^j(t^j, t^{-j}, s^j(t^j), s^{-j}(t^{-j})) \geq E_{D^j} J^j(t^j, t^{-j}, x^j, s^{-j}(t^{-j})), \tag{12}$$

for all $x^j \in K^j$, $t^j \in T^j$, and for all $j \in \{1, \ldots, N\}$.

Example 7. Let us consider a simple two-player game, with a payoff matrix (utilities rather than losses) given as in Figure 1. In this table, the green colors correspond to the utility of the first player and the blue colors correspond to the utility of the second player. The first player chooses the row of the matrix and the second player chooses the column of the matrix. Furthermore, both players are maximizing their utility. Verify that (defect, defect) is the Nash equilibrium in this case.

Now consider the situation where player 2 can have two types. She can be either selfish or nice. In the selfish case, the payoff matrix is as per Figure 1. In the nice case, the payoff matrix is as per Figure 2. Player 1 does not know the type of player 2. How shall we define the concept of Nash equilibrium? One way to handle this uncertainty is to assume a probability for each type of player 2. We can say that with probability d player 2 is selfish and with probability $1 - d$ player 2 is nice. Determine the Bayesian Nash equilibrium.

[1]recall that for two events A, B, Bayes' rule gives probability of event A conditioned on event B as follows: $P(A|B) = \frac{P(A \cap B)}{P(B)}$

		player 2	
		cooperate	defect
player 1	cooperate	3, 3	1, 2
	defect	2, 1	0, 0

Figure 2: Payoff matrix if player 2 is nice

Solution. First, note that $T^2 = \{\text{selfish}, \text{nice}\}$ and player 1 does not have types. Hence, we look for a strategy for player 2 for each of her types and an action for player 1. What is the dominant strategy of player 2 if she is selfish? what about the case when she is nice?
You can verify that if player 2 is selfish, her dominant strategy is to defect (D) and if she is nice, her dominant strategy is to cooperate (C). Note that this dominant strategy is $s^2 : T^2 \to \{C, D\}$. In other words, it is a strategy for each of player 2's type. Player 2 knows her type and accordingly will play her dominant strategy. Player 1 on the other hand has only one type. Hence, we need to determine its optimal action (not type-dependent). She also does not know the type of player 2, only the probability of her being nice or selfish. The expected utility of player 1 from cooperating is $d \times 0 + (1-d) \times 3 = 3 - 3d$, and her expected utility from defecting is $d \times 1 + (1-d) \times 2 = 2 - d$. Hence, player 1 should cooperate for $d \leq 12$ and should defect otherwise.

Exercise 17. Consider a first-price auction with two players, call them Alice and Bob. Alice's true valuation of the item is $a \in [0, 1]$ and Bob's true valuation of the item is $b \in [0, 1]$. We assume these true valuations are independently and uniformly distributed on the interval $[0, 1]$. Verify that $s^1(a) = \frac{a}{2}$ and $s^2(x) = \frac{b}{2}$ is a Bayesian Nash equilibrium bidding strategy for Alice and Bob, respectively. What is the Bayesian Nash equilibrium in this setting under the second-price auction?

Solution. Without loss of generality (due to symmetry), let us compute the best response strategy of Alice considering that Bob bids $b/2$. The utility of Alice is $J^a = a - x$ if she wins the bid and $J^a = 0$ otherwise. Hence, the expected utility of Alice is

$$E_{D^a} J^a(a, b, x, \frac{b}{2}) = E_{D^a}\{(a - x)\mathbf{1}_{z>\frac{b}{2}}(x)\},$$

where $\mathbf{1}_C(x)$ is the indicator function of the set C. Since x and a are known to Alice, the only probabilistic variable is b, the type of Bob. Hence, expected utility of Alice is $(a-x) \times E_{D^a}\{\mathbf{1}_{z>\frac{b}{2}}(x)\}$. Now, since a and b are independent, $D(b|a)$ becomes the probability density of b. Hence, expected utility of Alice becomes equivalent to $(a - x) \times$ (Probability that $x > \frac{b}{2}$). In words, $(a - x) \times$ Probability that Alice wins). This was a lengthy discussion for a statement that may seem intuitive. Nevertheless, the point was to practice using the framework introduced, since the framework will become handy for more complex setups.

Let's compute probability that Alice wins. This is the probability that $x > b/2$, which is equal to the probability of $b < 2x$. Considering that b is uniformly distributed on $[0, 1]$, this probability is given as follows:

$$P(b < 2x) = \begin{cases} 0 & \text{if } x \leq 0, \\ 2x & \text{if } 0 \leq x \leq 1/2, \\ 1 & \text{if } x \geq 1/2. \end{cases}$$

Since utility of Alice is $a - x$ if she wins, it follows that she must compare the maximum of $a - x$ for $x \geq 1/2$ or the maximum of maximize $(a - x)2x$ subject to $0 \leq x \leq 1/2$. The first maximization gives the utility of $a - 1/2$ for $x = 1/2$ and the second one gives $a^2/2$ for $x = a/2$. From comparing these two results, it follows that the optimal solution is given by $x = a/2$. Hence, Alice's best response to Bob's bid of $b/2$ is to bid $a/2$. By symmetry, the same argument holds for Bob. Hence, $(a/2, b/2)$ is the Bayesian Nash equilibrium strategy.

For the second-price auction, we already saw that bidding truthfully is dominant strategy Nash equilibrium for each player. The fact that this is a dominant strategy, means that each player has no incentive to deviate from truthful bidding regardless of types (and hence corresponding actions) of the other players. Hence, Alice would bid a, while Bob would bid b.

The above analysis can be generalized to the first-price auction with N players, where each player's valuation is independently and identically distributed. The analysis becomes a little more involved, but there is closed form solution for the uniform distribution.

From the auctioneer perspective, which auction generates more profit? In our simple two-player example with uniform distribution of valuations on $[0, 1]$, in the case of first-price auction, the auctioneer's profit is $\max\{a/2, b/2\}$, with a and b uniformly distributed. In the second price auction, since the Nash equilibrium strategy is bidding truthfully, then the auctioneers revenue is $\min\{a, b\}$ (since the highest bidder gives the second high price).

Exercise 18. Verify that the expected profit of the auctioneer in both first-price and second-price auction under the two player setup above is $1/3$.

Solution. Hint: Start by computing the expected value of $c = \max\{a/2, b/2\}$ and $c = \min\{a, b\}$ for a, b being uniformly distributed on $[0, 1]$.
Let us start with the first-price function. Let $z = \max\{a, b\}$. Consider the cumulative distribution function of the random variable z, denoted by $F(z \leq y)$. Now, $z \leq y \iff a \leq y \wedge b \leq y$, and since a and b are independent, $F(z \leq y) = P(a \leq y)P(b \leq y) = y^2$. Hence, the density of this function of this random variable is $2y$ and that of c is y. It follows that expected value of c is $\int_0^1 y^2 y dy = 1/3$. Similarly, you can compute the density of the random variable $c = \min\{a, b\}$. Here, it might be easier to compute $F(c \geq y) = 1 - F(c \leq y)$, since $c \geq y \iff a \geq y \wedge b \geq y$. Then, verify that expected value of this random variable is $1/3$.

The above result is an instance of the celebrated *Revenue Equivalence Principle*. This theorem shows that for all mechanisms satisfying certain assumptions the expected utility of the auctioneer at a Bayesian Nash equilibrium is the same. For more details, see Chapter 9 of [11] and the original paper of Myerson [21]. Note that in practice, players may not have independent and identically distributed valuations. However, analyzing equilibria in auctions under partial information is a very challenging problem. This problem has been receiving increasing attention from the computer science community due to online auctions.

5 Summary and further reading

There are several topics we have not had a chance to discuss in this course. In the references, we suggest additional text books for game theory as well as references for auctions. We will briefly comment on two of the most relevant research topics.

5.1 Learning Nash equilibria

Related to the incomplete information setting, a very active research topic is how do players learn Nash equilibria in iterative games. Note that if players cannot learn to play Nash equilibria, then any theoretical analysis of outcome of the game based on the Nash equilibrium concept is questionable. In an auction such as those in electricity markets or for adverts on internet, players have generally less information than that assumed by a Bayesian game. In particular, they do not know how many other players there are, their strategy spaces, and sometimes they don't even fully understand the auction being run, specially in the case when there are multiple items or indivisible items and complex constraints (recall electricity market procurement constraints). So, players cannot compute their Bayesian Nash equilibrium strategy. Nevertheless, these auctions are repeated over and over. In each run, players are gathering information about the game based on their past observations. How should they use these observations and update their strategies?

This is the topic of learning Nash equilibria with limited information. In general, first one needs to assume the specific feedback players receive after each iteration of the game. Then, to devise a learning rule and prove that this rule converges to an equilibrium. Both latter steps are very difficult and sometimes one can only answer these questions only in simulation. Under assumption of convexity of the game, this topic is addressed in [6]. For more general games with dynamic population, you can read [23].

5.2 Mechanism design

From the auctioneer or system operator's perspective, how should the market be designed to maximize social welfare? In order to do so, we need to incentivize bidders to bid truthfully. It is known that the VCG mechanisms are the only individually rational dominant strategy incentive compatible and efficient mechanisms. We have already discussed the issue with collusion and shill bidding. Another serious problem is the budget deficit in two-sided VCG markets. This is the case where both consumers and producers submit bid and the auctioneer ends up with a budget loss. In fact, there are also impossibility results for achieving budget-balance in a two-sided market, while ensuring incentive compatibility and individual rationality, see [22]. Given that it is theoretically impossible to ensure all desired properties of a mechanism, we might need to relax certain properties of the mechanism, such as incentive compatibility. The analysis of modified mechanism becomes extremely challenging. It might be better to resort to simulation to understand equilibrium strategies and the auction revenue/loss [26].

6 Problem sets

Problem 1 (Potential games)**.** Consider the convex game $\Gamma(\mathcal{N}, K, \{J^j\}_{i\in\mathcal{N}})$. Suppose there exists a single function $P : \mathbb{R}^{Nn} \to \mathbb{R}$, such that $\nabla_{x^j} P(x) = \nabla_{x^j} J^j(x)$, for all $i \in \mathcal{N}$.

1. Show that finding the Nash equilibrium is reduced to a single optimization problem.
2. Suppose additionally the game satisfies the conditions we derived for existence and uniqueness of Nash equilibrium. Propose an alternative algorithm using the potential function to compute the Nash equilibrium of the game.

Note: In this case, the game is referred to as a *potential game*. This class of games has been widely studied due to its applicability (economics, traffic, etc.). Furthermore, the potential function enables several classes of provably convergent algorithms for computing the equilibria [8].

Problem 2 (An aggregative game)**.** In this problem, we consider an extension of the Cournot game we discussed in the previous lecture. We have N players. The decision space of each player is denoted by $K^i \subset \mathbb{R}^n$. We assume player j's cost function is given by $J^j(x) = p(\sum_{j=1}^N x^j)x^j$, where $p : \mathbb{R}^n \to \mathbb{R}^n$ is a price function. Consider the case in which p is linear $p(x) = Cx + c$, with $C \in \mathcal{S}^n$ and $c \in \mathbb{R}^n$.

1. Derive conditions on C to ensure existence and uniqueness of the Nash equilibria.
2. Assume the conditions for existence and uniqueness of the Nash equilibrium as you derived above is satisfied. Derive a matrix A such that the iterative algorithm based on the skewed projection operator $\Pi_{K,A}$ converges to a Nash equilibrium.
3. Show that the above is a potential game.

Note: Since p and hence, each player's objective function, depends on the sum of all players' strategies, this falls in the class of the so-called *aggregative games*.

Problem 3 (Computation of Nash equilibria)**.** Consider the setting of Problem 2. Let $N = 10$, $n = 3$, $K^i = \{x \in \mathbb{R}^3 \,|\, 0 \le x \le 1, \mathbf{1}^T x = 1\}$, where $\mathbf{1} \in \mathbb{R}^n$ is the vector with unit entries. Let $C = I_3$, where I_n is the identity matrix in $\mathbb{R}^{n\times n}$. Choose $c \in \mathbb{R}^3$ randomly.

1. Design an algorithm that converges to a Nash equilibrium. Plot the iterates of your algorithm.
2. Suppose each player wants to compute the iterates x^j_{t+1} without knowing strategies of other players and still ensure x_t converges to a Nash equilibrium. Would that be possible? If yes, explain how. If not, what is the minimum information each player needs for local computation?
3. Now, consider the cost function $\theta(x) := \sum_{i=1}^N J^j(x)$. Compute the solution of the optimization problem $\min_{x\in K} \theta(x)$. Denote the solution by x^o. Compute PoA $:= \frac{\theta(x^o)-\theta(x^\star)}{\theta(x^o)}$, where $x^\star$ denotes the Nash equilibrium computed in the previous part.
   Note: PoA measures the distance between a socially optimal solution versus a Nash equilibrium and is referred to as *price of anarchy*.
4. Repeat the previous step for a few increasing values of N, such as $N = 100, 300, 1000$. What do you observe about the price of anarchy as N increases? What do you observe about convergence speed of your algorithm as N increases?

7 Appendix

We define some notation and basic background result. Let $x^j \in \mathbb{R}^n$. We denote $x = (x^1; x^2; \ldots; x^N) \in \mathbb{R}^{Nn}$ as the stacked vector; $x^{-j} \in \mathbb{R}^{(n-1)N}$ is constructed from x by removing the j-th vector from the stack. We let $\mathcal{S}^n_+$ and $\mathcal{S}^n_{++}$ denote the set of symmetric positive semi-definite and positive definite matrices in $\mathbb{R}^{n\times n}$, respectively. Recall that a positive definite matrix $A \in \mathcal{S}^n_{++}$ has n real positive eigenvalues $\lambda_i \in \mathbb{R}^n_{>0}$. Hence, the eigenvalues can be ordered. In this case, we denote the smallest and largest eigenvalues of A by $\underline{\lambda}$ and $\bar{\lambda}$, respectively.

A *norm* on $\mathbb{R}^n$ (in general on a vector space over a field) is a function $\|.\| : \mathbb{R}^n \to \mathbb{R}_{\ge 0}$ that satisfies three axioms: (1) $\|ax\| = |a|\|x\|, \forall a \in \mathbb{R}$, (2) $\|x+y\| \le \|x\| + \|y\|$, (3) $\|x\| = 0 \iff x = 0$. The Euclidean or $2-$norm is $\|x\|_2 := (x^T x)^{1/2} = (x_1^2 + x_2^2 + \cdots + x_n^2)^{1/2}$. For any $A \in \mathcal{S}^n_{++}$, $\|x\|_A := (x^T A x)^{1/2}$ is a well-defined norm (verify this statement). We denote the inner product $\langle x, y\rangle_A = x^T A y$. From the Cauchy-Schwarz inequality it follows that $|\langle x, y\rangle_A| = |x^T A y| \le \|x\|_A \|y\|_A$.

A set $D \subset \mathbb{R}^n$ is *convex* if $\forall x, y \in D$, $\forall t \in [0,1]$, $tx + (1-t)y \in D$. A function $f : D \to \mathbb{R}$ is convex if D is a convex set and $\forall x, y \in D$, $\forall t \in [0,1]$,

$$f(tx + (1-t)y) \leq tf(x) + (1-t)f(y).$$

If in the above, the inequality is strict $\forall t \in (0,1)$, then the function is *strictly convex*. A differentiable function $f : D \to \mathbb{R}$ is convex if and only if D is convex and

$$f(y) \geq f(x) + \nabla f(x)^T (y - x), \quad \forall x, y \in D. \tag{13}$$

If in the above, the inequality is strict $\forall t \in (0,1)$, then the function is *strictly convex*. A twice differentiable function $f : D \to \mathbb{R}$ is convex if and only if D is convex and its Hessian $\nabla^2 f(x)$ is positive semi-definite $\forall x \in D$, and is strictly convex if $\nabla^2 f(x)$ is positive definite $\forall x \in D$ [1].

Exercise 19. Let $f : \mathbb{R}^n \to \mathbb{R}$ be a continuously differentiable convex function, with derivate $\nabla f : \mathbb{R}^n \to \mathbb{R}^n$. Let $K \subset \mathbb{R}^n$ be a convex set. Consider the following optimization problem

$$\begin{aligned} \min \quad & f(x) \\ \text{s.t.} \quad & x \in K. \end{aligned}$$

Show that $x^\star$ is an optimal solution above if and only if $x^\star \in \text{SOL}(K, \nabla f)$.

Solution. First, suppose $x^\star \in \text{SOL}(K, \nabla f)$, which is equivalent to $\nabla f(x^\star)^T (y - x^\star) \geq 0$, $\forall y \in K$. This, combined with Inequality (13) implies that $f(y) \geq f(x^\star)$, $\forall y \in K$. Hence, $x^\star$ is a minimum. To show that if $x^\star$ is minimum then, $\nabla f(x^\star)^T (y - x^\star) \geq 0$, $\forall y \in K$, we can prove by contradiction. In particular, suppose there exists $z \in K$ such that $\nabla f(x^\star)^T (z - x^\star) < 0$. We can then show that for sufficiently small $t \in (0,1]$, $f(tz + (1-t)x^\star) < f(x^\star)$, which contradicts optimality of $x^\star$, since $tz + (1-t)x^\star$ is a feasible solution (by convexity of set K) with a lower objective [1].

Theorem 2 (Banach fixed point theorem)**.** *Let $(C, \|.\|)$ be a Banach space (complete metric space). Let $f : C \to C$. Suppose that f is a contraction, that is, there exists $\tau \in [0,1)$ such that*

$$\|f(x) - f(y)\| \leq \tau \|x - y\|, \quad \forall x, y \in C. \tag{14}$$

Then, f has a unique fixed point, that is, there exists a unique $x^\star \in C$ such that $x^\star = f(x^\star)$. Furthermore, starting from any $x_0 \in C$, the iteration

$$x_{t+1} = f(x_t),$$

converges to $x^\star$.

Recall that a Banach space is a complete normed vector space, for example, $\mathbb{R}^n$ and any compact subsets of $\mathbb{R}^n$ are Banach spaces. Note that in the definition of the contraction (14) the norm is not specified. It is sufficient to show the contraction property for any well-defined norm on C.

The following theorem is a generalization of Brower's fixed point theorem to set-valued maps with a compact domain.

Fact 1. [Kakutani's fixed point theorem] Let $K \subset \mathbb{R}^n$ be convex compact. Let $\Omega : K \to 2^K$ be a lower semicontinuous map whose range is convex closed subsets of K. Then, there exists $x_0 \in K$ such that $x_0 \in \Omega(x_0)$.

The following theorem provides conditions on well-behavedness of set of minimizers of a function as a parameter changes. It is used in the proof of the theorem above on existence of Nash equilibria. It also has applications in optimal control theory. Note that the statement of the theorem can be generalized but the following version suffices for our problem.

Fact 2. [The Minimum theorem] Let $K \subset \mathbb{R}^n$ be compact and $f : X \times Y \to \mathbb{R}$ be continuous on $X \times Y$, and convex in Y for each fixed x. Then, for each point $x \in X$, $\Omega(x) = \arg\min_{y \in K}\{f(x, y) \mid y \in K\} \subset Y$ is lower semicontinuous and $\Omega(x) \subset Y$ is a compact convex set.

References

[1] S. Boyd and L. Vandenberghe, *Convex Optimization*, Cambridge University Press, 2004.

[2] J. B. Rosen, *Existence and uniqueness of equilibrium points for concave n-person games*, Econometrica: Journal of the Econometric Society, 1965.

[3] S. Kakutani, *A generalization of Brouwer's fixed point theorem*, Duke mathematical journal, 8(3), 457-459, 1941

[4] F. Facchinei and J. Pang, *Finite-dimensional variational inequalities and complementarity problems*, Springer Science & Business Media, 2007.

[5] D. Paccagnan, M. Kamgarpour and J. Lygeros, *On Aggregative and Mean Field Games with Applications to Electricity Markets*, Proceedings of the European Control Conference, 2016.

[6] T. Tatarenko and M. Kamgarpour, *Payoff-Based Approach to Learning Generalized Nash Equilibria in Convex Games*, to appear in IEEE Transactions on Control Systems, available at `https://arxiv.org/abs/1703.04113`, 2018.

[7] B. Gentile, F. Parise, D. Paccagnan, M. Kamgarpour and J. Lygeros, *Nash and Wardrop equilibria in aggregative games with coupling constraints*, to appear in IEEE Transactions of Automatic Control Problem, `https://arxiv.org/abs/1702.08789`, 2018.

[8] D. Monderer and S. Shapley. *Potential games*, Games and economic behavior 14.1 (1996): 124-143.

[9] G. Scutari, S. Barbarossa, and D. P. Palomar, *Potential games: A framework for vector power control problems with coupled constraints*, IEEE International Conference on Acoustics Speech and Signal Processing, volume 4, pages 241-244, May 2006.

[10] D. Bertsimas, Dimitris, V. Gupta, and I. Paschalidis, *Data-driven estimation in equilibrium using inverse optimization*, Mathematical Programming 153, no. 2 (2015): 595-633.

[11] N. Nisan, T. Roughgarden, E. Tardos, V. V. Vazirani, eds., 2007. *Algorithmic game theory.* Cambridge university press.

[12] L. M. Ausubel and P. Milgrom, *The lovely but lonely Vickrey auction*, Combinatorial auctions, 17, 22-26, 2006.

[13] P. Cramton, *Spectrum auction design*, Review of Industrial Organization, 42(2), 161-190, 2013.

[14] D. Easley and J. Kleinberg, *Networks, Crowds, and Markets: Reasoning About a Highly Connected World*, `http://www.cs.cornell.edu/home/kleinber/networks-book/`, 2010.

[15] Y. Shoham and K. Leyton-Brown, *Multiagent systems: Algorithmic, game-theoretic, and logical foundations*, Cambridge University Press, 2008.

[16] O. Karaca, P. G. Sessa, N. Walton, and M. Kamgarpour, *Designing coalition-proof reverse auctions over continuous goods*, `https://arxiv.org/abs/1711.06774`, 2017.

[17] E. Wolfstetter, *The Swiss UMTS spectrum auction flop: Bad luck or bad design?*, 2001.

[18] P. Klemperer, *How (not) to run auctions: The European 3G telecom auctions*, European Economic Review, 46(4-5), 829-845.

[19] J. C. Harsanyi, *Games with incomplete information played by "Bayesian" players, I-III Part I. The basic model*, Management science. 1967 14(3):159-82.

[20] M. J. Osborne and A. Rubinstein, *A course in game theory*, MIT press; 1994 Jul 12.

[21] R. B. Myerson, *Mechanism design*, in Allocation, Information and Markets 1989, pp. 191-206, Palgrave Macmillan, London.

[22] R. B. Myerson and M. A. Satterthwaite, *Efficient mechanisms for bilateral trading*, Journal of economic theory, 29(2):265-81, 1983.

[23] T. Lykouris, V. Syrgkanis, E. Tardos, *Learning and efficiency in games with dynamic population*, in Proceedings of the twenty-seventh annual ACM-SIAM symposium on Discrete algorithms, pp. 120-129, Society for Industrial and Applied Mathematics, 2016.

[24] A. Dreves, F. Facchinei, C. Kanzow and S. Sagratella, *On the solution of the KKT conditions of generalized Nash equilibrium problems*, SIAM Journal on Optimization, 21(3):1082-108, 2011.

[25] F. Facchinei and C. Kanzow, *Generalized Nash equilibrium problems*, 4OR. 5(3):173-210, 2007.

[26] B. Bünz, B. Lubin and S. Seuken, *Designing core-selecting payment rules: A computational search approach*, `http://www.ifi.uzh.ch/ce/publications/Core_Payment_Rules.pdf`, 2018.