CS639: Algorithmic Game Theory & Learning

University of Wisconsin–Madison, Spring 2026 *Instructor: Kirthevasan Kandasamy*

Homework 6. Due 05/01/2026, 11.59 pm

Instructions:

1. Homework is due on Canvas by 11:59 pm on the due date. Please plan to submit well before the deadline. Refer to the course website for policies on late submission.
2. Homework must be typeset using appropriate software, such as LaTeX. Handwritten and scanned submissions will **not** be accepted.
3. Your solutions will be evaluated on correctness, clarity, and conciseness.
4. Unless otherwise specified, you may use any result we have already discussed in class. Clearly state which result you are using.
5. If you use any external references, please cite them in your submission.
6. **Collaboration:** You may collaborate in groups of size up to 3. If you collaborate, please indicate your collaborators at the beginning of the homework. Even if you collaborate, *you must write the solution in your own words*.

1 Residency Matching

We wish to match medical school graduates (students) to hospital residency positions. Let $\mathcal{H}$ denote the set of hospitals, each with exactly one open position. The hospitals are partitioned into two groups $\mathcal{H} = \mathcal{H}_{\mathrm{G}} \cup \mathcal{H}_{\mathrm{R}}$. Here, $\mathcal{H}_{\mathrm{G}}$ are rural government-run hospitals, and $\mathcal{H}_{\mathrm{R}}$ are regular (competitive, urban) hospitals which students generally prefer over $\mathcal{H}_{\mathrm{G}}$. Government hospitals have no preferences over students—they are happy to receive anyone.

The students $\mathcal{S}$ are also partitioned into two groups $\mathcal{S} = \mathcal{S}_1 \cup \mathcal{S}_2$. Students in $\mathcal{S}_1$ have completed government medical service; as a reward, the government **guarantees** that every student in $\mathcal{S}_1$ will be matched to some hospital. Students in $\mathcal{S}_2$ did not complete government service and receive no such guarantee. We have $|\mathcal{S}_1| \leq |\mathcal{H}_{\mathrm{G}}|$ which allows the government to guarantee matches for $\mathcal{S}_1$, but $|\mathcal{S}| > |\mathcal{H}|$, so some students in $\mathcal{S}_2$ will necessarily go unmatched.

Each student $i \in \mathcal{S}$ has strict preferences p_i over $\mathcal{H} \cup \{\varnothing\}$, where $\varnothing$ denotes being unmatched. All students prefer any hospital to being unmatched, i.e., $h \succ_i \varnothing$ for all i, h. Each urban hospital $h \in \mathcal{H}_{\mathrm{R}}$ has strict preferences p_h over all students and does not differentiate between $\mathcal{S}_1$ and $\mathcal{S}_2$. Government hospitals have no preferences. A matching $\mu : \mathcal{S} \to \mathcal{H} \cup \{\varnothing\}$ assigns each student to at most one hospital, with each hospital matched to at most one student.

[10 pts] Design a mechanism for this residency matching problem that asks students to report their bids $\{b_i\}_{i \in \mathcal{S}}$ indicating their preference ordering over hospitals, and computes a matching μ. Your mechanism should satisfy the following properties.

1. *Guaranteed match for government-service students.* For a truthful student $i \in \mathcal{S}_1$, regardless of the bids b_{-i} submitted by the others, $\mu(i) \neq \varnothing$.
2. *Stability.* A pair (i, h) with $\mu(i) \neq h$ is a *blocking pair* for a matching μ if *(a)* $h \succ_i \mu(i)$, and *(b)* h is unmatched, or $h \in \mathcal{H}_{\mathrm{R}}$ and h prefers i to its current match. (For government hospitals, which have no preferences, *(b)* reduces to h being unmatched.) A matching is *stable* if it has no blocking pairs. The mechanism should return a stable matching when all students report truthfully.
3. *Dominant strategy incentive-compatibility (DSIC) for students.* Regardless of the bids b_{-i} submitted by other students, the best strategy for any student i is to report truthfully, i.e., $b_i = p_i$.

Your solution may build on an algorithm we studied in class. You do not need to prove these properties from scratch but should sufficiently prove/explain why these properties are satisfied. You may assume that the regular hospitals will report their preferences truthfully to the mechanism designer.

2 Myerson's lemma for ad auctions

In this question, you will consider the following variant of ad auctions we saw in homework 4, where we allow the platform to not assign an ad slot to advertisers. The platform (mechanism designer) has K ad slots with publicly known click through rates $c_1 \geq c_2 \geq \cdots \geq c_K$. The clickthrough rate of a slot is the (average) number of times an ad in that slot gets clicked in that slot (per unit time interval). There are $n (> K)$ advertisers (agents), who each have a private value $v_i \in \mathbb{R}_+$ for a single click of an ad. That is, if advertiser $i \in [n]$ is placed in slot $k \in [K]$, the value she derives is $v_i c_k$. The advertising platform will solicit bids $b_i \in \mathbb{R}_+^n$ representing the value derived from a click from the advertisers. It will then allocate *at most one* advertiser to each slot, and charge payments $p \in \mathbb{R}^n$.

We can write an outcome as $\omega = (\omega_1, \ldots, \omega_k)$ where $\omega_k \in \{0\} \cup [n]$. Here, if $\omega_k = i$ means that advertiser i assigned to slot k, and if $\omega_k = 0$, it means that no advertiser is assigned to slot k. An advertiser cannot be assigned to more than one slot, i.e for any $j \neq k$ $\omega_j \neq \omega_k$ whenever $\omega_j \neq 0$. We can write agent i's value for outcome ω as $v_i \cdot \sum_{k=1}^{K} c_k \mathbb{1}(\omega_k = i)$. Therefore, this is a single-parameter environment.

1. **[1 pts]** *(Single parameter environment.)* Show that this is a single parameter environment.
2. **[8 pts]** *(Welfare-maximizing ad auctions.)* Apply Myerson's lemma to derive a truthful auction which maximizes welfare and show that this is precisely the same as the one obtained via the VCG mechanism in homework 4. Is it surprising that both auctions are the same?

For the remainder of this problem, you may assume that the bidder valuations $\{v_i\}_i$ are drawn from a regular distribution D with cdf F and pdf f. Let $\phi(v) = v - \frac{1-F(v)}{f(v)}$ denote the virtual value function.

3. **[6 pts]** *(Revenue-optimal ad auctions)* Apply Myerson's lemma to derive a truthful mechanism to maximize the expected revenue (i.e., when bidder valuations are drawn from D) when agents are reporting truthfully.
4. **[6 pts]** *(The value of more bidders)* Computing the revenue-optimal auction with n bidders requires knowing the distribution D, which may not always be known to a seller. Show that the VCG mechanism for the online ad auction, i.e the auction in part 2 with $n + K$ bidders, has no smaller expected revenue than the revenue-optimal auction in part 3 with n bidders.

3 Posted price mechanisms

Consider a seller who wishes to sell a single item. The seller will choose some price p. A buyer with value v, drawn from some distribution D, will purchase the item if $v \geq p$, i.e if she values the item more than the price chosen by the seller. The expected revenue is $\mathrm{rev}(p) = p\mathbb{P}_{v\sim D}(v \geq p)$. If the seller chooses a small price, then the buyer will likely buy the item, but the revenue will be small on purchase. If he chooses a higher price instead, then there will be a large revenue if there is a purchase, but the probability of purchase will be small. The optimal price is simple the price that maximizes the revenue, i.e $p^\star = \mathrm{argmax}_p \, \mathrm{rev}(p)$.

This is called a posted price mechanism (PPM), referring to the fact that the seller posts a price and the buyer purchases it if it is individually rational for her. The PPM is much simpler and more practical than a sealed bid auction, and reflects day-to-day purchases in the real world.

1. **[3 pts]** *(Myerson's optimal auction for single bidders)* There is an alternative interpretation of the posted price mechanism which can be derived by viewing it as a single-item single-bidder auction. Assume that the same seller wishes to design a sealed-bid auction to maximize his expected revenue. The Myerson's optimal auction will solicit a bid b from the buyer, and sell it to the buyer at the reserve price $\phi^{-1}(0)$ if $b > \phi^{-1}(0)$.

   Show that $\phi^{-1}(0) = p^\star$. You may assume that D is regular and that rev is differentiable.

2. *(Learning an optimal reserve price.)* Consider a single-item single bidder auction, similar to the one described in part 1, repeated over T rounds. On each round t, the seller chooses a reserve price p_t. A *new* buyer with value $v_t \in [0, 1]$ appears, *truthfully* bids this value, and receives the item at price p_t if $v_t \geq p_t$. The regret R_T of a revenue-maximizing seller compares the actual revenue earned by the seller to the single best reserve price in hindsight. We have,

$$R_T = \max_{p\in[0,1]} \sum_{t=1}^{T} p\mathbb{1}(v_t \geq p) - \sum_{t=1}^{T} p_t \mathbb{1}(v_t \geq p_t). \tag{1}$$

   In this question, you will design a policy to achieve sublinear regret. In Chapter 4, we studied several algorithms for minimizing regret that operated over a discrete set of K actions. Here, however, we have a continuous set of prices[1] in $[0, 1]$. To handle this, we will first consider a discrete set of K prices, design a regret-minimizing algorithm in this discrete set, and the control the error due to discretization. For this, let $\overline{P} = \left\{\frac{1}{2K}, \frac{3}{2K}, \frac{5}{2K}, \dots, \frac{2K-1}{2K}\right\}$ be a discretization of the set of possible prices $[0, 1]$. We can then decompose the regret as follows,

$$R_T = \underbrace{\max_{p\in[0,1]} \sum_{t=1}^{T} p\mathbb{1}(v_t \geq p) - \max_{p\in\overline{P}} \sum_{t=1}^{T} p\mathbb{1}(v_t \geq p)}_{\text{regret due to discretization}} + \underbrace{\max_{p\in\overline{P}} \sum_{t=1}^{T} p\mathbb{1}(v_t \geq p) - \sum_{t=1}^{T} p_t \mathbb{1}(v_t \geq p_t)}_{\triangleq \overline{R}_T} \tag{2}$$

   Here, the first term in the RHS is the regret due to discretization and the second term $\overline{R}_T$ is the regret of the seller's choices of p_t when compared to the best fixed price in the discretization $\overline{P}$.

   (a) **[3 pts]** *(Regret on the discrete set)* Using the Hedge algorithm or otherwise, outline a procedure that achieves $\mathbb{E}[\overline{R}_T] \in \mathcal{O}(\sqrt{T\log(K)})$ regret on the discrete set, where the expectation $\mathbb{E}$ is with respect to

[1] Note that as $v_t \in [0, 1]$, we can assume that the optimal prices is also in $[0, 1]$ without loss of generality.

any randomness in the algorithm.

(b) **[4 pts]** *(Controlling the regret due to discretization.)* Show that the regret due to discretization, i.e first term in the RHS of (2), can be bounded by T/K.

(c) **[2 pts]** *(Final bound on the regret.)* Show that by choosing the size of the discretization to be $K = T$, we can achieve $\mathbb{E}[R_T] \in \mathcal{O}(\sqrt{T \log(T)})$.

3. **[3 pts]** *(Truthfulness over multiple rounds.)* In the previous question we assumed that bidders will bid truthfully. This is because the auction is truthful in a single round, and we have a *new* buyer on each round for whom truth-telling is a dominant strategy after the seller has chosen a reserve price.

   However, say that the *same* buyer repeatedly appears on all rounds, but has a different value for the item on different days based on her personal needs. If the seller uses the algorithm developed in part 2, is it still a dominant strategy for the buyer to reveal her bid truthfully on each round? If 'yes', then outline a proof, and if 'no', then outline a strategy the buyer could use to obtain an advantage?

4. **[8 pts]** *(Learning an optimal price.)* We will now consider a multi-round version of the PPM stated at the beginning of this question. On each round t, the seller chooses a price p_t. A *new* buyer with value $v_t \in [0, 1]$ appears, and will purchase the item if $v_t \geq p_t$. In particular, unlike in part 2, the buyer does not place a bid before purchasing the item. The regret of the seller is defined exactly the same as in (1), except now p_t should be interpreted as the sale price and not the reserve price of an auction.

   Hint: To answer this question, first revisit Chapter 4.2 on adversarial bandits. Then, outline a procedure to obtain sublinear regret. You may follow a similar discretization strategy to the one in part 2, but will need to choose a different value for K to get an optimal bound on the regret.

   You should achieve $\tilde{\mathcal{O}}(T^{2/3})$ regret, where the $\tilde{\mathcal{O}}$ notation hides any log factors. When compared to part 2, the seller has limited feedback, which results in worse regret. In particular, previously the buyer revealed their value (bid). Now, the seller only observes if the chosen price p_t was larger (no purchase) or smaller (there was a purchase) than the buyer's value.

4 Axiomatic bargaining

Recall the bargaining problem with n agents. Let $\mathcal{U} \subset \mathbb{R}_+^n$ be a closed and strictly convex set of possible utilities the agents could achieve via collaboration. As we did in class, we will assume that the disagreement point is $d = \mathbf{0}$ and that we are only interested in individually rational and translation invariant solutions. Recall the following four axioms for axiomatic bargaining.

1. *Pareto-efficiency:* $f(\mathcal{U})$ is Pareto-optimal in $\mathcal{U}$. That is, for all $u \in \mathcal{U}$, if there exists an agent $i \in [n]$ such that $u_i > f_i(\mathcal{U})$, then there exists another agent $j \in [n]\backslash\{i\}$ such that $u_i < f_i(\mathcal{U})$.
2. *Symmetry:* If $\mathcal{U}$ is a symmetric set[2], then $f_i(\mathcal{U}) = f_j(\mathcal{U})$ for all i, j.
3. *Scale invariance:* For any $\alpha \in \mathbb{R}_n^+$, let $\phi_\alpha(u) = (\alpha_1 u_1, \ldots, \alpha_n u_n)$ and let $\phi_\alpha(\mathcal{U}) = \{\phi_\alpha(u); u \in \mathcal{U}\}$. Then, f is scale invariant if $f(\phi_\alpha(\mathcal{U})) = \phi_\alpha(f(\mathcal{U}))$.
4. *Independent of irrelevant alternatives:* Let $\mathcal{U}' \subset \mathcal{U}$. If $f(\mathcal{U}) \in \mathcal{U}'$, then $f(\mathcal{U}) = f(\mathcal{U}')$.

State if the four axioms are *"satisfied"* or *"not satisfied"* by the following two bargaining solutions. If you answer *"satisfied"*, you should provide a proof and if you answer *"not satisfied"*, you should provide a counter example.

1. **[8 pts]** *(Utilitarian solution:)* First, let us define the social welfare (a.k.a utilitarian welfare) $W^{\mathrm{soc}}(u) = \sum_{i=1}^n u_i$. The utilitarian solution stipulates that we choose the point which maximizes the utilitarian welfare, i.e $f^{\mathrm{U}}(\mathcal{U}) = \mathrm{argmax}_{u \in \mathcal{U}} W^{\mathrm{soc}}(u)$.

[2] By symmetric set we mean that $\mathcal{U}$ is unaltered under any permutation of the axes. That is, for all permutations π of $[n]$, we have $(u_1, \ldots, u_n) \in \mathcal{U} \implies (u_{\pi(1)}, \ldots, u_{\pi(n)}) \in \mathcal{U}$.

Hint 1: In Nash bargaining, we used the fact that a *strictly* concave function has a unique maximum over a closed (not necessarily strictly) convex set. It is also true that a (not necessarily strictly) concave function has a unique maximum over a closed *strictly* convex set. Note that a linear function is concave but not strictly concave.

Hint 2: In class, we mentioned, but did not prove that the utilitarian solution does not satisfy scale invariance. To construct a counter-example, let $n = 2, \mathcal{U} = \{u \in \mathbb{R}^2_+; u_1^2 + u_2^2 \le 1\}$, and choose an appropriate α. You may use the method of Lagrange multipliers or any software (e.g., Mathematica) to compute the utilitarian solution.

2. **[8 pts]** *(Kalai-Smordinsky (KS) solution:)* First define $u_i^\circ = \max_{u \in \mathcal{U}} u_i$. The KS solution $f^{\text{KS}}(\mathcal{U})$ chooses the unique point in $\mathcal{U}$ which is Pareto-optimal and is proportional to $(u_1^\circ, \ldots, u_n^\circ)$, i.e $\frac{f_i^{\text{KS}}(\mathcal{U})}{f_j^{\text{KS}}(\mathcal{U})} = \frac{u_i^\circ}{u_j^\circ}$.

   Note that unlike the egalitarian, utilitarian, and Nash bargaining solutions, the KS solution is not obtained by maximizing a welfare function. To motivate the KS solution, assume that agent i has all the bargaining power in the group. Then, she will choose the allocation which maximizes her own utility, i.e u_i°. The KS solution states that final allocation should be proportional to these 'utopian' values the agents could obtain, while also being Pareto-optimal. We have illustrated the KS solution for two agents below.

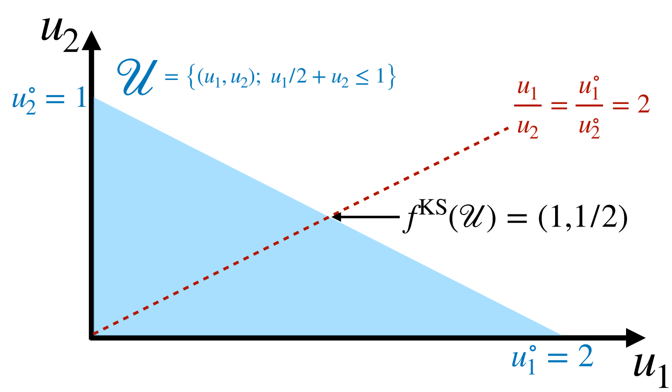

5 Coalitional games

1. **[11 pts]** *(Banzhaf power index)* Consider an n player coalitional game, with characteristic function v. The following quantity, called the Banzhaf power index, is sometimes used to distribute the payoffs.

$$\psi_i(v) = \frac{1}{2^{n-1}} \sum_{S \subset [n] \setminus \{i\}} (v(S \cup \{i\}) - v(S)) \tag{3}$$

   You will note that this is similar to the alternative formula for the Shapley value, in that it sums over all subsets $S \subset [n] \setminus \{i\}$ and considers i's contribution to S via $v(S \cup \{i\}) - v(S)$. However, instead of weighting them with different coefficients, this simply takes their average.

   Recall Shapley's four axioms for coalitional games:

   (a) *Efficiency/feasibility:* We have, $\sum_{i=1}^n \psi_i(v) = v([n])$.

   (b) *Dummy:* If $v(S \cup \{i\}) = v(S)$ for all $S \subset [n] \setminus \{i\}$, then $\psi_i(v) = 0$.

   (c) *Additivity:* For any two characteristic functions u, v, we have $\psi(u + v) = \psi(u) + \psi(v)$.

   (d) *Symmetry:* If $v(S \cup \{i\}) = v(S \cup \{j\})$ for all $S \subset [n] \setminus \{i, j\}$, then $\psi_i(v) = \psi_j(v)$.

   Identify which of the axioms is satisfied by the Banzhaf power index. Justify your answers with proofs and/or counter-examples.

2. For the two coalitional games shown below, *(i)* identify all solutions in the core, *(ii)* compute the Shapley value for all agents, and *(iii)* compute the Banzhaf power index for all agents.

   Hint: You may find it easier to use the Shapley axioms to compute the Shapley value instead of the formula.

(a) **[5 pts]** *(3-player majority game)* There are $n = 3$ players, and $v(S) = \mathbb{1}(|S| \geq 2)$. That is, the value is 1 as long as there are a majority of the players.

(b) **[7 pts]** *(T–veto game)* Here, a subset $T \subseteq [n]$ of the players hold the power. The characteristic function[3] is given by, $v(S) = \mathbb{1}(T \subseteq S)$.

[3] Note that this is the same as the w_T function we saw in the uniqueness proof of the Shapley value.