CS639: Algorithmic Game Theory & Learning

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

Homework 5. Due 04/17/2026, 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 One-sided matching

There are n agents who each bring a good for trade. Each agent i has her preferences p_i, i.e. a *strict* total ordering over the n goods brought by all agents. If i prefers j's good to k's, we write $j \succ_i k$. An allocation is a bijection $\pi : [n] \to [n]$. If $\pi(i) = j$, it means j's good was allocated to i.

A mechanism for trading ask agents to report their bids $\{b_i\}_{i=1}^n$ (preference ordering) and computes an allocation. Recall the top-trading cycles (TTC) algorithm we discussed in class.

Algorithm 1 Top trading cycles (TTC)

```
Given: bids (reported preferences) {b_i}_{i=1}^n
S_t ← [n].                                                        # unallocated agents
t ← 1.                                                            # round index
π(i) ← NULL    for all i ∈ [n].                                   # allocation
while |S_t| > 0 do                                  # while there are unallocated agents
    Let G_t = (S_t, E_t) be the directed graph where E_t contains directed edges
        from each i ∈ S_t to the agent with i's most preferred good in S_t according
        to the bids.
    Let C_t be all edges which form cycles in G_t.
    For each (i, j) ∈ C_t, set π(i) ← j.
    Let S_{t+1} ← S_t \ {i; (i, j) ∈ C_t}                         # Clear all cycles
end while
Return π.
```

You are encouraged to read Section 10.4 of KP before attempting this question.

1. **[3 pts]** *(Termination)* Show that TTC terminates in a finite number of steps.
2. **[3 pts]** *(Individual rationality)* Show that TTC returns an individually rational allocation. That is, for all agents i, regardless of the bids b_{-i} submitted by the other agents, we have $\pi(i) \succeq_i i$ if agent i truthfully reports $b_i = p_i$.
3. **[5 pts]** *(Pareto-efficient)* An allocation π is Pareto-efficient if no subset of agents can exchange the goods allocated to them so that all of them obtain a more preferred good. Formally, we say that an allocation π' ($\neq \pi$) Pareto-dominates another allocation π if for all $i \in \{j; \pi(j) \neq \pi'(j)\}$, we have $\pi'(i) \succ_i \pi(i)$. An allocation π is Pareto-efficient if there is no allocation that Pareto-dominates π.

   Show that, when all agents report truthfully, TTC returns a Pareto-optimal allocation.
4. **[5 pts]** *(DSIC)* Show that the TTC algorithm is DSIC. That is, regardless of the bids b_{-i} submitted by the other agents, the best strategy for an agent is to report $b_i = p_i$.
5. **[5 pts]** *(Uniqueness)* Indicate if the following statement is true or false: *When all agents report truthfully, the TTC allocation is the* unique *allocation that is both individually rational and Pareto-efficient.* If true, provide a proof. If false, provide a counterexample.
6. **[0 pts]** *(Reading exercise on kidney exchange)* Read Section 1 of TRA Lecture 10, available at https://timroughgarden.org/f13/l/l10.pdf.

2 Properties of fair resource allocation and Max-min fairness

In this question, you will prove some properties about fair sharing and the max-min fairness (MMF) algorithm that we discussed, but did not prove, in class.

1. *(Any Pareto-efficient allocation has zero resource loss.)* First let us recall the definition of Pareto-efficiency of a resource allocation $x \in \mathbb{R}_+^n$, where $\sum_{i=1}^n x_i \leq 1$. An allocation x is said to be Pareto-dominated by another allocation x' if $\widetilde{u}_i(x_i') \geq \widetilde{u}_i(x_i)$ for all $i \in [n]$, and $\widetilde{u}_j(x_j') > \widetilde{u}_j(x_j)$ for some $j \in [n]$. An allocation x is said to be Pareto-efficient if no other allocation Pareto-dominates x.

Additionally, recall the definition of the resource loss $\ell(x) = \min(\ell^{\rm ud}(x), \ell^{\rm or}(x)+\ell^{\rm ur}(x)) > 0$, where $\ell^{\rm ur}(x) = 1-\sum_{i=1}^n x_i$ are the unallocated resources, $\ell^{\rm or}(x) = \sum_{i=1}^n (x_i - d_i)^+$ are the resources allocated over an agent's demand, and $\ell^{\rm ud}(x) = \sum_{i=1}^n (d_i - x_i)^+$ is the unmet demand of the agents.

Assume that agent utilities are *strictly* increasing up to the demand for all i. That is, for all agents i we have $\widetilde{u}_i(x_i) = \widetilde{u}_i(d_i)$ for all $x_i > d_i$ and $\widetilde{u}_i(x_i) < \widetilde{u}_i(x'_i)$ for all $x_i < x'_i$ such that $x_i < x'_i \leq d_i$. In class we saw that if $\ell(x) = 0$ for an allocation x, then x is Pareto-efficient.

[5 pts] Show that if the allocation x is Pareto-efficient, then $\ell(x) = 0$.

2. *(Overstating your resource demand does not help in MMF.)* Recall the DSIC proof for the Max-min fairness mechanism $M^{\rm MMF}$. Fix the bids (reported demands) b_{-i} of all agents except i and the true demand d_i of agent i. In class we showed that under-stating the demand can hurt the agent, i.e if agent i reported $b_i < d_i$, then $u_i(b_i, b_{-i}; M^{\rm MMF}) \leq u_i(d_i, b_{-i}; M^{\rm MMF})$. Here, $u_i(b; M) = \widetilde{u}_i(x_i(b; M))$, where $x_i(b; M)$ is the allocation returned to agent i by a mechanism M when the bids are b.

   [5 pts] Show that over-stating the demand does not help the agent in MMF, i.e if she reported $b_i > d_i$, then $u_i(b_i, b_{-i}; M^{\rm MMF}) = u_i(d_i, b_{-i}; M^{\rm MMF})$.

3 Mechanisms for resource allocation

Recall the setting for fair resource allocation, where we have n agents sharing a resource of size 1. Each agent has a demand. User i's utility increases up to her demand d_i, but does not increase thereafter. Let us assume that all agents have an equal endowment, i.e $e_i = 1/n$ for all i. In this question, we will look at the following four policies and determine if they satisfy our required desiderata for a fair resource allocation mechanism.

1. **Equal Split ($M^{\rm ES}$):** $M^{\rm ES}$ simply allocates $1/n$ resources to each agent, regardless of the reported demands.
2. **Proportional to reported demand ($M^{\rm PD}$):** $M^{\rm PD}$ allocates proportional to the bids (reported demands). That is, $x_i = \frac{b_i}{\sum_{j=1}^n b_j}$.
3. **Max-Min Fairness ($M^{\rm MMF}$):** $M^{\rm MMF}$ follows the max-min fairness policy we discussed in class.
4. **Minimum Demand First ($M^{\rm MDF}$):** $M^{\rm MDF}$ chooses the agent with the smallest bid b_i (breaking ties arbitrarily), and allocates b_i to that agent. It proceeds in this fashion until no resources are left. We can describe it formally via the following algorithm.

Algorithm 2 $M^{\rm MDF}$

```
Given: bids {b_i}_{i=1}^n
x_i ← 0 for all i ∈ [n].                    # allocations
r ← 1.                                      # remaining resources
for i in ascending order of b_i do
    If b_i < r:
        x_i ← b_i.    r ← r − b_i.
    Else
        x_i ← r.    Break
end for
```

[26 pts] Fill out the following table indicating if the policies satisfy (✓) or do not satisfy (✗) the three desiderata: efficiency, individual rationality (IR), and dominant-strategy incentive compatiblity (IC). If you indicate ✗, you need to provide a counter-example, choosing specific values for the number of agents n, the true demand $\{d_i\}_i$, and bids $\{b_i\}_n$ and show that the property is not satisfied. If you indicate ✓, you should provide a proof/explanation. To help you get started, we have partially filled out this table and provided the explanations.

Policy	M^{ES}	M^{PD}	M^{MMF}	M^{MDF}
Efficiency	(1) ✗	(4)	(7)	(10)
IR	(2)	(5)	(8) ✓	(11)
DSIC	(3)	(6)	(9)	(12)

See below for the explanations and counter-examples.

(1) M^{ES} is not efficient. Suppose there are two agents and $b_1 = d_1 = 0.2$ and $b_2 = d_2 = 0.8$. Then, agent 1 is allocated 0.3 more than she needs which could have been used by agent 2.

(8) M^{MMF} is IR, as we saw in class (N.B. If we proved a result in class, you can simply state so).

4 Deriving the VCG mechanism

In this question, you will derive the VCG mechanism for two problems. You will work with the following set up.

Set up. Let Ω be a set of outcomes. There are a set of n agents, who each have a valuation function $v_i : \Omega \to \mathbb{R}_+$. Here $v_i(\omega)$ denotes the value agent i has for an outcome $\omega \in \Omega$. Let $\mathcal{V} = \{f : \Omega \to \mathbb{R}_+\}$ denote all possible agent valuation functions. We will assume that the mechanism designer's valuation is zero for all outcomes. A mechanism M is a process to solicit bids $b = \{b_i\}_{i=1}^n$, where $b_i \in \mathcal{V}$, and choose an outcome $\omega(b; M)$ and a payment vector $p(b; M)$ based on the submitted bids.

The VCG mechanism. Let $W(\omega) = \sum_{i=1}^n b_i(\omega)$ be the reported social welfare and let $W_{-i}(\omega) = \sum_{j\neq i} b_j(\omega)$ be the reported social welfare without agent i. The VCG mechanism M^{VCG} stipulates that we choose the outcome which maximizes W, i.e $\omega(b; M^{\mathrm{VCG}}) = \mathrm{argmax}_{\omega\in\Omega} W(\omega)$. The payment for each agent is the externality her presence causes to the others, i.e., $p_i(b; M^{\mathrm{VCG}}) = \max_{\omega\in\Omega} W_{-i}(\omega) - W_{-i}(\omega(b; M^{\mathrm{VCG}}))$.

1. **[5 pts]** *(Single-item auctions)* In a single item auction, an auctioneer has a single item that he wishes to sell. There are n buyers who each have a private value $\widetilde{v}_i \in \mathbb{R}_+$ for the item. The auctioneer will solicit bids $\widetilde{b}_i \in \mathbb{R}_+$ from each buyer, choose a winner, and charge each buyer some payment.

   Specify the space of outcomes Ω, the valuation functions v_i and bid functions b_i for this problem and derive the VCG mechanism. Show that this is the same as the second-price (Vickrey) auction.

2. **[7 pts]** *(Online advertising)* An advertising 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 $\widetilde{v}_i \in \mathbb{R}_+$ 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 $\widetilde{v}_i c_k$. The advertising platform will solicit bids $\widetilde{b}_i \in \mathbb{R}_+^n$ representing the value derived from a click from the advertisers. It will then allocate exactly one advertiser to each slot, which we can write as $\omega = (\omega_1, \ldots, \omega_k)$ where $\omega_k \in [n]$ is the advertiser assigned to slot k and $\omega_j \neq \omega_k$ for $j \neq k$. Finally, it will charge payments $p \in \mathbb{R}^n$.

   Specify the space of outcomes Ω, the valuation functions v_i and bid functions b_i for this problem and derive the VCG mechanism.

5 Friends of VCG

1. **[6 pts]** *(Groves mechanism, Affine maximizer auctions)* Consider the setting under the paragraph **Set up** in problem 4. Let us consider the following mechanism, which is closely related to the VCG mechanism. First define an reported affine welfare function $W^{\alpha,\beta}(\omega) = \sum_{j=1}^n \alpha_j b_j(\omega) + \beta(\omega)$, where $\alpha_i \in \mathbb{R}_+$ and $\beta : \Omega \to \mathbb{R}$. Similarly define $W_{-i}(\omega) = \sum_{j\neq i} \alpha_j b_j(\omega) + \beta(\omega)$ be the reported affine social welfare without agent i. Finally, for each $i \in [n]$, let $h_{-i} : \mathcal{V}^{n-1} \to \mathbb{R}_+$ be an arbitrary function of all bids except for i.

   Consider a mechanism $M^{\alpha,\beta,h}$ which chooses an outcome to maximize the reported affine welfare, i.e $\omega(b; M^{\alpha,\beta,h}) =$

$\operatorname{argmax}_{\omega\in\Omega} W^{\alpha,\beta}(\omega)$. The payment for each agent is,

$$p_i(b; M^{\alpha,\beta,h}) = \frac{1}{\alpha_i}\Big(h_{-i}(b_{-i}) - W_{-i}(\omega(b; M^{\alpha,\beta,h}))\Big).$$

Show that this mechanism is DSIC.

Hint: You may adapt the DSIC proof for the VCG mechanism we studied in class.

N.B: Here, the α_i's can be thought of as weights to prioritize some agents over the others, and $\beta(\omega)$ is usually referred to as a booster function to prioritize some outcomes over the others. This defines a class of DSIC mechanisms, defined below,

$$\Big\{M^{\alpha,\beta,h};\ \alpha\in\mathbb{R}_+,\ \beta\in\{f:\Omega\to\mathbb{R}\},\ h\in\{f:\mathcal{V}^{n-1}\to\mathbb{R}\}^n\Big\}.$$

Note that when we use the following two choices, we obtain the VCG mechanism.

(i) Set $\alpha_i = 1$ and set $\beta(\omega)$ to be the value of the mechanism designer.

(ii) Set $h_{-i}(b_i) = \max_{\omega\in\Omega} W^{\alpha,\beta}_{-i}(\omega)$.

If we only use choice (i), the class $\{M^h; h\}$ is called the Groves mechanisms, and if we only use choice (ii), the class $\{M^{\alpha,\beta}; \alpha,\beta\}$ is called affine maximizer auctions.

2. *(Buyer as mechanism designer)* In the VCG mechanism we studied in class, the mechanism designer acted as a seller or platform: agents submitted bids representing how much they *valued* receiving a good or service, and the mechanism chose an outcome and charged payments to the agents. In many important settings the roles are reversed: a mechanism designer (e.g., a government agency) wishes to *procure* goods or services from a set of agents (suppliers), each of whom has a private cost for providing their service. For instance, a government may solicit bids from construction firms to build infrastructure, or a chip manufacturer may procure components from suppliers to assemble a chip. In each case, the buyer has a publicly known value for each possible combination of services and wishes to decide which agents to hire and how much to pay them so as to maximize welfare.

Set up. We have a set of n agents (suppliers), a mechanism designer (buyer), and a set of possible outcomes $\Omega = \Omega_1\times\cdots\times\Omega_n$. Here, every outcome $\omega\in\Omega$ can be written as $\omega = \{\omega_i\}_{i\in[n]}$, where $\omega_i\in\Omega_i$ is the outcome relevant to agent i. For example, when a government wishes to contract firms to build roads, ω_i could be the subset of roads that firm i is assigned to build, and $c_i(\omega_i)$ is firm i's cost for building that subset. Agent i has a *cost function* $c_i:\Omega_i\to\mathbb{R}_+$, where $c_i(\omega_i)$ denotes agent i's cost for providing her service under outcome ω_i; these costs are private information. Let $\mathcal{C}_i\subseteq\{f:\Omega_i\to\mathbb{R}_+\}$ be the space of all agent cost functions for agent i. The mechanism designer has a valuation function $v_0:\Omega\to\mathbb{R}_+$, where $v_0(\omega)$ is her value for outcome ω; this is publicly known. The social welfare of an outcome ω is $W(\omega) = v_0(\omega) - \sum_{i=1}^n c_i(\omega_i)$.

In a mechanism $M = (\omega, p)$, each agent i submits a bid $b_i\in\mathcal{C}_i$ representing her cost function (not necessarily truthfully). Based on the submitted bids $b = \{b_i\}_{i\in[n]}$, the mechanism chooses an outcome $\omega(b; M) = \{\omega_i(b; M)\}_{i\in[n]}$ and a payment vector $p(b; M)\in\mathbb{R}^n$, where $p_i(b; M)$ is the amount *paid to* agent i by the mechanism designer. Agent i's utility u_i and the mechanism designer's utility u_0 are:

$$u_i(b; M) = p_i(b; M) - c_i\big(\omega_i(b; M)\big), \qquad u_0(b; M) = v_0\big(\omega(b; M)\big) - \sum_{i=1}^n p_i(b; M).$$

Note the sign conventions compared to the setting in class: agents *receive* payments (rather than making them), utility for agents is payment minus cost (rather than value minus payment), and welfare is the buyer's value minus costs (rather than a sum of values). Note that the welfare is equal to the sum of everyone's utilities.

Null outcome. We assume that each Ω_i contains a "null" outcome $\omega_\varnothing$ corresponding to agent i not participating (for example, in the road-building setting, $\omega_\varnothing$ means firm i is not assigned any roads), with $c_i(\omega_\varnothing) = 0$ for all i.

The VCG mechanism for procurement. Each agent $i \in [n]$ submits a bid $b_i \in \mathcal{C}_i$ representing her cost function. The mechanism designer chooses the outcome $\omega^\star$ that maximizes the reported welfare and *pays* each agent i, an amount p_i as shown below. We have:

$$\omega(b; M) = \operatorname*{argmax}_{\omega \in \Omega} \left(v_0(\omega) - \sum_{i=1}^{n} b_i(\omega_i) \right) \quad \triangleq \omega^\star.$$

$$p_i(b; M) = \left(v_0(\omega^\star) - \sum_{j \neq i} b_j(\omega_j^\star) \right) - \max_{\omega \in \Omega;\ \omega_i = \omega_\varnothing} \left(v_0(\omega) - \sum_{j \neq i} b_j(\omega_j) \right).$$

The first term in p_i is the reported welfare at $\omega^\star$ without i. The second term is the maximum welfare achievable by the buyer and the other agents when i does not participate ($\omega_i = \omega_\varnothing$). Thus, agent i is paid the difference: how much additional welfare she creates for the rest of the system by participating.

[10 pts] Properties. Clearly, the above mechanism is efficient, in that it maximizes the welfare when all agents report truthfully. Prove that it also satisfies the following two properties.

(a) **Individual rationality (IR).** A truthful agent has non-negative utility regardless of the others' bids. That is, for all $b_{-i} \in \mathcal{C}_{-i}$, we have $u_i(c_i, b_{-i}; M) \geq 0$. Equivalently, the agent is always paid at least her true cost, $p_i(c_i, b_{-i}; M) \geq c_i\big(\omega_i(c_i, b_{-i}; M)\big)$.

Hint: You will need to use the existence of $\omega_\varnothing$, and the fact that $c_i(\omega_\varnothing) = 0$.

(b) **Truthfulness (DSIC).** Reporting truthfully is a dominant strategy for each agent. That is, for all $b_i \in \mathcal{C}_i$ and $b_{-i} \in \mathcal{C}_{-i}$, we have $u_i(c_i, b_{-i}; M) \geq u_i(b_i, b_{-i}; M)$.