Problem Set 3

Question 1. Prove Gibbard-Satterthwaite Theorem by following the procedures below.

1. A social choice function f is *monotone* if $f(\prec_1, \ldots, \prec_i, \ldots, \prec_n) = a \neq a' = f(\prec_1, \ldots, \prec_i', \ldots, \prec_n)$ implies that $a' \prec_i a$ and $a \prec_i' a'$. That is, if the social choice changed from a to a' when a single voter i changed his vote from $\prec_i$ to $\prec_i'$ then it must be because he switched his preference between a and a'. Prove that a social choice function is strategy-proof if and only if it is monotone.

2. Let $S \subseteq A$ and $\prec \in L$. Denote by $\prec^S$ the order obtained by moving all alternatives in S to the top in $\prec$. If f is monotone and onto, prove that $f(\prec_1^S, \ldots, \prec_n^S) \in S$ for any $\prec_1, \ldots, \prec_n \in L$ and any $S \subseteq A$.

3. Consider the following extension of a social choice function f to a social welfare function F: $F(\prec_1, \ldots, \prec_n) = \prec$, where $a \prec b$ if and only if $f(\prec_1^{\{a,b\}}, \ldots, \prec_n^{\{a,b\}}) = b$. Prove that F is a valid social welfare function by showing that 1) for any $a, b \in A$, we have either $a \prec b$ or $b \prec a$ (but not both); 2) $a \prec b$ and $b \prec c$ implies $a \prec c$ for any $a, b, c \in A$.

4. If f is monotone, onto, and nondictator, prove that F satisfies unanimity, independence of irrelevant alternatives, and nondictatorship.

Question 2. In a reverse auction, a buyer wants to buy an item from one of many sellers. In this auction, the sellers are the bidders. Each bidder i submits a bid $b_i > 0$ which is the price of the item. The buyer buys the item from the bidder with the lowest bid. Bidder i wins if the buyer buys the item from him. The valuation of the bidders are as follows: $v_i(i \text{ wins}) = -w_i$ for certain $w_i > 0$ and $v_i(j \text{ wins}) = 0$ for all $j \neq i$.

1. Formulate the reverse auction as a social choice problem, and give a VCG mechanism with Clarke pivot rule such that it is a dominant strategy for each bidder i to bid $b_i = w_i$.

2. Consider a generalized version of reverse auction where the sellers are the edges in a directed graph $G = (V, E)$ with a source s and a sink t. Suppose there is at least one path from s to t in G. Instead of buying from only one seller as in the original setting, the buyer buys multiple edges forming a path from s to t. Following the setting in the reverse auction, each seller i submits a bid $b_i > 0$ which is the price of the corresponding edge. The buyer buys the path with the minimum sum of prices.

   For the two graphs below, compute each edge's payment in VCG mechanism with the Clarke pivot rule. The numbers on the edges are the bids.

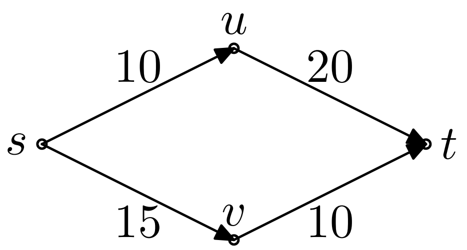

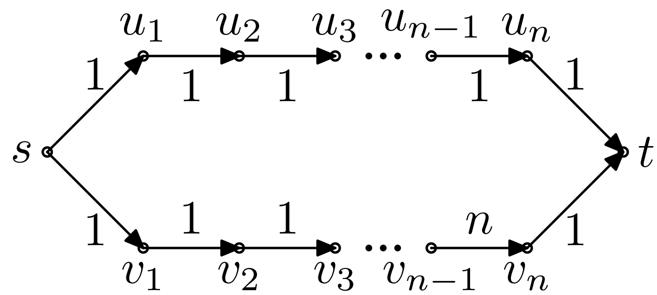

Question 3. This question is based on the paper "Mechanism Design for Fair Division: Allocating Divisible Items without Payments'' by Richard Cole, Vasilis Gkatzelis, Gagan Goel, published in EC'13.

Consider a set of m items to be allocated to a set of n agents. Each item is *divisible*, meaning that it can be divided into arbitrarily small pieces, which are then allocated to different agents. An *allocation* $x = (x_{ij})_{i\in[n],j\in[m]}$ of these items to the agents defines the fraction x_{ij} of each item j that each agent i will be receiving. We allow *partial allocations*, i.e., it is allowed that some portion of an item is unallocated ($\sum_{i=1}^{n} x_{ij} < 1$ for some j). Each agent i has a *valuation function* v_i, and his valuation for allocation x is given by $v_i(x) = \sum_{j=1}^{m} v_{ij}x_{ij}$, where v_{ij} denotes agent i' value for item j.

The *Nash Social Welfare* of an allocation x is defined by the *product* of the agents' utilities:

$$\mathcal{NSW}(x) = \sum_{i=1}^{n} v_i(x),$$

while recall that the *social welfare* is defined with the *sum*. It is well-known that the allocation x that maximizes the Nash Social Welfare can be computed in a polynomial time under this setting by standard convex optimization tools. It is also well-known that the maximizer to the Nash Social Welfare is unique in some sense: for x^* and $x^\dagger$ that both maximize the Nash Social Welfare, we have $v_i(x^*) = v_i(x^\dagger)$ for each agent i.

The mechanism designed in the paper is described as follows.

- Each agent i reports his valuation $(v_{i1}, \ldots, v_{im})$.
- Compute x^* that maximizes the Nash Social Welfare.
- For each agent i, remove this agent and compute the Nash Social Welfare maximizing allocation x^*_{-i} that would arise in his absence.
- Let x be the (partial) allocation such that $x_{ij} = f_i \cdot x^*_{ij}$, where

$$f_i = \frac{\prod_{i'\neq i} v_{i'}(x^*)}{\prod_{i'\neq i} v_{i'}(x^*_{-i})}.$$

- Output allocation x.

We have seen in the class that a strategy-proof mechanism is difficult to design if payment/money is not involved. The mechanism designed in this paper is a multiplicative variant of the VCG mechanism that "simulate" the payment by "resource burning". At the final stage of the mechanism, some fraction of the allocation to each agent is thrown away, and the thrown away portion of the items can be viewed as the payment.

1. Prove that the allocation x returned by the mechanism is feasible by showing that $f_i \in [0, 1]$ for each $i \in [n]$.

2. Prove that the mechanism is strategy-proof.

3. (Extra) Prove that $v_i(x) \geq \frac{1}{e} \cdot v_i(x^*)$ for every agent i. This shows that the mechanism has some guarantees for allocation efficiency. This part may be challenging, and you may want to read the paper.