CS 598RM: Algorithmic Game Theory, Fall 2021
HW 3 Solutions

---

1. Consider the following allocation rules for a single-parameter environment auction. For each of these, prove if they are implementable or not.

   (a) (3 points) The allocation rule that maximizes the social welfare among all feasible allocations. That is, for bids $b_1, b_2 \ldots, b_n$, $x(b) = \arg\max_x \sum_i b_i x_i$.

   (b) (3 points) Suppose the set of feasible allocations is $X = \{x \in [0,1]^n \mid \sum_i x_i = 1\}$. Give each agent a fraction of the item proportional to the distance of their bid from the average of all bids (normalized to 1), and additionally have $x_i = 0$ when $b_i = 0$. That is, if $b_1, b_2, \ldots, b_n$ are the bids of all agents, and if we let $b_{avg} = \sum_i b_i/n$, then the allocation $x(b)$ satisfies for all agents i, j where $x_i, x_j > 0$, $\frac{x_i}{x_j} = \frac{|b_i - b_{avg}|}{|b_j - b_{avg}|}$, and $\sum_i x_i = 1$.

Solution.

By Myerson's Lemma, we know that an allocation rule is implementable if and only if it is monotone.

(a) The allocation is monotone, and thus implementable. To see this, fix an agent i and bids from all other agents $\boldsymbol{b}_{-i}$. We need to show that i's allocation, x_i is non-decreasing in i's bid, b_i. For a given bid profile, we have $x(\boldsymbol{b}) = \sum_i b_i x_i$. Since this is a linear function in $\boldsymbol{b}$, we can view the problem of computing the social welfare-maximizing allocation as a linear function maximization problem, where the bids b_j correspond to the weights. We need to show that for all i, for all $\boldsymbol{b}_{-i}$, if $b'_i > b_i$ then $x_i(b'_i, \boldsymbol{b}_{-i}) \geq x_i(b_i, \boldsymbol{b}_{-i})$.
Letting x' and x be the SW maximizing allocations at $(b'_i, \boldsymbol{b}_{-i})$ and $(b_i, \boldsymbol{b}_{-i})$, respectively, we have

$$\begin{cases} \sum_j b'_j x'_j \geq \sum_j b_j x'_j \\ \sum_j b_j x_j \geq \sum_j b'_j x_j \end{cases} \iff (b'_i - b_i)x'_i \geq 0 \geq (b'_i - b_i)x_i,$$

where the last inequalities follow since $\boldsymbol{b}_{-i}$ does not change. Since $b'_i - b_i > 0$, we get that $x_i(b'_i, \boldsymbol{b}_{-i}) \geq x_i(b_i, \boldsymbol{b}_{-i})$, which implies that the allocaiton rule is monotone.

(b) The allocation is not monotone, and thus not implementable. To see this, we need to show that there exists an instance in which i's bid b_i increases, but i's allocation, x_i decreases. Consider an instance with just three agents, 1, 2 and 3. Let $b_{avg} = \frac{b_1+b_2+b_3}{3}$, and assume that $0 < b_1 = b_2 < b_3$. Fix b_1 and b_3. We know that the allocation $x(\boldsymbol{b})$ satisfies $x_2 = \frac{|b_2 - b_{avg}|}{|b_1 - b_{avg}|} x_1 = \frac{|b_2 - b_{avg}|}{|b_3 - b_{avg}|} x_3$, and $x_1 + x_2 + x_3 = 1$, which implies $x_1 = x_2$ and $x_3 = 1 - 2x_2$. By $x_2 = \frac{|b_2 - b_{avg}|}{|b_3 - b_{avg}|} x_3 = \frac{|b_2 - b_{avg}|}{|b_3 - b_{avg}|}(1 - 2x_2)$, we also get that

$$x_2 = \frac{\frac{|b_2 - b_{avg}|}{|b_3 - b_{avg}|}}{1 + 2\frac{|b_2 - b_{avg}|}{|b_3 - b_{avg}|}} = \frac{|b_2 - b_{avg}|}{|b_3 - b_{avg}| + 2|b_2 - b_{avg}|} = \frac{|b_2 - b_{avg}|}{|b_1 - b_{avg}| + |b_2 - b_{avg}| + |b_3 - b_{avg}|}. \tag{1}$$

Suppose 2 deviates and bids $b'_2 = b_{avg} > b_2$. Then, we have $b'_{avg} = \frac{b_1+b'_2+b_3}{3} > b_{avg}$, but notice that $b'_2 - b_2 > b'_{avg} - b_{avg} \Leftrightarrow b'_{avg} - b'_2 < b_{avg} - b_2$. We have that the allocation $x'(\boldsymbol{b'})$ satisfies $x'_2 = \frac{|b'_2 - b'_{avg}|}{|b_1 - b'_{avg}|} x'_1 = \frac{|b'_2 - b'_{avg}|}{|b_3 - b'_{avg}|} x'_3$, and $x'_1 + x'_2 + x'_3 = 1$, which implies

$$\frac{|b_1 - b'_{avg}|}{|b'_2 - b'_{avg}|} x'_2 + x'_2 + \frac{|b_3 - b'_{avg}|}{|b'_2 - b'_{avg}|} x'_2 = 1 \Leftrightarrow$$

$$x'_2 = \frac{1}{1 + \frac{|b_1 - b'_{avg}|}{|b'_2 - b'_{avg}|} + \frac{|b_3 - b'_{avg}|}{|b'_2 - b'_{avg}|}} = \frac{|b'_2 - b'_{avg}|}{|b_1 - b'_{avg}| + |b'_2 - b'_{avg}| + |b_3 - b'_{avg}|} \tag{2}$$

Notice that since b_1 and b_3 don't change, we have $|b_1 - b_{avg}| + |b_3 - b_{avg}| = |b_1 - b'_{avg}| + |b_3 - b'_{avg}|$. Also recall that $|b'_2 - b'_{avg}| < |b_2 - b_{avg}|$. Together with (1) and (2), these imply that $x'_1 < x_1$, and thus the allocation rule is not monotone.

2. Recall the sponsored search (keyword) auction: k ad slots are on sale on a search website, where probability of getting i^{th} slot clicked is α_i. We know that $\alpha_1 > \alpha_2 > \cdots > \alpha_k > 0$. Given bids $b_1 \geq b_2 \geq \ldots b_n$ of n agents, agent i gets slot i for $i \leq k$, and the remaining agents do not get anything. Note that this allocation rule is monotone. Compute the Myerson's payment for this sponsored search auction.

Solution.

We are given $\alpha_1 > \alpha_2 \cdots > \alpha_k$, and bids $b_1 \geq b_2 \geq \ldots b_n$. The allocation rule which assigns slot i to agent i for $i \leq k$ and nothing to the remaining agents is monotone. Hence we can apply Myerson's payment rule:

$$p(b_i, \mathbf{b}_{-i}) = \sum_{j=1}^{\ell} z_j \cdot [\text{jump in } x_i(\cdot, \mathbf{b}_{-i}) \text{ at } z_j], \tag{3}$$

where $z_1, \ldots, z_\ell$ are the breakpoints of $x_i(\cdot, \mathbf{b}_{-i})$ in $[0, b_i]$.

Consider agent i. While her bid is less than b_{k+1} agent i does not get a slot. When her bid is just above b_{k+1}, her utility changes to α_k. Thus, b_k is a breakpoint and the jump at b_k is $\alpha_k - \alpha_{k+1}$, where $\alpha_{k+1} = 0$.

The next breakpoint occurs at b_{k-1} and the jump here is $\alpha_{k-1} - \alpha_k$. Continuing similarly, one can note that b_j is a breakpoint for each $k \geq j > i$ and the jump at j is $\alpha_j - \alpha_{j+1}$. Since we are considering agent i, there are no more breakpoints. Therefore, applying (3) gives:

$$p(b_i, \mathbf{b}_{-i}) = \sum_{j=i}^{k} b_{j+1}(\alpha_j - \alpha_{j+1}),$$

for agents $i \leq k$, and $p(b_i, \mathbf{b}_{-i}) = 0$ for agents $i > k$.

3. (VCG: Combinatorial Auction)

   (a) (5 points) We want to auction m heterogeneous items among n unit demand bidders, where bidder i has value v_{ij} for item j, and her value for set $S \subset \{1, \ldots, m\}$ is $\max_{j \in S} v_{ij}$. Design a polynomial-time algorithm to implement the VCG auction.

   (b) Consider combinatorial auctions for m items among n bidders. The goal of this exercise is to study special cases in which the VCG mechanism can be implemented efficiently. Suppose that each buyer sends as a bid a valuation vector of $2^m - 1$ numbers (a value for each subset of items, since $v_i(\phi) = 0$ for all i).

      i. (2 points) Suppose that $m = 2$. Show that the VCG mechanism can be implemented in time $O(n^2)$. Provide both the allocation and the payments computed by the VCG mechanism.

      ii. (3 points) Suppose that $m = k$, for some fixed $k > 0$. Show that the VCG mechanism can be implemented in time $O(n^2 \cdot f(k))$, for some function f, where $f(k)$ does not depend on n. Notice that for fixed k, this is polynomial in n. Provide both the allocation and the payments computed by the VCG mechanism.

Solution.

(a) For unit-demand bidders, the VCG auction can be seen as:

   i. Collect bids $b_1, \ldots, b_n$.

   ii. Compute a maximum-weight matching in the complete bipartite graph $(N \cup M, E)$, where the weight of an edge $\{i, j\}$ is v_{ij}, and allocate each item j to its corresponding bidder in the matching.

   iii. Charge each bidder i their externality, i.e. the difference between the weight of the maximum-weight matching computed in Step (ii), and a maximum-weight matching in the complete bipartite graph $(N_{-i} \cup M, E)$, which is the previous graph with the bidder i removed.

   As one can see from the above, we can easily implement the VCG algorithm for uni-demand bidders by computing $(n+1)$ maximum-weight matchings in bipartite graphs, which can be done in polynomial time using a max flow algorithm.

(b) We show how to implement the VCG mechanism in time $O\left(n^2 \cdot (2^m - 1)^2\right)$, using dynamic programming. Note that if we can compute the social welfare maximizing allocation in time $O(n \cdot f(m))$, then we are done, as this will give the allocation, and then for each bidder i, the computation of their payment p_i requires computing the social-welfare again but after excluding i. Since we need to do this for all bidders, the payment computation will add an extra factor of n, and thus, the overall running time will be $O(n^2 \cdot f(m))$.

   First, we compute the optimal allocation. Let M denote the set of m items. Consider a 2-dimensional array $A(i, S)$, of size $(n+1) \times 2^m - 1$. Intuitively, $A(i, S)$ corresponds to the value of assigning the items in S to the first i bidders, for $0 \leq i \leq n$. The return value of our algorithm is $A(n, M)$. We know that $A(i, \emptyset) = 0$ for all i, and that $A(0, S) = 0$, for all sets S. Consider the following recursive formula for $A(i, S)$:

$$A(i, S) = \max_{U \subseteq S} \{v_i(U) + A(i-1, S \setminus U)\}$$

In words, the value of $A(i, S)$ is equal to the value received by the best possible assignment of some items $U \subseteq S$ to i, and the remaining items to the first $i-1$ bidders. Notice that $A(i, S)$ takes $2^{|S|}$ time to compute. The order in which we fill in A is ascending in i, and for each i, we fill in $A(i, S)$ for all S first, in ascending cardinality order, before proceeding to $i+1$. The total running time of our allocation algorithm is

$$\sum_{i=1}^{n} \sum_{k=1}^{m} \binom{m}{k} 2^k = n(3^m - 2^m) = n(2^{\log_2 3m} - 2^m) \leq n(2^m - 1)^2$$

where the last inequality holds for all $m \geq 1$.

Next, we compute the payments of the VCG mechanism. Suppose that bidder i is allocated set S_i. The VCG payment scheme yields

$$p_i(\boldsymbol{v}) = \sum_{j \neq i} v_j(S'_j) - \sum_{j \neq i} v_j(S_j)$$

where S'_j is the set allocated to j, if i did not participate in the auction. To compute this, we need to rerun our dynamic programming algorithm without i, and compute $A(n-1, M)$. This will take at most $(n-1)(2^m - 1)^2$ time, and we need to do it for all bidders i, for a total running time of

$$n(n-1)(2^m - 1)^2 \leq \left(n(2^m - 1)\right)^2$$

Thus, the VCG mechanism can be implemented in time $O\left(n^2(2^m - 1)^2\right)$.

Notice that, for $m = 2$, this running time is $O(n^2)$, while for $m = k$, for some fixed $k > 0$, this running time is $O(n^2 \cdot f(k))$, where $f(k) = 2^k - 1$ is independent of n.

4. (a) (5 points) This problem considers a variation of the Bulow-Klemperer theorem. Consider selling $k \geq 1$ identical items (with at most one given to each bidder) to bidders with valuations drawn iid from F, where F is a regular distribution (i.e., the corresponding ϕ^{-1} is a monotonically increasing function). Prove that for every $n \geq k$, the expected revenue of the Vickrey auction (with no reserve) with $n+k$ bidders is at least that of the Myersons optimal auction for F with n bidders.

   [Hint: Myersons optimal auction will be Vickrey with reserve $\phi^{-1}(0)$, i.e., discard bids below $\phi^{-1}(0)$, give the item to k highest bidders and charge them max$(k+1)^{th}$ highest bid, $\phi^{-1}(0)$]

   (b) (5 points) Consider the reverse auction we briefly talked about in class: A denotes the set of bidders who are willing to sell the spectrum they hold. We say that set $S \subset A$ is feasible, if we can repack $A \setminus S$ in the available range given that S is acquired. Clearly, the set $F \subseteq 2^B$ of feasible sets is upward closed (i.e., supersets of feasible sets are again feasible).

   - Initialize $S = A$.
   - While there is a bidder $i \in S$ such that $S \setminus \{i\}$ is feasible:
     (*) Delete some such bidder i from S such that $S \setminus \{i\}$ is feasible.
   - Return S.

   Suppose we implement the step (*) using a scoring rule, which assigns a number to each bidder i. At each iteration, the bidder with the largest score (whose deletion does not destroy feasibility of S) gets deleted. The score assigned to a bidder i can depend on is bid, the bids of other bidders that have already been deleted, the feasible set F, and the history of what happened in previous iterations. (Note a score is not allowed to depend on the value of the bids of other bidders that have not yet been deleted.) Assume that the scoring rule is increasing – holding everything fixed except for b_i, is score is increasing in its bid b_i. Then, show that the allocation rule above is monotone: for every i and b_{-i}, if i wins with bid bid b_i then she will keep winning with any bid less than b_i.

Solution.

(a) Suppose we want to sell $k \geq 1$ identical items to bidders (with at most one item given to each bidder) with valuations drawn iid from a regular distribution F. Fix $n \geq k$. As in the proof of the Bulow-Klemperer theorem, we define a new auction $\mathcal{A}$ in the following way:

   - Simulate Myerson's optimal auction on bidders $1, \ldots, n$.
   - If not all k items were allocated in the previous step, give the remaining items to bidders $n+1, n+2, \ldots$ for free.

   Note that the expected revenue of $\mathcal{A}$ with $n+k$ bidders is exactly the expected revenue of Myerson's optimal auction for F with n bidders. Thus, it remains to be shown that a Vickrey auction with $n+k$ bidders achieves at least the same revenue as the auction $\mathcal{A}$. Again, following the proof of Bulow-Klemperer, we will argue that the Vickrey auction with $n+k$ bidders maximizes expected revenue over all auctions that allocate all k items. As expected revenue and expected virtual welfare are equivalent, the optimal auction that always allocates the items must give the items to the k bidders with the highest virtual valuations. Since F is regular, the virtual valuation function φ is increasing, so the bidders with the k highest virtual valuations must also have the k highest valuations. In a Vickrey auction, the bidders with the k highest valuations are guaranteed to win the auction. Thus, a Vickrey auction on $n + k$ bidders will give the items to the bidders with the highest virtual valuations, hence a Vickrey auction optimizes expected revenue among auctions which guarantee all k items are allocated. From this, it follows that the expected revenue of a Vickrey auction on $n + k$ bidders is at least the expected revenue of auction $\mathcal{A}$, which is in turn at least the expected revenue of Myerson's optimal auction on n bidders. Hence the expected revenue of a Vickrey auction on $n + k$ bidders is at least the expected revenue of Myerson's optimal auction on n bidders.

(b) If i wins with the bid b_i, then, at any iteration of the greedy selection procedure, either

- i can't be removed because it leads to an infeasible outcome;
- i can be removed but doesn't have the highest score among the bidders who can be safely removed.

Let's now assume i's bid decreased to b_i', while other bidders' bids remained the same, and let's imagine that for some reason the procedure decided to eliminate i at some iteration k. Notice that outcomes of all previous iterations couldn't change they simply didn't depend on i's bid. But if, after $k - 1$ iterations of the procedure with the original bids, removing i leads to an infeasible outcome, then it is still infeasible after $k - 1$ equivalent iterations of the procedure with smaller b_i'. On the other hand, if i can be removed at iteration k of the original procedure, but her score is too small, then she clearly shouldn't be removed in the new procedure since her score is even smaller (scoring is an increasing in b_i, other factors remain the same).

So, the described situation is impossible, i.e., i can't lose by decreasing her bid, and the allocation rule is indeed monotone.

5. (Public good game) Consider a game among n agents, where each agent has some initial money endowment, say agent i has $\$m_i$. There is a magic-pot (government) to which if they put certain amount of money (taxes) it gets multiplied by a factor $\beta > 1$ and gets redistributed among the agents equally (benefits of government facilities). That is if agent i puts $g_i \in [0, m_i]$ in the pot, then out comes $\beta \sum_{i=1}^{n} g_i$ and every agent is allocated $\frac{\beta}{n} \sum_{i=1}^{n} g_i$. Strategy of each agent is to choose $g_i \in [0, m_i]$, the amount to contribute to the magic-pot. When $g = (g_1, \dots, g_n)$ is the strategy profile played, the net utility of the agents is:

$$U_i(g) = (m_i - g_i) + \frac{\beta}{n} \sum_{i=1}^{n} g_i$$

(a) (3 points) What is the social welfare maximizing strategy, *i.e.*, $g^* \in \arg\max_{g: g_i \in [0, m_i]} U_i(g)$.

(b) (2 points) Suppose $\beta < n$ and strategies of all the agents except i is fixed to g_{-i}. What is the best response strategy of agent i?

(c) (3 points) For $\beta < n$, using the previous part, construct a Nash equilibrium. Is it a dominant strategy Nash equilibrium?

(d) (2 points) Construct a Nash equilibrium of this game when $\beta > n$.

Solution.

(a) Let $g = (g_1, \dots, g_n)$ be a strategy profile. We have

$$\begin{aligned}
\sum_{i=1}^{n} U_i(g) &= \sum_{i=1}^{n} \left((m_i - g_i) + \frac{\beta}{n} \sum_{j=1}^{n} g_j \right) \\
&= \sum_{i=1}^{n} (m_i - g_i) + \beta \sum_{j=1}^{n} g_j \\
&= \sum_{i=1}^{n} m_i + (\beta - 1) \sum_{i=1}^{n} g_i \ .
\end{aligned}$$

In order to maximize $\sum_{i=1}^{n} U_i(g)$, since $\beta > 1$, we need to maximize $\sum_{i=1}^{n} g_i$. Therefore, $g_i^* = m_i$ and

$$g^* = (m_1, \dots, m_n) \ .$$

(b) Let $U_i(g, g_{-i})$ be the utility of agent i when the strategy profile of all the agents except i is g_{-i}. We have

$$\begin{aligned}
U_i(g, g_{-i}) &= \sum_{i=1}^{n} (m_i - g_i) + \frac{\beta}{n} \sum_{j=1}^{n} g_j \\
&= \sum_{i=1}^{n} \left(m_i - \left(1 - \frac{\beta}{n}\right) g_i \right) + \frac{\beta}{n} \sum_{\substack{j \neq i \\ j=1}}^{n} g_j \\
&= \sum_{i=1}^{n} \left(m_i - \left(1 - \frac{\beta}{n}\right) g_i \right) + C
\end{aligned}$$

where C is a constant. Since $\beta < n$, the best strategy for agent i in order to maximize $U_i(g, g_{-i})$, is to minimize g_i. Therefore, $g_i = 0$ is the best response.

(c) $g = (0, \ldots, 0)$ is a dominant strategy Nash Equilibrium because from part (b), we have

$$U_i(0, g_{-i}) \geq U_i(g, g_{-i}) \quad \forall i, g_{-i}, g \ .$$

(d) We have

$$\begin{aligned} U_i(g, g_{-i}) &= \sum_{i=1}^{n}(m_i - g_i) + \frac{\beta}{n}\sum_{j=1}^{n} g_j \\ &= \sum_{i=1}^{n}\left(m_i + \left(\frac{\beta}{n} - 1\right) g_i\right) + \frac{\beta}{n}\sum_{\substack{j \neq i \\ j=1}}^{n} g_j \\ &= \sum_{i=1}^{n}\left(m_i + \left(\frac{\beta}{n} - 1\right) g_i\right) + C' \end{aligned}$$

where C' is a constant. Since $\beta > n$, the best strategy for agent i in order to maximize $U_i(g, g_{-i})$, is to maximize g_i. Therefore, $g_i = m_i$ is the best response. By the similar reasoning to part (c) the following strategy profile is a Nash Equilibrium.

$$g = (m_1, \ldots, m_n)$$

6. (Bonus questions)

   (a) In the VCG auction question, give an asymptotic upper bound on the number of items m (with respect to the number of agents n) such that the VCG mechanism can be implemented in time that is polynomial in n. In particular, show that the VCG mechanism can be implemented in time $O(n^2 \cdot f(m))$, for some function f, where $f(m)$ does not depend on n, and give an asymptotic upper bound on m such that this running time is polynomial in n.

   [Hint: Use dynamic programming]

   (b) (Iterated Public goods games) Consider the following iterated variant of the public goods game (Players play a single shot pubic goods game for multiple rounds). Find the Nash equilibrium strategy of each player for each round, and reason whether their strategies for different rounds will be different. We assume that after each round, the payoff to all players together is twice the total investment. That is, if they invest amounts $x_1, x_2, \cdots, x_n$, then the total payoff is $2\sum_i x_i$. Again, let m_i be the money endowment/money of agent i before the game starts.

      i. The total payoff of each round is distributed equally only over the players who invested more than 3/4-value of the average of all investments, and the remaining players get back their investment. That is, if $x_1, x_2, \cdots, x_n$ are the investments of all players, then every player gets back their investment, and the additional payoff $\sum_i x_i$ is distributed equally among all players i whose investment x_i is equal or higher than $(3/4) \cdot \sum_i x_i$.

      ii. In every round, the payoff is distributed in proportion to their total wealth at the beginning of the round. That is, if players start with endowments $w_1, w_2 \cdots, w_n$ and invest $x_1, x_2 \cdots, x_n$, then their wealth after the round is $w_i' = w_i - x_i + 2\frac{w_i}{\sum_{k=1}^{n} w_k} \cdot (\sum_i x_i)$ for all i. Assume that all players start with equal endowments, and play more than one round.

Solution.

(a) In Problem 3, we've shown that the VCG mechanism can be implemented in time $O(n^2 \cdot f(m))$, where $f(m) = 2^m - 1$. Next, we give an asymptotic upper bound on m such that this running time is polynomial in n.

Let $m = O(\log n)$. This implies that there exists a $c > 0$ such that, for large enough n, we have $m \leq c \cdot \log n$. Notice that, for large enough n,

$$f(m) = (2^m - 1)^2 \leq \left(2^{c \cdot \log n} - 1\right)^2 \leq (n^c - 1)^2 \leq n^{2c}$$

Thus, in this case, $f(m) = O(n^{2c})$, and the total running time required to implement the VCG mechanism is

$$n^2 \cdot n^{2c} = n^{2c+2}$$

which is polynomial in n From the above, it is also easy to see that for $m = \omega(\log n)$, the total running time required to implement the VCG mechanism would not be polynomial.

(b) (Solution, and proof sketch)

i. First note that, for any invest profile $(x_1, \ldots, x_n)$ of the agents, at most one agent can have her $x_k \geq 3/4 \sum_{i=1}^{n} x_i$. Therefore, everyone, except at most one, will only get back their investment. Now if for some agent a, $x_a > 0$ in a round but $x_a < 3/4 \sum_i x_i$, then agent a is better off waiting for a round where no one is investing. On the other hand, if an agent has too much money to start with then she is better off investing with someone else with less money.

   Putting the above together, it boils down to case analysis based on the ratio of money across agents. For the simple case, where, $m_i = 1$ for all agents i, the following is a Nash equilibrium: The game goes on for n rounds. For $i \in \{1, \ldots, n\}$, in the ith round, agent i invests $x_i = 1$ while others invest 0. This way, each agent gets 2 units of utility in total, and can not improve on this by deviating unilaterally.

ii. Without loss of generality, let $m_i = 1$ for all $i \in \{1, \ldots, n\}$. Using similar reasoning as single shot public good game with $\beta = 2$, it follows that no agent would invest in the first round. Then, since their wealth is unchanged after the first round, by induction, it follows that at NE no agent invests any money in any round.

   NE through a grim-trigger strategy: Note that, if all the agents invest all their wealth in the first round, then they all double their wealth. Thus, the following is a NE: starting with investing everything in the first round, each agent invest everything in a round if everyone invested all of their wealth in the previous round, otherwise invest nothing from now onwards. The latter threat forces each individual to cooperate and invest all her wealth in every round.

   There may be other such NE that take into account plays in the previous rounds.