University of Illinois **Fall 2017**

ECE 586BH: Problem Set 5: Problems and Solutions
Games with incomplete information, multistage games, VCG mechanisims

Due: Tuesday, November 14, at beginning of class
Reading: Course notes, Section 4.3 and Chapters 5 & 6

1. **[The lemon problem]**
A seller has a used car to sell and a buyer is interested in buying it. Suppose the value θ of the car to seller is random and uniformly distributed on the interval $[0,1]$. The seller knows θ, which can be thought of as the type of the seller. The buyer does not know θ but it is assumed that the value of the car to the buyer is $a + b\theta$, for some constants a and b such that $a \geq 0$ and $1 < a + b < 2$. These constraints ensure that the car is worth more to the buyer than to the seller for any θ, so, payment aside, the socially optimal outcome would be for the car to be transferred from seller to buyer. But can they agree on a price? Suppose the seller, knowing θ, calculates a reserve price $r = r(\theta)$, which is the minimum amount the seller would sell the car for. Suppose the buyer selects a price $p \in [0,1]$ and offers to pay p for the car. If $p \geq r$, the car is sold for price p. Otherwise the car is not sold. The payoff for the seller is $u_s(p,r,\theta) = \theta\mathbf{1}_{\{p<r\}} + p\mathbf{1}_{\{p\geq r\}}$ and the payoff for the buyer is $u_b(p,r,\theta) = (a + b\theta - p)\mathbf{1}_{\{p\geq r\}}$.

   (a) Identify a weakly dominant strategy for the seller.
   Solution: For any θ, p, $u_s(p,r,\theta) \leq \max\{\theta, p\}$, with equality for $r = \theta$. Thus, $r = \theta$ is a weakly dominant strategy.

   (b) Assuming seller uses its weakly dominant strategy, find the price p the buyer should offer to maximize her/his expected payoff. (Suppose the buyer knows a, b, and the distribution of θ).
   Solution: Taking $r(\theta) = \theta$,

$$\mathbb{E}\left[u_b(p,r,\theta)\right] = \mathbb{E}\left[(a + b\theta - p)\mathbf{1}_{\{p\geq\theta\}}\right] = \int_0^p (a + b\theta - p)d\theta = ap + \frac{bp^2}{2} - p^2,$$

   and $\frac{d\mathbb{E}[u_b(p,r,\theta)]}{dr} = a + (b-2)p$. So the best response for the buyer is $p = \frac{a}{(2-b)}$. (This value is in $[0,1]$ by the assumptions on a, b.)

   (c) What happens in the special case $a = 0$ and $1 < b < 2$. Give a simple explanation.
   Solution: The optimal price offered is $p = 0$, so no sale is made with probability one. If buyer were to instead bid p with $0 < p \leq 1$, given the seller accepts, that is, given $\theta \leq p$, the conditional mean value of the car to player 2 would be $\frac{bp}{2}$, which would be less than the price p paid.

2. **[Trigger strategy with limited punishment]**
The payoff matrix for the prisoners' dilemma game is:

		player 2	
		C	D
player 1	C	$1,1$	$-1,2$
	D	$2,-1$	$0,0$

Consider repeated play of the game with discount factor δ. Consider the following limited trigger strategy $s^{T,k}$: Play C in the first stage. In the k stages following the first stage some player plays D, play D. Those k stages represent limited punishment. After those k stages reset the strategy and continue as from the beginning. Find the smallest value δ_k in the interval $[0, 1]$ such that $(s^{T,k}, s^{T,k})$ is a subgame perfect equilibrium for all δ with $\bar{\delta} \le \delta < 1$. Give numerical values for $k = 1$ and $k = 2$ and an equation to solve for $k \ge 1$. Justify your answer by appealing to the one step deviation principle.

Solution:

3. **[Shapley's version of rock-scissors-paper game, revisited]**
Consider Shapley's version of the rock-scissors-paper game:

	R	S	P
R	$0,0$	$1,0$	$0,1$
S	$0,1$	$0,0$	$1,0$
P	$1,0$	$0,1$	$0,0$

(a) Find the set of all correlated equilibria with maximum sum of payoffs. Justify your answer.

Solution: We seek correlated equilibria such that the sum of payoffs is 1 with probability one. So we look for probability distributions of the form shown, for nonnegative constants a through f summing to one:

	R	S	P
R	0	a	b
S	c	0	d
P	e	f	0

In order for player 1, when told to play R, to not have incentive to switch to S, requires $a \ge b$. Similarly, in order for player 1, when told to play S, to not have incentive to switch to P, requires $d \ge c$ Similarly, it must be that $e \ge f$. For similar scenarios involving player 2, correlated equilibrium requires $c \ge e$, $f \ge a$, and $b \ge d$. Combining all six inequalities requires $a \ge b \ge d \ge c \ge e \ge f \ge a$. Consequently, $a = b = c = d = e = f = \frac{1}{6}$ is the unique correlated equilibrium with maximum sum of payoffs. It gives sum of payoffs 1, and payoff vector $\left(\frac{1}{2}, \frac{1}{2}\right)$.

(b) Find all payoff vectors with maximum sum of payoffs that are in the Nash realization region for mixed strategies. Justify your answer.

Solution: The maxmin payoff value for any player is $\underline{v}_i = \frac{1}{3}$, which the player can achieve by using R, S, or P, each with probability $\frac{1}{3}$. The set of realizable payoff vectors with maximum sum of payoffs is obtained by considering all distributions of the form used in the solution of part (a), yielding any payoff vector of the form $(r, 1-r)$ with $0 \le r \le 1$. The Nash realization region is $\{(r, 1-r) : r > \frac{1}{3}, 1-r > \frac{1}{3}\}$, or, equivalently $\{(r, 1-r) : \frac{1}{3} < r < \frac{2}{3}\}$.

(c) Give a pair of behavioral strategies (σ_1, σ_2) for the repeated game with Shapley's R-S-P as stage game such that, as $\delta \to 1$, the payoff vector converges to $(\frac{1}{3}, \frac{2}{3})$. (Hint: Give deterministic scripts for both players to follow up until one player violates the script. Thereafter, players switch to the unique Nash equilibrium in mixed strategies.)

Solution: Following the hint, let the script be (R, P) at times $1, 4, 7, 10, \ldots$ and (R, S) at all other times. If some player does not follow this script in some stage, then the strategy of each player is to use R, S, or P with equal probability thereafter. For some $\underline{\delta}$ this is a Nash equilibrium for all δ with $\underline{\delta} < \delta < 1$, and its payoff to player 1 is: $(1-\delta)\left(1+\delta^3+\delta^6+\cdots\right) = \frac{1-\delta}{1-\delta^3} = \frac{1}{1+\delta+\delta^2}$ and to player 2 is $1 - \frac{1}{1+\delta+\delta^2}$, so the payoff vector converges to $\left(\frac{1}{3}, \frac{2}{3}\right)$ as $\delta \to 1$.

4. **[VCG allocation of a divisible good]**
Suppose one unit of a divisible good, such as wireless bandwidth or power output, is to be divided among n buyers. Suppose the value function of buyer i for quantity x_i is given by $v_i(x_i) = w_i \ln x_i$, where w_i is private information to buyer i with $w_i > 0$. A VCG mechanism is used to determine the allocation $x = (x_1, \ldots, x_n)$ (such that $x_i \geq 0$ with $\sum_i x_i = 1$) and payments $(m_1, \ldots, m_n)$ as a function of the bids. To reduce the amount of communication required, each buyer i submits a single positive scaler bid b_i, which the mechanism interprets as the value function $\tilde{v}_i(x_i) = b_i \ln x_i$. State the allocation and payment rules in as simple a form as possible. To be definite, use the payment rule as described in remark 6.4(c) of the notes. (Hint: To double check your answer you could make sure the payoff of a buyer i, i.e. $w_i \ln x_i(b) - m_i(b)$, is maximized with respect to b_i by setting b_i equal to the true value w_i.)

Solution: The VCG allocation maximizes the social welfare based on the reported functions. That is, it maximizes $\sum_{i=1}^n b_i \ln x_i$ over x subject to $\sum_i x_i = 1$. Introducing a Lagrange multiplier for the constraint and solving yields $x_i = \frac{b_i}{\sum_j b_j}$. That is, the allocation vector is proportional to the bid vector. Equivalently, $x_i = \frac{b_i}{b_i + B_{-i}}$, where $B_{-i} = \sum_{j:j\neq i} b_j$, which better shows how x_i depends on the bid b_i of buyer i. The payment m_i is the maximum social welfare for the other buyers minus their welfare under the allocation with buyer i participating:

$$
\begin{aligned}
m_i(b) &= \sum_{j:j\neq i} b_j \ln \frac{b_j}{B_{-i}} - \sum_{j:j\neq i} b_j \ln \frac{b_j}{B_{-i}+b_i} \\
&= B_{-i} \ln\left(1 + \frac{b_i}{B_{-i}}\right).
\end{aligned}
$$

5. **[Reverse auction VCG mechanism]**
Each day, a certain electric utility company purchases power production for the next day from a set of producers, such as nuclear power plants, coal fired power plants, or solar farms. Each producer i submits a bid, which is a function $p_i(q_i)$ for various values of $q_i \geq 0$. For each value of q_i, the bid $p_i(q_i)$ is the reported cost for the producer to supply q_i units of power for the next day. Suppose the utility company requires a total production of power $Q > 0$. Given the bids $(p_i(\cdot))_{i\in I}$, the utility must decide the quantity of power q_i^* to be provided by the i^{th} producer and the amount m_i to be paid to the i^{th} producer, for each i, such that $\sum_i q_i^* = Q$. This is called a *reverse auction* because there is a single buyer, buying from many sellers. The allocation mechanism should be incentive compatible (IC), so the producers report their true cost of production, individually rational (IR), so the producers should pay no more than their bids, and welfare maximizing, so the power Q is produced with the least cost possible.

   (a) Describe a VCG mechanism for this allocation problem. (Hint: There are two ways to approach this problem. One way is to map to the VCG framework and map back. The cost of production for producer i is the opposite of value, so the value of an allocation $(q_1, \ldots, q_n)$ for producer i as $v_i(q_i) = -p_i(q_i)$. Similarly, the payments to the producers are equivalent to negative payments from the producers to the utility, so $\bar{m}_i(q)$ defined by $\bar{m}_i(q) = -m_i(q)$ could be thought of as the equivalent payment from the producer to the utility. The other way may is to use the basic VCG philosophy in this new context without mapping back and forth.)

Solution: Given the bids, the utility selects the allocation (q_i) that maximizes social welfare, or in other words, is socially efficient. That is,

$$
\begin{aligned}
q^* = \arg\min &\sum_i p_i(q_i) \\
&\text{with respect to } q \text{ subject to} \\
&q_i \geq 0 \\
&\sum_i q_i = Q.
\end{aligned}
$$

The payment to producer i is the minimum sum of reported prices if producer i were not present minus the sum of prices paid to other producers when producer i is present. (This is the increase in total welfare of the other players due to the presence of producer i.):

$$
\begin{aligned}
m_i = \min &\sum_{j:j\neq i} p_j(q_j) - \left(\sum_{j:j\neq i} p_j(q_j^*) \right) \\
&\text{with respect to } q \text{ subject to} \\
&q_j \geq 0 \\
&\sum_{j:j\neq i} q_j = Q.
\end{aligned}
$$

(b) Consider an example with three producers and the following bids: $p_1(q_1) = \begin{cases} 2q_1 & 0 \leq q_1 \leq 1 \\ 6q_1 - 4 & q_1 \geq 1 \end{cases}$, $p_2(q_2) = 4q_2$ for all $q_2 \geq 0$, and $p_3(1) = 1$ and $p_3(2) = 2.5$.

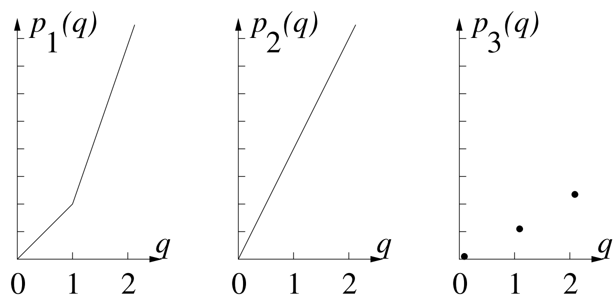

The technology of the third producer limits production to either 0, 1, or 2 units of power. Note that the total amount of power produced must be exactly Q. Describe the allocations (q_1^*, q_2^*, q_3^*) and the payments to the three producers if the total quantity of power required is $Q = 1.5$, for the mechanism you gave in part (a). (Note: The payments may seem to be higher than reasonable.)

Solution: Quantities $q^* = (0.5, 0, 1)$ give the minimum sum of prices, 1+0+1=2.

If producer 1 were not participating, the quantity vector would be (0, 0.5, 1) with sum of prices 0 + 2 + 1=3. Thus, the payment to producer 1 is m_1= (2+1) - (0+1) = 2.

If producer 3 were not participating, the quantity vector would be (1, 0.5, 0) with sum of prices 2 + 2 + 0=4. Thus, the payment to producer 3 is m_3= (2+2) - (1+0) = 3.

(c) Repeat part (b), with the same bids, for $Q = 3$.

Solution: Quantities $q^* = (1, 0, 2)$ give the minimum sum of prices, 2+0+2.5=4.5.

If producer 1 were not participating, the quantity vector would be (0, 1, 2) with sum of prices 0 + 4 + 2.5=6.5. Thus, the payment to producer 1 is m_1= (4+2.5) - (0+2.5) = 4.

If producer 3 were not participating, the quantity vector would be (1, 2, 0) with sum of prices 2 + 8 + 0=10. Thus, the payment to producer 3 is m_3= (2+8) - (2+0) = 8.