These problems are for ***practice only*** *and are not to be turned in. You are responsible for this material for the final exam. You should be able to do the first two problems now and should be able to answer the remaining two problems after the final week's lectures.*

Problem 1. Alice and Bob are trying to meet for lunch. They can each go to the Cafe or the Diner. Alice's office is near the Cafe, so she knows the exact length of time w it would take to wait in line at the cafe. Bob's office is far from the Cafe, so all he knows is that w is distributed $U[0,2]$. All else equal, Alice would be equally happy eating at the Cafe and the Diner, but Bob prefers eating at the Cafe by an amount b that varies from day to day: assume that Bob knows the exactly value of b, while Alice knows only that b is distributed $U[0,3]$, independently of w. In addition, Alice and Bob get a benefit of 1 from having lunch together. Summarizing, with Alice as player 1 and Bob as player 2 the payoff matrix is

$$\begin{array}{ccc} & C & D \\ C & 1-w, 1-w+b & -w, 0 \\ D & 0, -w+b & 1, 1 \end{array}$$

(a) Formally model this situation as an incomplete information game.

Solution. The set of players is given by $\{A, B\}$. The set of types is given by $\Theta = \Theta_A \times \Theta_B$, where $\Theta_A = [0,2]$ and $\Theta_B = [0,3]$. The players' sets of actions is given by $A_A = A_B = \{C, D\}$. The players' payoffs are given by

$$u_A(a,\theta) = \begin{cases} 1-w & \text{if} \quad a = (C,C), \\ -w & \text{if} \quad a = (C,D), \\ 0 & \text{if} \quad a = (D,C), \\ 1 & \text{if} \quad a = (D,D), \end{cases}$$

where $\theta = (w, b)$, and

$$u_B(a,\theta) = \begin{cases} 1-w+b & \text{if} \quad a = (C,C), \\ 0 & \text{if} \quad a = (C,D), \\ -w+b & \text{if} \quad a = (D,C), \\ 1 & \text{if} \quad a = (D,D). \end{cases}$$

The prior p is given by the uniform distribution over Θ.

Find a BNE, and prove that it is unique. How often do Alice and Bob have lunch together?

Solution. We first argue that any BNE must be in threshold strategies. Let $s_i(\theta_i)$ denote the action chosen by type θ_i of player i in a BNE. The expected payoff to type w of Alice from choosing C is given by

$$\mathbb{P}\left(s_B(\theta_B) = C\right)(1-w) - \mathbb{P}\left(s_B(\theta_B) = D\right)w,$$

whereas her payoff from choosing D is given by

$$\mathbb{P}\left(s_B(\theta_B) = D\right).$$

Since the payoff from C is a strictly decreasing function of w and the payoff from D is independent of w, in any BNE, there exists some $w^* \in [0,2]$ such that Alice chooses C if and only if $w \leq w^*$. The expected payoff to type b of Bob from choosing C is given by

$$\mathbb{P}\left(s_A(\theta_A) = C\right)(1 - \mathbb{E}[w|s_A(\theta_A) = C] + b) + \mathbb{P}\left(s_B(\theta_B) = D\right)(-\mathbb{E}[w|s_A(\theta_A) = D] + b),$$

whereas her payoff from choosing D is given by

$$\mathbb{P}\left(s_A(\theta_A) = D\right).$$

The payoff to Bob from C is strictly increasing in b while his payoff from D is independent of b. Therefore, there exists some b^* such that Bob chooses C if and only if $b \geq b^*$.

There are several cases to consider. First suppose that $w^* = 2$. Then, Alice always chooses C. The payoff to type b of Bob from C is then given by b, whereas the payoff from D is given by 0. Therefore, Bob must always choose C. This implies that the payoff to type w of Alice from choosing C is given by $1-w$, and the payoff from choosing D is given by 0. Thus, Alice would choose C if and only if $w \leq 1$, a contradiction.

Next suppose that $w^* = 0$, that is, Alice always chooses D. The payoff to type b of Bob from C is then given by $-1+b$, while the payoff from D is given by 1. So Bob must choose C if an only if $b \geq b^* = 2$. This implies that the payoff to type w of Alice from choosing C is given by $\frac{1}{3}(1-w) - \frac{2}{3}w = \frac{1}{3} - w$, while her payoff from choosing D is given by $\frac{2}{3} > \frac{1}{3} - w$. This confirms the assumption that $w^* = 0$. Thus, a threshold strategy with $w^* = 0$ and $b^* = 2$ is a BNE. We next argue that this is the only BNE.

Next suppose that $b^* = 0$, that is, Bob always chooses C. Then the payoff to type w of Alice from choosing C is given by $1-w$, whereas her payoff to choosing

D is given by 0. Thus, Alice chooses C if and only if $w \leq 1$. The payoff to type b of Bob from choosing C is thus given by $1/2(1/2+b)+1/2(-3/2+b) = -1/2+b$, whereas his payoff from choosing D is given by $1/2$. Therefore, Bob chooses C if and only if $b \geq 1$, a contradiction.

Next suppose that $b^* = 3$, that is, Bob is always choosing D. Then the payoff to type w of Alice from choosing C is given by $-w$, whereas her payoff from choosing D is given by 1. Therefore, Alice will always choose D. But then the payoff to type b of Bob from choosing C is given by $-1+b$, whereas his payoff from choosing D is given by 1. Thus, Bob will choose C if and only if $b \geq 2$, a contradiction.

The only remaining case is where $w^* \in (0,2)$ and $b^* \in (0,3)$. If $w^* \in (0,2)$, Alice needs to be indifferent between C and D when $w = w^*$, that is,

$$\mathbb{P}\left(s_B(\theta_B) = C\right) - w^* = \mathbb{P}\left(s_B(\theta_B) = D\right) = 1 - \mathbb{P}\left(s_B(\theta_B) = C\right).$$

This implies that

$$w^* = 2\mathbb{P}\left(s_B(\theta_B) = C\right) - 1 = 2\frac{3-b^*}{3} - 1.$$

Likewise, Bob needs to be indifferent between C and D when $b = b^*$, that is,

$$\frac{w^*}{2}\left(1 - \frac{w^*}{2} + b^*\right) + \frac{2-w^*}{2}\left(-\frac{2+w^*}{2} + b^*\right) = \frac{2-w^*}{2}.$$

Solving the above the equations for w^* and b^*, we get $w^* = -1$ and $b^* = 3$, a contradiction.

In the unique BNE, Alice chooses D with probability 1 and Bob chooses D with probability 2/3, so Alice and Bob have lunch together 2/3 of the time.

Problem 2. Consider a seller who must sell a single good. There are two potential buyers, each with a valuation for the good that is drawn independently and uniformly from the interval $[0, 1]$. The seller will offer the good using a second-price sealed-bid auction, but he can set a "reserve price" of $r \geq 0$ that modifies the rules of the auction as follows: If both bids are below r then neither bidder obtains the good and it is destroyed. If both bids are at or above r then the regular auction rules prevail. If only one bid is at or above r then that bidder obtains the good and pays r to the seller.

(a) Compute the seller's expected revenue as a function of r.

Solution. By an argument identical to the argument in lecture notes, it is weakly dominant for the bidders to bid their true valuations. We next compute the expected revenue of the seller. With probability r^2, both valuations are below r, and the seller gets zero. With probability $2r(1-r)$, only one valuation is above r, and the seller gets r. With probability $(1-r)^2$, both valuations are above r, and the sellers gets the minimum of the two valuations. The expected value of the minimum of the two valuations conditional on both being above r is equal to the expected value of the minimum of two random variables uniformly distributed over $[r, 1]$. By the formula in the notes, this is exactly equal to $\frac{2r+1}{3}$. Therefore, the expected revenue of the seller is given by

$$0r^2 + r(2r(1-r)) + \frac{2r+1}{3}(1-r)^2 = \frac{1}{3} + r^2 - \frac{4}{3}r^3.$$

(b) What is the optimal value of r for the seller?

Solution. Taking the first-order condition with respect to r and setting the derivative equal to zero we get $r = 0$ and $r = 1/2$. Checking the second-order conditions, we find that $r = 0$ is a local minimum while $r = 1/2$ is the global maximum of the function over $[0, 1]$. Therefore, the optimal value of r is equal to $r = 1/2$, leading to an expected revenue of $5/12$ to the seller.

(c) Intuitively, why does the seller benefit from setting a non-zero reserve price?

Solution. Intuitively, setting a non-zero reserve price is like adding an additional bidder who always bids r. Since the expected payoff to the seller is increasing in the number of bidders, adding a bidder increases the payoff to the seller.

Problem 3 (Herding model)**.** Jackson, Problem 8.6, pg. 254.

Solution. Let's start by computing the posterior probability of action B being bad conditional upon person 3 observing (B, B).

If action B were bad, this would happen with probability $(1-p)^2 + (1-p)p/2$: either both signals come up "good" indepeendently with probability $(1-p)$ or the first comes up "good," the second "bad," and the coin flip goes to "good." On the other hand, conditional upon the action being "good", this happens with probability $p^2 + p(1-p)/2$. Thus, by Bayes' rule, the probability of "good" conditional upon seeing (B, B) is

$$\frac{p^2 + (1-p)p/2}{(1-p)^2 + (1-p)p + p^2} \tag{1}$$

Now let's suppose the third person sees (B, B) and their signal comes up "bad." Their posterior probability of "good" being the truth is

$$\frac{\mathbb{P}(\text{good}|\text{history})\mathbb{P}(\text{signal bad}|\text{good})}{\mathbb{P}(\text{signal bad})} \propto p^2(1-p) + p(1-p)^2/2.$$

Their posterior probability of "bad" being the truth is

$$\frac{\mathbb{P}(\text{bad}|\text{history})\mathbb{P}(\text{signal bad}|\text{bad})}{\mathbb{P}(\text{signal bad})} \propto p(1-p)^2 + p^2(1-p)/2.$$

which is smaller! (We are ignoring the denomonator in (1) above and the denomonator $\mathbb{P}(\text{signal bad})$ from Bayes rule, which are the same for both terms we are comparing). Thus, even if person 3's signal is "bad," they believe "good" is the most likely state of the world, and therefore would also play B.

Persons $4, 5, \ldots$ have no additional information, since all preceding persons after the first two will play B with probability 1 independently of their signal, so, by the exact same calculation, they will also play B.

Problem 4 (DeGroot learning)**.** Consider the DeGroot learning model with N agents with initial belief vector $x(0) = (x_1(0), \ldots, x_N(0))$ and an $N \times N$, non-negative, row stochastic matrix T such that, for every period t, we have

$$x(t) = Tx(t-1).$$

(a) Suppose that $N = 3$ and

$$T = \begin{pmatrix} \frac{3}{5} & \frac{1}{5} & \frac{1}{5} \\ \frac{1}{4} & \frac{1}{2} & \frac{1}{4} \\ 0 & \frac{1}{3} & \frac{2}{3} \end{pmatrix}.$$

What properties of this matrix guarantee that, for any initial belief vector $x(0)$, the limit belief $x^* = \lim_{t\to\infty} x(t)$ is well-defined? Compute x^* as a function of $x(0)$.

Solution. The right-stochastic matrix T is aperiodic and strongly connected, so we know from lecture that there is a unique limiting belief x^* that depends only on $x(0)$. Moreover, it is given by $s^\top x(0)$ where the weight vector s solves

$$sT = s \iff s(T - I) = 0$$

and also satisfies $\sum_i s_i = 1$. Solving this linear system of equations by hand, or plugging them into Wolfram Alpha (or a similar tool) gives us $s = (\frac{5}{22}, \frac{8}{22}, \frac{9}{22})$.

(b) Suppose that $N = 3$ and

$$T = \begin{pmatrix} 0 & \frac{1}{3} & \frac{2}{3} \\ 1 & 0 & 0 \\ 1 & 0 & 0 \end{pmatrix}.$$

As a function of $x(0)$, compute $x(t)$ for every $t \geq 1$. What is going on?

[Hint: First compute T^2 and T^3, then compute $x(1)$, $x(2)$, and $x(3)$. You should notice a pattern.]

Solution. We can compute that

$$T^2 = \begin{pmatrix} 1 & 0 & 0 \\ 0 & \frac{1}{3} & \frac{2}{3} \\ 0 & \frac{1}{3} & \frac{2}{3} \end{pmatrix}, \quad T^3 = T = \begin{pmatrix} 0 & \frac{1}{3} & \frac{2}{3} \\ 1 & 0 & 0 \\ 1 & 0 & 0 \end{pmatrix}$$

Proceeding by inducton, we have $T^k x = T^2 x$ if k is even and $T^k x = Tx$ if k is odd—the system is periodic. This means that $x(t) = x(2) = T^2 x(0)$ for t even and $x(t) = x(1) = Tx(0)$ for t odd. The beliefs look like

$$(x(0), Tx(0), T^2x(0), Tx(0), T^2x(0), Tx(0), T^2x(0), \ldots).$$

(c) Prove that, for any N, if there exists an agent i such that $T_{ii} = 1$ and $T_{ji} > 0$ for all $j \neq i$, then $x_j^* \equiv \lim_{t\to\infty} x_j(t)$ is well-defined and equal to $x_i(0)$ for all $j \neq i$.

[Hint: Let $\Delta(t) = \max_{j\in N} |x_i(t) - x_j(t)|$ *and let* $\underline{T} = \min_{j\neq i} T_{ji}$*. Prove that* $\Delta(t+1) \leq (1 - \underline{T})\Delta(t)$ *for all* t*. Show that this implies that each* $x_j(t)$ *must converge to* $x_i(0)$ *as* $t \to \infty$*.]*

Solution. As suggested by the hint, let's define $\underline{T} = \min_k T_{ki} > 0$ and $\Delta(t) = \max_k |x_k(t) - x_i(t)|$. Firstly, notice that since $T_{ii} = 1$ and the rows of T sum to 1, $T_{ij} = 0$ for $j \neq i$ and we must have $x_i(t) = x_i(0)$ for all t by matrix multiplication. Then we can compute

$$\begin{aligned}
|x_i(t+1) - x_j(t+1)| &= |x_i(t) - x_j(t+1)| \\
&= |x_i(t) - (Tx(t))_j| \\
&= \ x_i(t) - \sum_{k=1}^{n} T_{jk} x_k(t)
\end{aligned}$$

Since $\sum_{k=1}^{n} T_{jk} = 1$ this can be rewritten as

$$= \ \sum_{k=1}^{n} T_{jk}(x_i(t) - x_k(t))$$

By the triangle inequality $|a + b| \leq |a| + |b|$, this is at most

$$\leq \sum_{k=1}^{n} T_{jk} |x_i(t) - x_k(t)|$$

Since $\sum_{k\neq i} T_{jk} = 1 - T_{ji} \leq 1 - \underline{T}$ and $\Delta(t) = \max_k |x_k(t) - x_i(t)|$, this is at most

$$\leq (1 - \underline{T})\Delta(t).$$

Since the above holds for every $j \neq i$, we can deduce that $\Delta(t+1) \leq (1-\underline{T})\Delta(t)$ which means that, since $(1-\underline{T}) < 1$, we must have $\Delta(t) \downarrow 0$ as needed. We conclude that for each k, $x_k(t) \to x_i(0)$ as $t \to \infty$.