Microeconomic Theory II
Preliminary Examination Solutions

Exam date: August 8, 2016

1. **(30 points)** Suppose the government (through its agency, the Internal Revenue Service—IRS) wants to ensure compliance with the tax code. A taxpayer, in filling out a tax return, has a choice, be truthful (T) or lie (L). The IRS can audit the return to determine whether the taxpayer was truthful on his return(A), or can decide not to audit (N) . The taxpayer's payoffs are given by $u_t(T,A)=2$, $u_t(T,N)=3$, $u_t(L,A)=-2$, and $u_t(L,N)=4$. The government's payoffs are given by $u_g(T,A)=-1$, $u_g(T,N)=1$, $u_g(L,A)=0$, and $u_g(L,N)=-2$. (These payoffs capture the idea that the taxpayer has the natural ranking $LN \succ TN \succ TA \succ LA$, while the IRS has the natural ranking $TN \succ LA \succ TA \succ LN$.)

   Suppose the taxpayer does not know if he will be audited when filling out the return, and the IRS cannot condition the audit decision on specifics in the return. In other words, the taxpayer and IRS are effectively making their decisions simultaneously.

   (a) What is the normal form of this game? **[5 points]**

   Solution: The normal form is

		IRS	
		A	N
taxpayer	T	$2,-1$	$3,1$
	L	$-2,0$	$4,-2$

■

   (b) What are the Nash equilibria of this game? **[5 points]**

   Solution: The game has no pure strategy equilibria. To calculate the mixed strategy equilibrium, let q be the probability the taxpayer tells the truth (T) and p be the probability that the IRS audits (A). The taxpayer is indifferent if

$$\begin{aligned} 2p+3(1-p) &= -2p+4(1-p) \\ \iff 3-p &= 4-6p \\ \iff p &= 1/5. \end{aligned}$$

   The IRS is indifferent if

$$\begin{aligned} -q &= q-2(1-q) \\ \iff 4q &= 2 \\ \iff q &= 1/2. \end{aligned}$$

   Hence, the unique Nash equilibrium is $\sigma^* = (\frac{1}{2}\circ T + \frac{1}{2}\circ L,\ \frac{1}{5}\circ A + \frac{4}{5}\circ N)$. ■

Suppose now the IRS can commit to a probability of auditing the return *before* the taxpayer fills out his return, and that the taxpayer *observes* this probability. (This commitment may involve the early hiring of IRS auditors, for example.)

(c) What is the normal form of this game? **[5 points]**

Solution: The IRS commits to $p \in [0,1]$, the probability of A, audit. The taxpayer chooses whether to cheat as a function of the probability committed to, i.e., $s : [0,1] \to \{T, L\}$. Hence, the strategy space for the IRS is $S_{IRS} = [0,1]$ and the strategy space for the taxpayer is the set of all (measurable) functions $s : [0,1] \to \{T, L\}$. The payoffs in the normal form are, for the taxpayer,

$$u_T(s,p) = p\,u_t(s(p),A) + (1-p)u_t(s(p),N)$$

and for the IRS,

$$u_{IRS}(s,p) = p\,u_{IRS}(s(p),A) + (1-p)u_{IRS}(s(p),N).$$

■

(d) What is the backward induction solution of this game? Is it unique? **[10 points]**

Solution: The taxpayer has T as a best reply iff

$$\begin{aligned} 2p + 3(1-p) &\geq -2p + 4(1-p) \\ \iff 3 - p &\geq 4 - 6p \\ \iff p &\geq 1/5, \end{aligned}$$

and has L as a best reply iff $p \leq 1/5$. Let s^* denote any strategy for the taxpayer satisfying

$$s^*(p) = \begin{cases} T, & \text{if } p > 1/5, \\ L, & \text{if } p < 1/5. \end{cases}$$

Backward inducting, the IRS's payoff is

$$u_g(s^*, p) = \begin{cases} -2(1-p), & \text{if } p < 1/5, \\ 1 - 2p, & \text{if } p > 1/5. \end{cases}$$

The IRS's payoff is increasing in p for $p < 1/5$ and decreasing in p for $p > 1/5$, and is strictly larger in a neighborhood of $1/5$ for $p > 1/5$ (since the taxpayer is truthful). The IRS does not have a well-defined best response if the taxpayer lies at $p = 1/5$ with positive probability since the IRS strictly prefers to induce truthfulness with the smallest possible audit probability. Consequently, the unique backward induction solution is $(1/5, \hat{s})$, where $\hat{s}(p) = s^*(p)$ for all $p \neq 1/5$ and $\hat{s}(1/5) = T$. ■

(e) Will an announcement by the IRS of intended auditing probabilities serve the same function as the commitment? **[5 points]**

Solution: An announcement cannot serve the same function, because the taxpayer understands that, if the taxpayer were to believe any announcement, the government has no incentive to carry out the announcement. Thus, there is a unique subgame perfect equilibrium distribution over $\{T, L\} \times \{A, N\}$ given by the equilibrium from part b (the announcement is ignored, and so each announcement corresponds to a different equilibrium).

■

2. **(30 points)** Consider the following infinitely repeated game with three players. In each period, player 1 is simultaneously playing a prisoners' dilemma with player 2 and another prisoners' dilemma with player 3 (so that within the stage game, all players simultaneously choose actions). The stage games are given by

player 1 \ player 2	E_2	S_2
E_1	$3,3$	$-1,4$
S_1	$4,-1$	$0,0$

player 1 \ player 3	E_3	S_3
E_1	$2,2$	$-1,4$
S_1	$4,-2$	$0,0$

(Note that the two games are *not* the same!) All players have a common discount factor $\delta \in (0,1)$. Player 1's payoffs are the *sum* of payoffs in the two games.

Suppose there is a technological restriction that forces player 1 to choose the same action in the two different games (i.e., player 1 plays E_1 against player 2 if and only if he does so against player 3).

(a) Suppose there is perfect monitoring of all players' past actions. Describe the "grim-trigger" strategy profile that induces the outcome path in which $E_1E_2E_3$ is played in every period, and which is subgame perfect for large δ. What is the smallest value of δ for which the profile is a subgame perfect equilibrium? [Remember to provide support for your answer.] **[5 points]**

Solution: Grim trigger is that player i plays E_i in the first period and thereafter as long as $E_1E_2E_3$ always played, and otherwise play S_i, $i = 1,2,3$. After a deviation, a static Nash equilibrium is being played, and so those incentive constraints are trivially satisfied for all δ. The incentive constraints on the equilibrium path are

$$5 \geq 8(1-\delta) \iff \delta \geq \frac{3}{8} \qquad \text{for player 1,}$$
$$3 \geq 4(1-\delta) \iff \delta \geq \frac{1}{4} \qquad \text{for player 2, and}$$
$$2 \geq 4(1-\delta) \iff \delta \geq \frac{1}{2} \qquad \text{for player 3.}$$

Thus the profile is an equilibrium iff $\delta \geq \frac{1}{2}$. ■

Suppose now that while player 1 observes the past actions of players 2 and 3, players 2 and 3 only observe the past actions of player 1. Players 2 and 3 do not observe each others past actions. Grim trigger is now: player 1 plays E_1 in the first period and then as long as $E_1E_2E_3$ always played, otherwise play S_1; player i plays E_i in the first period and then as long as i observes only E_1E_i, otherwise play S_i, $i = 2, 3$.

(b) For what values of δ is the grim trigger profile just described a Nash equilibrium? Why is every Nash equilibrium of this game subgame perfect? **[5 points]**

Solution: The incentive constraints for all players on the equilibrium path are unchanged from part 2(a), and so we again require $\delta \geq \frac{1}{2}$.

Every Nash equilibrium is subgame perfect because the only subgame is the original game (players 2 and 3 have nontrivial information after period 1 that cannot be "disentangled." ■

(c) Suppose δ is such that the grim trigger profile is a Nash equilibrium.

i. Give an intuitive description of the specific restrictions that sequential rationality imposes on the grim trigger profile.. **[5 points]**
Solution: Sequential rationality requires that it be optimal for all players to follow the profile after a deviation. In particular, after a deviation by player 2, say, if player 1 continues to play E_1, then player 3 does not know a deviation has occurred and continues to play E_3. The profile requires that it be optimal for player 1 to play S_1 after a deviation by player 2, and so trigger future S_3 by player 3. ■

ii. Prove that the grim trigger profile is not sequentially rational for large δ by showing that one of the players has a profitable one shot deviation. **[5 points]**
Solution: Neither player 2 nor 3 has a profitable one shot deviation for large δ: Suppose player 1 has deviated. Then player i, $i = 2, 3$, expects S_1 thereafter, and so S_i is optimal for large δ.
Consider now player 1's history that is reached by a deviation by player 2. Then player 1 is supposed to play S_1 thereafter, and since player 3 plays E_3 (having not observed the deviation by 2) for one period (after which S_3 is played), this yields a payoff of

$$4(1-\delta).$$

On the other hand, ignoring the deviation and playing E_1 implies player 3 does not realize a deviation has occurred and so 3 will play E_3 for one more period. The one-shot deviation yields the payoff

$$(1-\delta) + 4\delta(1-\delta),$$

which is strictly larger than $4(1-\delta)$ (and so the deviation is profitable) iff $\delta > \frac{3}{4}$. ■

iii. For what values of δ is the profile sequentially rational? **[5 points]**
Solution: From the earlier analysis, we need $\delta \geq \frac{1}{2}$, so that players 2 and 3 do not have any profitable one shot deviations. This is also sufficient so that player 1 does not have a profitable one shot deviation on the equilibrium path.
We need $\delta \leq \frac{3}{4}$ so that player 1 does not have a profitable one shot deviation after a deviation by player 2.
We finally need to rule out a profitable one shot deviation for player 1 after a deviation by player 3. Suppose player 3 deviates. Then player 1 is supposed to play S_1 thereafter, and since player 2 plays E_2 (having not observed the deviation by 3) for one period (after which S_2 is played), this yields a payoff of

$$4(1-\delta).$$

On the other hand, ignoring the deviation and playing E_1 implies player 2 does not realize a deviation has occurred and so 2 will play E_2 for one more period. The one-shot deviation yields the payoff

$$2(1-\delta)+4\delta(1-\delta),$$

which is strictly larger than $4(1-\delta)$ (and so the deviation is profitable) iff $\delta > \frac{1}{2}$. and again, this is a profitable deviation iff $\delta > \frac{2}{3}$.)
Thus, the profile is sequentially rational for $\delta = \frac{1}{2}$. ∎

(d) Suppose now there is no technological restriction: player 1 is free to choose different actions in the two prisoners' dilemmas. Let a_1^j denote the action a_1 played by player 1 in the game with player j, $j = 2,3$. Maintain the observability assumption introduced just before part 2(b). Describe a grim trigger like subgame perfect strategy profile that supports $E_1^2 E_1^3 E_2 E_3$ in every period for large δ. Explain why your answer differs from your analysis in part 2(c). **[5 points]**

Solution: Player 1 can now treat the two prisoners' dilemma games independently. The profile specifies player 1 plays $E_1^2 E_1^3$ in the first period and then as long as $E_1^2 E_1^3 E_2 E_3$ always played, play S_1^i if ever S_i or S_1^i is played; player i plays E_i in the first period and then as long as i observes only $E_1^i E_i$, otherwise play S_i, $i = 2,3$.

Since we can treat the two prisoners' dilemma games independently, the analysis from part 2(a) applies game by game, and so the profile is subgame perfect iff $\delta \geq \frac{1}{2}$. ∎

3. **(30 points)** Suppose that the payoff to a firm from hiring a worker of type θ with education e at wage w is

$$f(e,\theta) - w = 3e\theta - w.$$

The utility of a worker of type θ with education e receiving a wage w is

$$w - c(e,\theta) = w - \frac{e^3}{\theta}.$$

The worker's ability is privately known by the worker. There are at least two firms. The worker (knowing his ability) first chooses an education level $e \in \mathbb{R}_+$; firms then compete for the worker by simultaneously announcing a wage; finally the worker chooses a firm. Treat the wage determination as in class, a function $w : \mathbb{R}_+ \to \mathbb{R}_+$ determining wage as a function of eduction.

Suppose the support of the firms' prior beliefs ρ on θ is $\Theta = \{\theta_L, \theta_H\}$ where $\theta_L = 1$ and $\theta_H = 3$.

(a) What is the full information efficient education level for each type of worker? **[5 points]**

Solution: The social surplus from a worker of type θ taking education level e is:

$$3e\theta - \frac{e^3}{\theta}.$$

Taking first order conditions,

$$3\theta - \frac{3e^2}{\theta} = 0,$$
$$\implies e = \theta,$$

i.e., the full information efficient education level for worker of type θ is $e = \theta$. ■

(b) Is there a perfect Bayesian equilibrium in which both types of worker choose their full information eduction level? Be sure to verify that all the incentive constraints are satisfied. **[5 points]**

Solution: The strategy profile $\{e, w\}$ constitutes a perfect Bayesian equilibrium, where

$$e(\theta) = \theta; w(e) = \begin{cases} 3e & \text{if } e \neq 3, \\ 9e & \text{if } e = 3, \end{cases}$$

with beliefs $\mu(e)$ which place all mass on $\theta = 1$ if $e \neq 3$, and all mass on $\theta = 3$ if $e = 3$.

To verify incentive compatibility observe that on path, a worker of type 1 has a net utility of 2. However, if she deviates to type 3's education level, she makes a net utility of 0, so this deviation is not profitable.

Similarly, a worker of type 3 on path has a net utility of 18, while by deviating to type 1's education level, she makes a net utility $\frac{8}{3}$. Therefore, this deviation is not profitable.

Next, to verify that this is a PBE, one should carefully check that neither type of worker wishes to deviate to any other education level.

For the low type—deviating to any education level $e' \neq 3$ will still induce a belief that the agent is of low type. By part a), $e = 1$ is education level for an agent of low type when the firms infer she is of low type. So any deviation is not profitable.

For the high type, deviating to any education level $e' \neq 3$ will induce a belief that the agent is of low type. The optimal such deviation solves $\arg\max 3e - e^3/3$. Taking first order conditions, we get $e' = \sqrt{3}$ is the optimal such deviation, plugging back in a simple calculation gets that $3\sqrt{3} - (\sqrt{3})^3/3 = 2\sqrt{3} < 18$, so no such deviation is profitable. ■

(c) Suppose e_L is the education level undertaken by θ_L, while e_H is the education level taken by θ_H in a separating Perfect Bayesian Equilibrium. What can you conclude about e_H and e_L? Be as precise as possible. **[10 points]**

Solution: In any separating PBE, the low type of agent takes his full information optimal education level, i.e. $e_L = 1$. Why? Suppose not, suppose there is a separating equilibrium where the low type of agent takes an education level e'_L different from 1. On observing e'_L, firms must believe that the agent is type θ_L for sure, and therefore offer a wage $3e'_L$. In a PBE, the most punishing response to an education level different from e_L and e_H is is that firms believe the worker is of type θ_L, so a deviation to the education level $e''_L = 1$ is profitable—we have already argued in part a) that

$$3e''_L - (e''_L)^3/\theta_L = \max_e 3e - e^3/\theta_L,$$

and since the objective function is strictly concave,

$$3e_L - e_L^3/\theta_L > 3e'_L - e'^3_L/\theta_L.$$

Next, note that since this is a separating equilibrium, the firms pay $9e_H$ when seeing e_H (and, as we argued before, 3 when seeing $e_L = 1$).

It is sufficient to consider beliefs in the following form:

$$\mu(e) = \begin{cases} 1 \circ \theta_H & \text{if } e = e_H \\ 1 \circ \theta_L & \text{otherwise} \end{cases}.$$

In this form of proposed equilibria, the low type worker gets 2 and the high type worker gets $9e_H - e_H^3/3$. Given the beliefs specified, the best possible deviation for type θ_L is e_H; to get the best possible deviation for type θ_H, we solve

$$\max_e 3e - \frac{e^3}{3}$$

The best possible deviation is to an education level $\sqrt{3}$ and the high type will get $2\sqrt{3}$. The sufficient and necessary conditions such that the strategies of both types are incentive compatible are thus

$$2 \geq 9e_H - e_H^3;$$

$$9e_H - e_H^3/3 \geq 2\sqrt{3}.$$

■

(d) Suppose the firms' prior beliefs ρ are that the worker has type θ_H with probability $\frac{2}{3}$, and type θ_L with probability $\frac{1}{3}$. Is there a pooling Perfect Bayesian equilibrium in this setting? If yes, describe a pooling PBE and argue that it is one. If not, why not? **[10 points]**

Solution: Yes, a pooling PBE exists.

Consider a putative PBE in which both types of worker take the same education level $e^\star$. Further firms' posterior on observing this education level is the same as their prior, if they observe any other education level their posterior puts all mass on θ_L.

Therefore if the firms observe $e^\star$, they believe the agent is of high type with probability $\frac{2}{3}$ and low type with probability $\frac{1}{3}$. Their expected payoff from hiring such a worker and offering him a wage of w is:

$$\begin{aligned}&\frac{2}{3}\times(3e^\star\theta_H)+\frac{1}{3}\times(3e^\star\theta_L)-w\\=&7e^\star-w.\end{aligned}$$

Since firms are competetive they make 0 profits, i.e. both offer a wage of $7e^\star$. For any other education level $e\neq e^\star$, they offer a wage of $3e$.

Note that a worker of type θ_L who takes the equilibrium education level gets a surplus of $7e^\star-e^{\star 3}$. His optimal education level were he to deviate is $e=1$ for a net surplus of 2. Therefore for our putative pooling equilibrium to be one, it must be the case that

$$7e^\star-e^{\star 3}\geq 2.$$

Similarly, if a θ_H type deviates, his optimal education level is $\sqrt{3}$, for a net surplus of $2\sqrt{3}$ Therefore, it must be the case that

$$7e^\star-\frac{e^{\star 3}}{3}\geq 2\sqrt{3}.$$

Picking $e^\star=2$ (say) clearly satisfies both inequalities. ■

4. **(30 points)** Mussa Ltd is faced with a single buyer, Mr. Rosen. It is known that Mr. Rosen has constant marginal utility θ for its product. In particular, if Mr. Rosen has a marginal utility of θ, buys q units of the product, and pays p, his net utility is:

$$u(q,p,\theta)=\theta q-p.$$

Mr. Rosen's outside option is normalized to 0.

Mussa Ltd has a cost function $c(q)=\frac{1}{2}cq^2$, so its total profit if it sells q units for p is:

$$\pi(p,q)=p-\frac{1}{2}cq^2.$$

(a) Suppose Mussa Ltd knows Mr. Rosen's marginal utility θ. Describe, as a function of θ, the quantity $q(\theta)$ and price $p(\theta)$ it offers Mr. Rosen to maximize its profit. **[5 points]**

Solution: Clearly, for any quantity q offered to Mr. Rosen, it is optimal for Mussa Ltd to charge exactly $p=q\theta$ to extract all the surplus/ leave Mr. Rosen indifferent between purchasing and not.

It remains to determine the quantity offered. By our previous argument, when Mr. Rosen has marginal utility θ, q offered should solve,

$$\max q\theta - \frac{1}{2}cq^2.$$

Taking first order conditions, we see that $q(\theta) = \frac{\theta}{c}$, and therefore that $p(\theta) = \frac{\theta^2}{c}$. ■

(b) Suppose now that Mussa Ltd does not know Mr. Rosen's marginal utility, and instead believes it to be drawn from the uniform distribution over $[0,1]$.

Due to complexity issues, Mussa Ltd offers a menu of 2 bundles: a bundle (q_1, p_1) intended for Mr. Rosen if his type $\theta \in [\theta_1, \theta_2)$, and a bundle (q_2, p_2) if his type is in $[\theta_2, 1]$, where $0 < \theta_1 < \theta_2 < 1$. The values of θ_1 and θ_2 are fixed (e.g., provided by the marketing department); you only need to supply the optimal bundles given the values of θ_1 and θ_2.

Describe carefully Mussa Ltd's expected profit maximization problem (i.e., what bundles (q_1, p_1) and (q_2, p_2) to offer) when he would like $[\theta_1, \theta_2)$ to purchase bundle 1, $[\theta_2, 1]$ to purchase bundle 2, and $[0, \theta_1)$ to purchase nothing. (**Hint:** the relevant IC/IR constraints here are that the intended bundle purchase should be optimal for those types among the available bundles and the option to purchase nothing.)

Solve for the profit maximizing bundle given θ_1, θ_2. **REMEMBER θ_1 AND θ_2 ARE FIXED.** You may assume that $\theta_1 + \theta_2 > 1$. **[25 points]**

Solution: Mussa Ltd's expected profit maximization problem can be written as:

$$\begin{aligned}
\max_{(q_1,p_1),(q_2,p_2)} & (\theta_2 - \theta_1)\left(p_1 - \frac{1}{2}cq_1^2\right) + (1-\theta_2)\left(p_2 - \frac{1}{2}cq_2^2\right), \\
& \text{s.t. "Incentive Comapatibility"}: \\
\forall \theta \in [\theta_2, 1] : & \theta q_2 - p_2 \geq \theta q_1 - p_1, \\
\forall \theta \in [\theta_1, \theta_2) : & \theta q_1 - p_1 \geq \theta q_2 - p_2, \\
\forall \theta \in [0, \theta_1) : & 0 \geq \max_{i=1,2}(\theta q_i - p_i) \\
& \text{"Individual Rationality"} \\
\forall \theta \in [\theta_2, 1] : & \theta q_2 - p_2 \geq 0, \\
\forall \theta \in [\theta_1, \theta_2) : & \theta q_1 - p_1 \geq 0.
\end{aligned}$$

To solve this problem, let us first consider the IC constraints. Rearranging the first two IC constraints, we have

$$\forall \theta \in [\theta_2, 1]: \quad \theta(q_2 - q_1) \geq p_2 - p_1 \quad \text{and} \tag{1}$$

$$\forall \theta \in [\theta_1, \theta_2): \quad p_2 - p_1 \geq \theta(q_2 - q_1). \tag{2}$$

The only way this is possible is if $q_2 - q_1 \geq 0$. Further, from (1) we have

$$\theta_2(q_2 - q_1) \geq p_2 - p_1,$$

and taking limits as $\theta \nearrow \theta_2$ in (2) gives

$$\theta_2(q_2 - q_1) \leq p_2 - p_1,$$

which implies

$$\theta_2(q_2 - q_1) = p_2 - p_1.$$

Similarly, for the third IC and second IR constraints to jointly be satisfied, it must be the case that:

$$p_1 = q_1\theta_1.$$

By observation, if $q_2 \geq q_1$, these prices satisfy the remaining constraints.
Plugging into the objective function, we have:

$$\begin{aligned}&(\theta_2-\theta_1)\left(\theta_1 q_1 - \frac{1}{2}cq_1^2\right) + (1-\theta_2)\left(\theta_1 q_1 + \theta_2(q_2-q_1) - \frac{1}{2}cq_2^2\right),\\ =&(\theta_2-\theta_1)\left((\theta_1-(1-\theta_2))q_1 - \frac{1}{2}cq_1^2\right) + (1-\theta_2)\left(\theta_2 q_2 - \frac{1}{2}cq_2^2\right)\end{aligned}$$

Ignoring the monotonicity constraint, we can maximize pointwise to get:

$$\begin{aligned}q_2 &= \frac{\theta_2}{c},\\ q_1 &= \frac{\theta_1+\theta_2-1}{c} \quad \left(< \frac{\theta_1}{c}\right).\end{aligned}$$

By observation, $q_1 < q_2$ so these quantities coulpled with the prices computed above constitute the optimal menu. ■