Economics 203: Final Review Session

Jeffrey Naecker

March 19, 2012

1 Logistics

- Your final exam for this class is from **7:00 pm to 10:00 pm on Wednesday, March 21, in EDUC128**.
- All problem sets have been graded and are avaible for pick-up in today's review session. I've also released solutions for all the problems, which are available on Coursework.
- I will hold extra office hours this week: Tuesday (tomorrow), from 1:00 to 3:00 pm, in Landau 218.

2 Tips and Reminders

- Don't forget the formula for the number of strategies a player must have. As soon as you've drawn the extensive form of a game, I recommend counting info sets and actions, and then calculating how big each player's strategy set should be.
- If you are asked to give an equilibrium, make sure you list strategies for every player. An equilibrium is more than just an outcome; it is a *complete contingent plan*. It must specify an action for every player at each of that player's info sets, even if that info set is not reached on the equilibrium path.
- For any WPBE/PBE/SE, you must specify beliefs in addition to strategies. For an information set that is just a singleton, beliefs are degenerate and you don't have to specify them. But for any non-singleton information set (that is, one that contains more than one node), you need to specify beliefs for than info set as part of any WPBE/PBE/SE.

3 Problems

Problem 1. *Problem Bank #98.*

Student/workers are of two types, L and H. There are at least two potential employers. The incremental profits earned from hiring a worker of type i is π_i, with $\pi_H > \pi_L$. Events unfold as follows: (1) Students observe their types (employers do not). (2) Students choose education, e, which is costly and observed by employers, but is not productive. (3) Employers name wages for prospective workers (currently students), w, contingent on observed education (Bertrand competition). (4) Students choose employers and become workers. The utility for a student/worker of type i receiving a wage w is given by $U(w, e) = w - c_i(e)$. The payoff for the employer of a type i student/worker paying a wage w is given by $\pi_i - w$. Assume: $c'_H(e) < c'_L(e)$ for all $e > 0$. Throughout, focus on pure strategy equilibria.

(a) Does this model satisfy the Spence-Mirrlees single-crossing property? Justify your answer.

(b) What level of education is obtained by a type L worker in a separating equilibrium? Derive expressions indicating the minimum and maximum levels of education obtained by a type H worker in separating equilibria.

(c) Derive expressions indicating the minimum and maximum levels of education obtained by all workers in pooling equilibria.

(d) Which equilibrium is likely to prevail? Justify your answer.

(e) Now suppose the government imposes a tax on education, t (so that the amount collected from an individual who obtains e years of education is te). Also suppose that the government keeps all of the tax revenue.

  (i) How does this tax alter the conditions for separating equilibria and pooling equilibria that you derived in parts (b) and (c)? (Derive new conditions.)

  (ii) Focus on the most efficient separating equilibrium. Derive an expression for the derivative of the utility of a type H worker with respect to t, the tax rate. Sign the derivative. Do the same for a type L worker. How does the tax affect the well-being of the workers?

(f) Suppose as in part (e) that the government imposes a tax on education, t (so that the amount collected from an individual who obtains e years of education is te). Now, however, suppose that the government distributes all revenue back to the workers. Because it cannot distinguish a type L worker from a type H worker, all workers must receive the same lump-sum payment. (Assume there are many workers so that each worker ignores the impact of his decision on the size of the lump-sum.)

  (i) How does this tax-transfer system alter the conditions for separating equilibria and pooling equilibria that you derived in parts (b) and (c)? (Derive new conditions.)

  (ii) Focus on the most efficient separating equilibrium. Derive an expression for the derivative of the utility of a type H worker with respect to t, the tax rate. Sign the derivative. Do the same for a type L worker. How does the tax affect the well-being of the workers? Are there conditions under which the tax yields a Pareto improvement?

Problem 2. *Problem Bank #57 (2011 Final).*

On the first day of class, the professor announces whether or not there will be a final exam. The possible announcements are Y and N, where Y stands for the announcement: "Yes, there will be an exam." N stands for the announcement: "No, there will not be an exam."

Suppose that after the announcement, the professor decides secretly (without telling the students), whether or not he will come to the final exam with a stack of exams (e) or a stack of pizzas for a party (p). Thus, the students only hear the professor's announcement (Y or N), but do NOT know whether they'll have an exam or a party (e or p). The students must decide during finals week whether or not to study (s for study, n for not study). For simplicity, assume there is one student.

(a) Draw the extensive form of this game, without any payoffs at the terminal nodes.

(b) What is each player's strategy set?

Suppose that the utilities of the professor and student depend only on whether the student studied and whether the professor wrote an exam. The announcement itself does not affect the utilities of either player. Suppose that they have the following payoffs for the given combinations of choosing whether or not to study and choosing whether or not to show up with a test:

	(e,s)	(e,n)	(p,s)	(p,n)
Professor	3	1	4	2
Student	3	1	2	4

(c) Add the payoffs to the terminal nodes of the extensive form.

(d) Does there exist a pure-strategy Nash equilibrium of this game in which the professor plays e at some information set? If so, describe such an equilibrium. If not, explain why not.

(e) Does there exist a subgame perfect Nash equilibrium of this game in which the professor plays e at some information set? If so, describe such an equilibrium. If not, explain why not.

(f) Give an example of a mixed-strategy Nash equilibrium in which the professor puts strictly positive weight on Y and N.

(g) Characterize the set of perfect Bayesian equilibria of this game.

(h) Characterize the set of sequential equilibria of this game.

Now suppose that the chair of the economics department frowns upon professors lying to students and will fire the professor if he lies; thus, the professor gets a payoff of 0 if his action does not match his announcement (regardless of whether or not the student studies). If the professor tells the truth, then his payoff is determined using the table above. The student's payoff is determined by the table above regardless of whether the professor has told the truth or lied.

(i) Draw the extensive form again, but modify the payoffs at the terminal nodes where appropriate. Note, the shape of the extensive form diagram hasn't changed.

(j) Now, using the new diagram, characterize the set of perfect Bayesian equilibria.

Problem 3. *Problem Bank #89.*

There are n people in a group. Each person has friends and enemies in the group (everyone's friends and enemies are common knowledge). We can illustrate these relationships with the following type of graph:

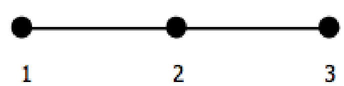

Here, a "link" indicates that people are friends (people are automatically linked with themselves). No link indicates that they are enemies. So, player 1 is friends with player 2 (and player 2 is friends with player 1), but player 1 is enemies with player 3 (and vice versa). Player 2 is friends with player 3 (and vice versa).

We will analyze a game that involves passing on a secret. The secret is "hard information": players simply choose whether or not to pass it on; they cannot lie about it. At first, only player 1 knows the secret. The game is played as follows:

1) Player 1 chooses a subset of her friends (possibly empty), to whom she passes on the secret. If the chosen subset is empty, the game ends. If it is non-empty, the game continues, as follows.

2) Those who just received the information choose subsets of their uninformed friends (possibly empty), to whom they pass on the secret. If the chosen subsets are all empty, the game ends. If it is non-empty, the game continues by repeating this step.

Let $x_i = 1$ if person i learns the secret by the end of the game, and $x_i = 0$ if she does not. Each person has the following utility function: $u_i(x_1, x_2, x_3, ...) = \alpha \sum_{j \in F_i} x_j - \beta \sum_{k \in E_i} x_k$ where $\alpha \geq 0$ and $\beta \geq 0$, and F_i and E_i are the sets of i's friends and enemies, respectively.

For parts a-d and f, assume that $\alpha = 1$ and $\beta = 2$.

(a) Suppose the group consists of only two people, who are friends. Note that player 2 has no choices. What is player 1's strategy set? Does player 1 have a dominant strategy? What is the SPNE of this game? (Warning: don't look for a trick or something deep here. It's just a simple first step.)

(b) Suppose the group consists of 3 people whose relationships correspond to the graph below. What are the SPNE of this game? What are the NE of this game?

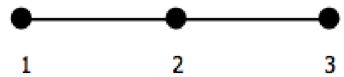

(c) Suppose the group consists of 4 people whose relationships correspond to the graph below. What are the SPNE of this game? What are the NE of this game?

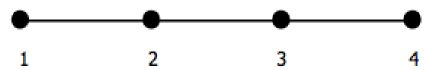

(d) Suppose the group consists of n people arrayed in a line, as above (see the graph below). What are the SPNE of this game? How does it depend on n?

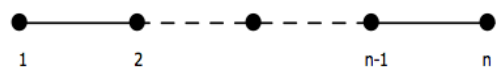

(e) Solve question (d) for arbitrary α and β. Show how the SPNE of the game depend on n, α and β. Your answer does not need to be completely formal, but try to characterize the way in which the equilibrium depends on the parameters.

(f) Now, once again, assume that $\alpha = 1$ and $\beta = 2$. Consider a repeated version of the game in problem (c) with 3 players:

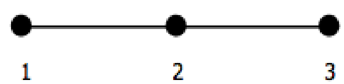

So, at every time t, player 1 learns a new secret and must decide whether to tell player 2 (who, if told, must decide whether to tell player 3). Each players' utility is the sum of the payoffs in each stage game discounted at the geometric rate δ.

(i) Suppose the number of repetitions is finite. Is there a SPNE in which player 1 tells player 2 the secret? If so, describe the equilibrium. If not, explain why.

(ii) Suppose the number of repetitions is infinite. Also suppose player 1 CANNOT verify whether player 2 tells player 3 the secret (Even so, player 1's utility still depends on whether player 3 learns the secret.). Is there a SPNE in which player 1 tells player 2 the secret? If so, what is the lowest discount factor for which such an equilibrium exists?

(iii) Suppose again that the number of repetitions is infinite. If player 2 passes the secret on to player 3, player 1 discovers that fact with probability p; moreover, player 2 always knows whether player 1 has made that discovery. Is there a SPNE in which player 1 tells player 2 the secret? If so, what is the lowest discount factor for which such an equilibrium exists (as a function of p)?

4 Solutions

Solution 1.

(a) Yes, the single-crossing property is satisfied. An easy way to see this is to note that indifference curves for type i are given by $w = c_i(e) + K$. Differentiating, we get $\frac{\partial w}{\partial e} = c_i'(e)$. Then the assumption $c_H'(e) < c_L'(e)$ for all e gets us

$$\frac{\partial w}{\partial e}|_{i=H} < \frac{\partial w}{\partial e}|_{i=L}$$

for all e, as required for the single-crossing property.

(b) First, some notation. In a separating equilibrium, let's say that L types choose education e_L and H types choose e_H. Since the types are identified, we know that $w_L \equiv w(e_L) = \pi_L$ and $w_H \equiv w(e_H) = \pi_H$.

Suppose $e_L > 0$ in a separating equilibrium. Then on the equilibrium path, L types get payoff $\pi_L - c_L(e_L)$. But note that they can deviate to $e = 0$, which reduces their costs and does not lower their wages (since $w(e) \in [\pi_L, \pi_H]$). Then this cannot be an equilibrium. Thus we must have $e_L = 0$.

Next, note that we cannot have e_H too high, or H types will then prefer to deviate to e_L because the cost of education is too great. Thus we need the payoff from staying on the equilibrium path to be greater than the payoff from that deviation: $\pi_H + c_H(e_H) \geq \pi_L - c_H(0)$. We'll assume from here on out that $c_i(0) = 0$ for $i \in \{L, H\}$. Then we see that we need $e_H \leq c_H^{-1}(\pi_H - \pi_L) \equiv \bar{e}_H$.

Finally, we cannot have e_H too low, or L types will then prefer to deviate to e_H because the cost of education is small enough relative to the higher wage. Thus we need the payoff from staying on the equilibrium path to be greater than the payoff from that deviation: $\pi_L \geq \pi_H - c_L(e_H)$. This gives $e_H \geq c_L^{-1}(\pi_H - \pi_L) \equiv \underline{e}_H$.

For any $e_H \in [\underline{e}_H, \bar{e}_H]$, there exist beliefs that support such a separating WPBE.

(c) Now we are considering pooling equilibria, where both types pick education level e^P and wages are w^P. Note that by Bayes' rule, we must have the firm's belief that the student is of type H when observing e^P be $\mu(e^P) = \lambda$, where λ is the population distribution of H types. Then we have $w^P = w(e^P) = \lambda\pi_H + (1-\lambda)\pi_L \equiv \pi^E$.

We can support a pooling equilibrium with wages

$$w(e) = \begin{cases} \pi_L & e < e^P \\ \pi^E & e \geq e^P. \end{cases}$$

(Note that wages and beliefs are 1-to-1 by the Bertrand assumption.) However, we need to ensure that under this wage schedule, low types will not want to deviate to zero education. That is, we need $\pi^E - c_L(e^P) \geq \pi_L$, or $e^P \leq c_L^{-1}(\pi^E - \pi_L) = c_L^{-1}(\lambda\pi_H - \lambda\pi_L) \equiv e^*$. If $e^P > e^*$, there are no beliefs that will support this pooling equilibrium.

(d) We can use the equilibrium dominance criterion (also known as the intuitive criterion) to eliminate all but the most efficient separating equilibrium. For a proof, see Doug's notes, section 7.3.

(e) (i) For the pooling equilibrium: Under the taxation scheme, our argument that $e_L = 0$ still holds. The intuition for the bounds on e_H is also the same, though we need to consider that type i's payoffs are now $w - c_i(e) - te$. This gives us $c_H(\bar{e}_H) + t\bar{e}_H = \pi_H - \pi_L$ and $c_L(\underline{e}_H) + t\underline{e}_H = \pi_H - \pi_L$.

For the pooling equilibrium: Wages are the same as above. Note that the intuition for e^* is the same, but again we need to consider the new form for the utility. This allows to find e^* by $c_L(e^*) + te^* = \pi^E - \pi_L = \lambda(\pi_H - \pi_L)$.

(ii) In the most efficient separating equilibrium, we have $e_H = \underline{e}_H$, where $\underline{e}_H$ is given by $c_L(\underline{e}_H) + t\underline{e}_H =$

$\pi_H - \pi_L$. We can implicitly differentiate this expression with respect to t:

$$\begin{aligned} c_L(\underline{e}_H) + t\underline{e}_H &= \pi_H - \pi_L \\ c_L'(\underline{e}_H)\frac{d\underline{e}_H}{dt} + t\frac{d\underline{e}_H}{dt} + \underline{e}_H &= 0 \\ \frac{d\underline{e}_H}{dt} &= \frac{-\underline{e}_H}{c_L'(\underline{e}_H) + t} \end{aligned}$$

Why do we need this expression? Because we need to calculate the derivative of $u_H = \pi_H - c_H(\underline{e}_H) - t\underline{e}_H$ with respect to t:

$$\begin{aligned} u_H &= \pi_H - c_H(\underline{e}_H) - t\underline{e}_H \\ \frac{du_H}{dt} &= -c_H'(\underline{e}_H)\frac{d\underline{e}_H}{dt} + t\frac{d\underline{e}_H}{dt} + \underline{e}_H \\ &= \frac{c_H'(\underline{e}_H) - t}{c_L'(\underline{e}_H) + t}\underline{e}_H + \underline{e}_H \\ &= \frac{c_H'(\underline{e}_H) - c_L'(\underline{e}_H)}{c_L'(\underline{e}_H) + t}\underline{e}_H < 0 \end{aligned}$$

And since $u_L = \pi_L$, we have that $\frac{du_L}{dt} = 0$. Thus the tax has an overall negative effect on the welfare of the workers.

(f) (i) Note that since the lump sum is the same for all workers, and there are so many workers that any individual cannot impact the size of government revenues, it must be that all of our equilibrium conditions do not change. This is because the lump sum amount will show up on either side of the equations that we use to derive the various constraints in the previous parts.

(ii) The lump sum simply shifts everyone's utility up by R. If we let $\tilde{u}_i$ be type i's utility under the transfer case, we have that $\tilde{u}_i = u_i + R$. Then we have $\frac{d\tilde{u}_L}{dt} = \frac{dR}{dt}$ and $\frac{d\tilde{u}_H}{dt} = \frac{du_H}{dt} + \frac{dR}{dt}$. We can calculate the derivative easily, since only the high types get any education, and their level of education is $\underline{e}_H$:

$$\begin{aligned} R &= \lambda t\underline{e}_H \\ \frac{dR}{dt} &= \lambda\underline{e}_H + \lambda t\frac{d\underline{e}_H}{dt} \\ &= \lambda\left[\underline{e}_H - \frac{t\underline{e}_H}{c_L'(\underline{e}_H) + t}\right] \\ &= \lambda\underline{e}_H\left[\frac{c_L'(\underline{e}_H)}{c_L'(\underline{e}_H) + t}\right] > 0 \end{aligned}$$

Thus $\frac{d\tilde{u}_L}{dt} > 0$. Note also that

$$\begin{aligned} \frac{d\tilde{u}_R}{dt} &= \frac{c_H'(\underline{e}_H) - c_L'(\underline{e}_H)}{c_L'(\underline{e}_H) + t}\underline{e}_H + \lambda\underline{e}_H\left[\frac{c_L'(\underline{e}_H)}{c_L'(\underline{e}_H) + t}\right] \\ &= \underline{e}_H\left[\frac{c_H'(\underline{e}_H) - (1-\lambda)c_L'(\underline{e}_H)}{c_L'(\underline{e}_H) + t}\right] \end{aligned}$$

Thus for λ sufficiently close to 1, we will have this expression be positive, and the net welfare effect on the workers of the tax will be positive.

Solution 2.

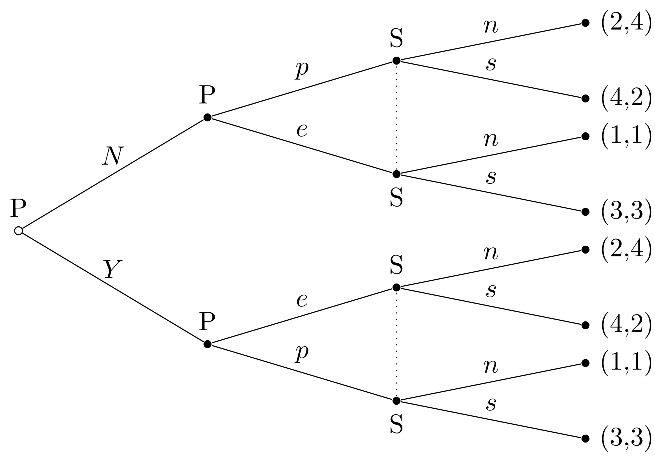

Figure 1: The extensive form for the game in 2(a).

(a) See Figure 1.

(b) The professor has 3 info sets: the root node, after he says Y, and after he says N. At each info set he chooses from among two actions, and so he will have 8 strategies. His strategy set is $S_P = \{Y, N\} \times \{e, p\} \times \{e, p\}$.

The student has just two information sets: After the professor says Y and after the professors says N. At each info set she has two possible actions, for a total of four strategies. We have $S_S = \{s, n\} \times \{s, n\}$.

(c) See Figure 1.

(d) I think the easiest and most reliable way to answer this question is to write down the normal form. It is large, but relatively fast to fill out since there is must repetition. From the normal form, you can immediately see that we have the following pure-strategy NE: (Ype, nn), (Ypp, nn), (Nep, nn), and (Npp, nn).

(e) Note that the game has two proper subgames: after Y and after N. These two subgames are identical, and the unique NE of each is (p, n). Thus the professor cannot play e at any point in a SPNE.

(f) From our normal form in (d), we can see that there are several possible ways to do this. One such MSNE: the professor strictly mixes between Ypp and Npp, and the student plays nn. Note that the student has no incentive to deviate, as she is getting her best payoff, 4. The professor also will not deviate, as he cannot do better than a payoff of 2 if the student is playing nn. Note that this is also a SPNE.

(g) First, note that we only need to consider strategies that are SPNE. The only SPNE are where the professor mixes between Ypp and Npp and the student plays nn with certainty. Thus, it remains to find beliefs that make our SPNE sequentially rational and that follow Bayes' rule.

First, some notation. Let μ_Y be the student's belief that the professor has decided to play e after announcing Y. Similarly, let μ_N be her belief that the professor has decided to play e after announcing N.

Note that the professor is clearly sequentially rational. Given the student's actions, the professors is indifferent between Ypp and Npp. For the student to be sequentially rational, she must be best-responding to her beliefs. So, for nn to be best response, we need $1\mu_Y + 4(1 - \mu_Y) \geq 3\mu_Y + 2(1 - \mu_Y)$, or $\mu_Y \leq \frac{1}{2}$, as well as $\mu_N \leq \frac{1}{2}$.

Next, consider Bayes' rule. If Y is played with positive probability, then the student's info set after Y is on the equilibrium path, and we can thus calculate that $\mu_Y = 0$. Similarly, if N is played with some probability, we must have $\mu_N = 0$.

These conditions on beliefs support $(\alpha Ypp + (1-\alpha)Npp, nn)$ as a *weak* perfect Bayesian equilibrium for any $\alpha \in [0,1]$. For a PBE, we also need a WPBE in every subgame. In this case there are two, as we noted earlier, and they are identical. In the subgames, the specified strategies are (p, n), so beliefs must be $\mu_Y = \mu_N = 0$ from Bayes' rule. (Note that for the strategies restricted to these subgames, the student's information set is always reached.)

So in summary, we have the following set of PBE: $\{(\alpha Ypp + (1-\alpha)Npp, nn), \mu_Y = \mu_N = 0 | \alpha \in [0,1]\}$.

(h) We know that the set of SE will be contained in the set of PBE. We need beliefs to be consistent, so consider a sequence δ^n of strictly mixed strategies. In particular, let δ^n_Y be the probability that the professor plays Y, and similar for δ^n_{eY} and δ^n_{pY}. Note also that we must have $\delta^n_{eY} \to 0$ and $\delta^n_{eN} \to 0$. From Bayes' rule we get

$$\mu_Y^N = \frac{\delta^n_Y \delta^n_{eY}}{\delta^n_Y \delta^n_{eY} + \delta^n_Y (1-\delta^n_{eY})} = \delta^n_{eY} \to 0.$$

Thus the only possible consistent beliefs are $\mu_Y = 0$. A similar procedure shows that we must have $\mu_N = 0$ as well.

Thus we have the following set of SE: $\{(\alpha Ypp + (1-\alpha)Npp, nn), \mu_Y = \mu_N = 0 | \alpha \in [0,1]\}$. This is the same as the set of PBE above.

(i) See Figure 2.

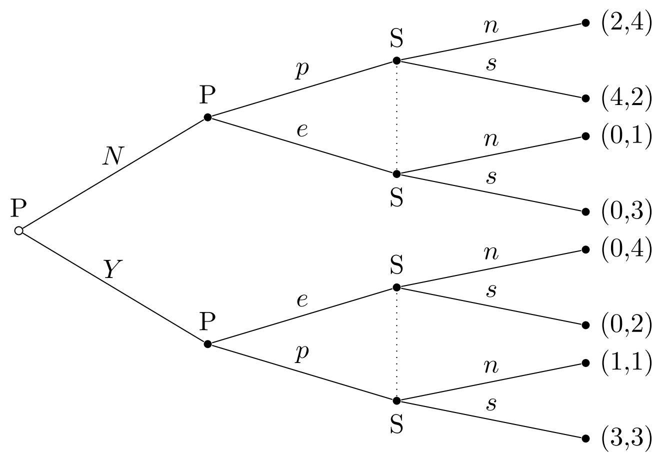

Figure 2: The extensive form for the game in 2(i).

(j) Note that in the Y subgame, the unique NE is (e, s), while in N subgame, the unique NE is (p, s). Given these continuation equilibria, the professor would rather choose Y at the beginning. Thus the unique SPNE is (Yep, sn). To find the possible PBE, we simply need to see which beliefs are allowed for this strategy profile.

First, Bayes' rule gives us that $\mu_Y = 1$. Next, note that the professor is sequentially rational, since Y then e gives him the highest payoffs given that the student is playing sn. The student is sequentially rational for playing s after Y given her beliefs that $\mu_Y = 1$. After N, we need her beliefs to be such that p is a best response: $3\mu_N + 2(1-\mu_N) \leq 1\mu_N + 4(1-\mu_N)$, or $\mu_N \leq \frac{1}{2}$.

However, as before we need to further check that the specific beliefs and strategies imply a WPBE in each subgame. This immediately requires $\mu_N = 0$ by Bayes' rule. Thus we have a unique PBE: $((Yep, sn), (\mu_Y = 1, \mu_N = 0))$.

Solution 3.

(a) Player 1's strategy set is $S_1 = \{T, NT\}$ where T stands for Tell and NT for Not Tell. Player 1 will play T in a SPNE, and this is also his dominant strategy.

(b) The strategy sets for players 1 and 2 are $S_1 = S_2 = \{T, NT\}$. By backwards induction, we see that the unique SPNE is (NT, T). We can find the NE by writing out the normal form:

	NT	T
NT	$(1,1)$	$(1,1)$
T	$(2,2)$	$(0,3)$

We can see that there is one PSNE: (NT, T). There is also a family of MSNE, where 1 plays NT and 2 mixes between NT and T with probability less that $\frac{1}{2}$ on NT. However, for the purposes of this question we will consider only pure strategies. (I believe this was the intent of the question, based on the past solutions.)

(c) By backwards induction, we see that the unique SPNE is (T, NT, T). To find the NE, we can look at the three-player normal form:

	NT	T
NT	$(1,1,-2)$	$(1,1,-2)$
T	$(2,2,-1)$	$(0,3,0)$

NT

	NT	T
NT	$(1,1,-2)$	$(1,1,-2)$
T	$(2,2,-1)$	$(0,3,1)$

T

The PSNE are (NT, T, NT), (T, NT, T), and (NT, T, NT).

(d) We can solve the general game by backwards induction. Note that player $n-1$ always tells, since he gains an extra informed friend by doing this. Knowing this, player $n-2$ chooses to not tell, since by telling he gains an informed friend but also and informed enemy. Given this, player $n-3$ faces the same situation as player $n-1$, and so we have the players continuing to alternate between T and NT. Thus we have the following SPNE: If n is even, all even number players play NT and all odd number players play T. If n is odd, all even number players play T and all odd number players play NT.

(e) Note that player $n-1$ will tell iff $\alpha \geq 0$; that is, if the value of an informed friend is non-negative. Assuming $\alpha \geq 0$, player $n-2$ will tell iff $\alpha \geq \beta$; that is, iff the value of an extra informed friend outweighs the value of an extra informed enemy. Assuming $\alpha \geq 0$ and $\alpha \geq \beta$, player $n-3$ will tell iff $\alpha \geq 2\beta$. We can continue this process until $\alpha \geq k\beta$ is not true for some k, then then we start over. The intuition here is that assuming all his successors will pass, a given player will pass only if the value of one informed friend is greater than the cost of many informed enemies.

More formally, let $\bar{k}$ be the integer such that $\alpha < \bar{k}\beta$ but $\alpha \geq (\bar{k}-1)\beta$. Then the following is a SPNE: $(\ldots, NT, \underbrace{T, \ldots TT}_{\bar{k} \text{ times}}, NT, \underbrace{T, \ldots TT}_{\bar{k} \text{ times}})$.

(f) (i) Note that the only (pure-strategy) NE of the stage game is (NT, T). If we have finite repetitions of the stage game, backwards induction will give us a unique (pure-strategy) SPNE, wherein each player plays his stage NE strategy in each repetition. Thus in this case there is no SPNE where player 1 tells player 2 the secret.

(ii) The lack of verification means that 1 has imperfect information in the game: He cannot tell in any stage whether 2 played NT or T. Thus a cooperative SPNE is not possible, since we would need player 1 to be able to punish player 2 for deviating. Such punishment is not possible if 1 cannot tell whether 2 deviated!

(iii) Now we can sustain a cooperative outcome with Nash reversion. Consider the following strategy profiles:

- Player 1 starts by playing T. If he ever observes 2 play T, he plays NT forever. Otherwise, he plays T.
- Player 2 starts by playing NT. If player 1 ever observes him play T, or if 1 plays NT at any point, then 2 plays T forever. Otherwise, he continues playing NT.

I propose that these constitute a NE for some discount factor δ. First, note that off the equilibrium path, they will play (NT, T). Clearly neither wishes to deviate, since this is the NE of the stage game.

The equilibrium path has the players playing (T, NT), so it remains to check whether either has a profitable one-shot deviation from this path:

- Player 1 is getting his best possible payoff (2) in every period, so he will clearly not want to deviate.
- On the path, player 2 is also getting a payoff of 2. If he deviates to playing T, he will get payoff 3 in that period. Player 1 will be informed of this with probability p, in which case player 2's payoff will be 1 for every stage thereafter. Player 1 will NOT be informed with probability $(1-p)$, in which case player 2 will go back to playing NT in subsequent periods. (Remember, we only need to check one-time deviations.) Thus 2 will not deviate if

$$\frac{1}{1-\delta}2 \geq 3 + p\frac{1}{1-\delta}1 + (1-p)\frac{\delta}{1-\delta}2,$$

or $\delta \geq \frac{1}{1+p}$.