14.126 Game Theory
Final Exam

You have 24 hours from the time you pick up the exam (i.e. you need to return your solutions by the same time next day). You can use any existing written source, but you cannot discuss the content of this exam with others. Questions are equally weighted. Good Luck.

1. This question tries to evaluate your understanding of the basic solution concepts.

   (a) (10 points) Consider a class of games $G_\theta = (N, S, u(\cdot;\theta))$, indexed by a payoff parameter θ which comes from a compact metric space. Assume that player set N and strategy space S are finite. State broad conditions on u under which whenever G_{θ^*} has a unique Nash equilibrium s^* in pure strategies, θ^* has a neighborhood η such that for each $\theta \in \eta$, G_θ has a unique Nash equilibrium s^θ in pure strategies and $s^\theta = s^*$.

   Answer: Assume that u is continuous in θ. Since S is finite, u is also continuous in s. Then, as proved in the lecture slides, the correspondence PNE of Nash equilibria in pure strategies is upper-hemicontinuous and non-empty. Since S is finite, this means that there exists a neighborhood η of θ^* such that for each $\theta \in \eta$, $PNE(\theta) \subseteq PNE(\theta^*)$. Since $PNE(\theta) \neq \varnothing$ and $PNE(\theta^*) = \{s^*\}$, this implies that $PNE(\theta) = \{s^*\}$ for each $\theta \in \eta$.

   (b) (15 points) Compute the sequential equilibria of the following Bayesian game. There are two players and a payoff parameter $\theta \in \{0, 1\}$. Player 1 knows the value of θ, and player 2 knows the value of a signal $x \in \{0, 1\}$. The joint distribution of θ and x is given by

   $$\Pr((\theta, x) = (0, 0)) = \Pr((\theta, x) = (1, 0)) = \Pr((\theta, x) = (1, 1)) = 1/3.$$

   The payoffs and actions are as in the following "tree":

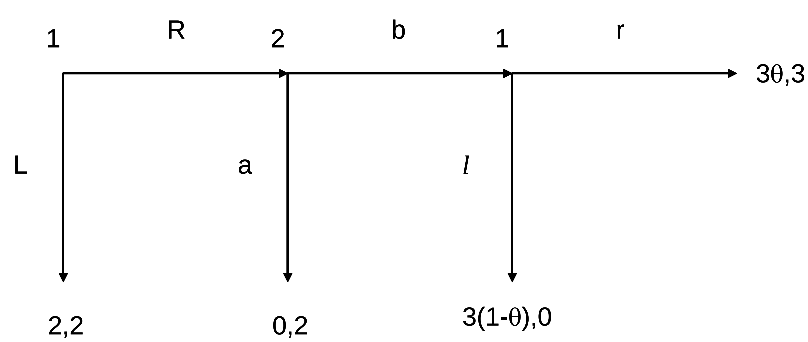

   Answer: The incomplete information game is as follows, where I also indicated the necessary moves in any sequential equilibrium, the moves that are obvious:

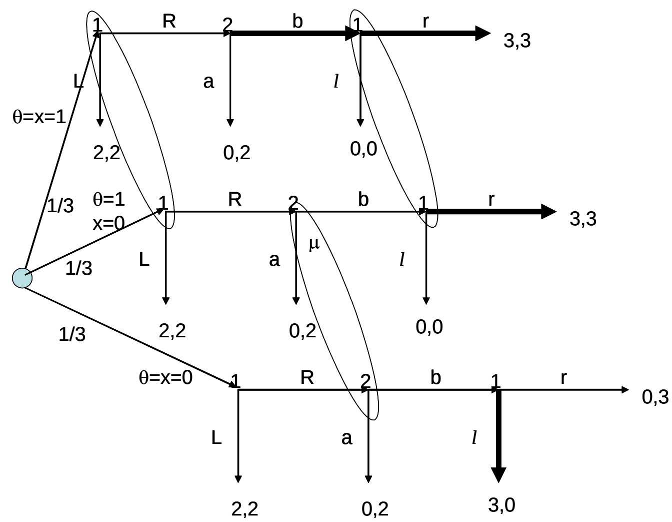

Write $\mu = \Pr(\theta = 1 | x = 0, R)$ for the probability player 2 assigns on her nontrivial information set. Write β for the probability that she plays b at that information set. Write $p = \Pr(R|\theta = 1)$ for the probability player 1 plays R on the upper branches. Similarly, write $q = \Pr(R|\theta = 0)$. Notice that player 1 assigns 1/2 to $x = 1$ on his first non-trivial information set ($\theta = 1$; empty history).

Now, at that information set, if player 1 plays R, he gets 3 with probability $1/2 + (1/2)\beta = (1+\beta)/2$ and 0 with the remaining probability. He gets 2 from playing L, so by sequential rationality

$$p = \begin{cases} 1 & \text{if } \beta > 1/3 \\ [0,1] & \text{if } \beta = 1/3 \\ 0 & \text{if } \beta < 1/3. \end{cases}$$

When $\theta = 0$, he gets 3 with probability β if he plays R and 2 for sure if he plays L. Hence, by sequential rationality

$$q = \begin{cases} 1 & \text{if } \beta > 2/3 \\ [0,1] & \text{if } \beta = 2/3 \\ 0 & \text{if } \beta < 2/3. \end{cases}$$

Finally sequential rationality for player 2 implies that

$$\beta = \begin{cases} 1 & \text{if } \mu > 2/3 \\ [0,1] & \text{if } \mu = 2/3 \\ 0 & \text{if } \mu < 2/3. \end{cases}$$

Towards computing sequential equilibria, let's rule out certain values for β. Firstly, if $\beta > 2/3$, then $p = q = 1$, and hence by consistency $\mu = 1/2 < 2/3$, showing that $\beta = 0$, a contradiction. Hence, $\beta \leq 2/3$. On the other hand, if $\beta \in (1/3, 2/3)$, then $p = 1$ and $q = 0$, so that $\mu = 0$ (by consistency), yielding $\beta = 0$, another contradiction. Hence, $\beta \in [0, 1/3] \cup \{2/3\}$. These values indeed correspond to the two components of sequential equilibria.

First component: $\beta = 2/3$; $\mu = 2/3$ (by sequential rationality of 2); $p = 1$, and $q = 1/2$ (by consistency).[1]

Second Component: $p = q = 0$; $\beta \leq 2/3$; $\mu \leq 2/3$.

Given the above equations, it is easy to verify that each combination in the above components yields a sequential equilibrium. There is a continuum of sequential equilibria, but there are only two sequential equilibrium outcomes.

2. Consider an infinitely repeated game with discount factor δ and the stage game

	C	D
C	$5, 5$	$0, 5+\theta$
D	$5+\theta, 0$	θ, θ

where $\theta \in \left[-\bar{\theta}, \bar{\theta}\right]$ for some $\bar{\theta} \gg 5$. All previous actions are common knowledge.

(a) (7 points) Assuming that θ is common knowledge and fixed throughout the repeated game, compute the best "strongly symmetric" subgame perfect equilibrium. (A strategy profile is strongly symmetric if the players' actions are identical for every given history.)

Answer: When $\theta \leq 0$, the best equilibrium is clearly, each player playing their dominant dominant C at all histories. Similarly, when $\theta \geq 5$, always playing the dominant action D is the best SPE (the best possible outcome). Now consider the case $\theta \in (0, 5)$. As in Abreu's theorem, we can take the outcome path after a deviation as the worst SPE, which is playing the static equilibrium (D, D) forever. On the path they play (C, C). In order this to be an equilibrium it must be that $5 \geq (5+\theta)(1-\delta) + \delta\theta$, i.e.,

$$5\delta \geq \theta.$$

Hence, when $\theta \leq 5\delta$, the the best strongly symmetric SPE is play (C, C) until someone deviates, and play (D, D) thereafter. It is always play (D, D) when $\theta > 5\delta$.

(b) (7 points) Assume that every day a new θ is drawn, and $(\theta_0, \theta_1, \ldots)$ are iid with uniform distribution on $\left[-\bar{\theta}, \bar{\theta}\right]$, where θ_t is the θ at t. Assume also that θ_t becomes common knowledge at t. Consider the following Grim Trigger strategy: on the path play C at t iff $\theta_t \leq \hat{\theta}$, and if any player deviates play C iff $\theta_t \leq 0$ thereafter. [The cutoff is $\hat{\theta}$ on the path and 0 off the path.] Find the conditions (on $\hat{\theta}$ and δ) under which each player playing the Grim Trigger strategy is a subgame-perfect equilibrium.

Answer: The value of being on the path is

$$f\left(\hat{\theta}\right) = \frac{1-\hat{\theta}}{2\bar{\theta}} \cdot \frac{\hat{\theta}+\bar{\theta}}{2} + \frac{\hat{\theta}}{2\bar{\theta}} \cdot 5.$$

The value of deviation is $f(0)$. For this to be an equiliubrium, for each $\theta \leq \hat{\theta}$, we must have

$$5(1-\delta) + \delta f\left(\hat{\theta}\right) \geq (5+\theta)(1-\delta) + \delta f(0)$$

[1] $\mu = p/(p+q)$, i.e., $2/3 = 1/(1+q)$.

i.e.,

$$\theta \le \frac{\delta}{1-\delta}\left[f\left(\hat{\theta}\right) - f(0)\right]. \tag{1}$$

Therefore, we must have

$$\hat{\theta} \le \frac{\delta}{1-\delta}\left[f\left(\hat{\theta}\right) - f(0)\right].$$

In order for playing D when $\theta > \hat{\theta}$, it suffices that $\hat{\theta} \ge 0$. For the best equilibrium, we consider the largest $\hat{\theta}$ with equality. For convenience, let's rewrite the equation as

$$g\left(\hat{\theta}\right) \equiv \frac{f\left(\hat{\theta}\right) - f(0)}{\hat{\theta}} = \frac{1-\delta}{\delta}. \tag{2}$$

Since f is concave, g is decreasing in $\hat{\theta}$, and hence there is a unique solution to this equation, which you can compute algebraically. Using (1) and the single deviation principle, one can easily show that this condition is also sufficient.

(c) (7 points) In part (b), assume instead that the players do not know θ_t, but each player i observes a signal $x_{i,t} = \theta_t + \varepsilon\eta_{i,t}$ at date t, where $\varepsilon > 0$ is small, the random variables $\left(\theta_t, \eta_{1,t}, \eta_{2,t}\right)$ are all stochastically independent, and each $\eta_{i,t}$ is uniformly distributed on $[-1, 1]$. Consider the following Grim Trigger strategy:

At the initial state, play C at t iff $x_{i,t} \le \hat{x}$. If any player plays D, switch to the punishment state, and remain there forever. At the punishment state, play C iff $x_{i,t} \le 0$.

Find the conditions under which each player playing the Grim Trigger strategy is a perfect Bayesian Nash equilibrium in the limit $\varepsilon \to 0$.

Answer: (This question was inspired by Sylvain Chassang's dissertation.) In the limit $\varepsilon \to 0$, the value of being at the initial state is $F(\hat{x})$ and the value of being in the punishment state is $f(0)$ (as in part (b)). The value $F(\hat{x})$ is different from $f(\hat{x})$ because of possibility of punisment on the path. To compute F, letting $\pi = \left(\hat{x} + \bar{\theta}\right) / \left(2\bar{\theta}\right)$ to be the probability of remaining on the path, we write

$$F(\hat{x}) = \pi \cdot [5(1-\delta) + \delta F(\hat{x})] + (1-\pi)\left[\left(\hat{x} + \bar{\theta}\right) / \left(2\bar{\theta}\right) + \delta f(0)\right]$$

and obtain

$$\begin{aligned} F(\hat{x}) &= \frac{\pi 5(1-\delta) + (1-\pi)\left(\hat{x} + \bar{\theta}\right) / \left(2\bar{\theta}\right) + (1-\pi)\delta f(0)}{1 - \delta\pi} \\ &= \frac{f(\hat{x}) + (1-\pi)\delta f(0)}{1 - \delta\pi}. \end{aligned} \tag{3}$$

Hence, the augmented game in the limit is

	C	D
C	$5(1-\delta) + \delta F(\hat{x}), 5(1-\delta) + \delta F(\hat{x})$	$\delta f(0), (1-\delta) x_2 + \delta f(0)$
D	$(5 + x_i)(1-\delta) + \delta f(0), \delta f(0)$	$(1-\delta) x_i + \delta f(0), (1-\delta) x_j + \delta f(0)$

.

Notice that probability that the other player plays C is $\Pr(x_j \leq \hat{x}|x_i)$. Hence, one must play C (in the limit) iff

$$\delta\left[F(\hat{x}) - f(0)\right]\Pr(x_j \leq \hat{x}|x_i) \geq (1-\delta)\,x_i. \tag{4}$$

By continuity (for $\varepsilon > 0$), we must have

$$\delta\left[F(\hat{x}) - f(0)\right]\Pr(x_j \leq \hat{x}|x_i = \hat{x}) = (1-\delta)\,\hat{x},$$

i.e.,

$$\delta\left[F(\hat{x}) - f(0)\right]/2 = (1-\delta)\,\hat{x},$$

yielding

$$\frac{F(\hat{x}) - f(0)}{\hat{x}} = 2 \cdot \frac{1-\delta}{\delta}. \tag{5}$$

Moreover, since $\Pr(x_j \leq \hat{x}|x_i)$ is decreasing in x_i, one can check from (4) that we have a Nash equilibrium of the augmented game, and the condition is indeed sufficient.

(d) (4 points) Comparing the cutoffs, briefly discuss your findings.

Answer: Players have the option value of getting 5 on the equilibrium path. Hence, uncertainty about the future values of θ makes the equilibrium path more attractive, making enforcement easier, and allowing more cooperation in part (b). On the other hand, even the slightest asymmetric information leads to substantial "fear of miscoordination" in equilibrium. This is reflected by the fact that the right-hand side of (5) doubled, reflecting the fact that at the cutoff there is a probability 1/2 that there will be miscoordination. Moreover, the initial state becomes less attractive because there will need to be punishment on the path. This also makes enforcement of more difficult, leading less efficient outcomes. Indeed, from (5), we can write

$$F(\hat{x}) - f(0) = \frac{(1-\delta)\left(f(\hat{x}) - f(0)\right)}{1-\delta\pi},$$

and the condition (5) becomes

$$g(\hat{x}) = 2\frac{1-\delta\pi}{\delta}.$$

Notice that the right hand side increased in two ways: $1-\delta$ increased to $1-\delta\pi$ because of possibility of punishment and it is doubled because of misccordination fear. Overall, since g is decreasing the cutoff become much smaller: $\hat{x} \ll \hat{\theta}$. There is a substantial range of parameters for which cooperation becomes unfeasible as we introduce slight incomplete information.

3. This question is about monotone transformations of supermodular games.

(a) (5 points) For a lattice $(S, \geq)$, give an example of functions $u : S \to \mathbb{R}$ and $f : \mathbb{R} \to \mathbb{R}$ such that u is supermodular, f is strictly increasing, but $f \circ u$ is not supermodular.

Answer: Take $S = (0,1)^2$, $u(x,y) = (x+y)^2$, a supermodular function, and $f = \log$, so that $f(u(x,y)) = 2\log(x+y)$, a submodular function.

(b) (12 points) Consider a supermodular game $G = (N, S, u)$ (with order $\geq$) and strictly increasing function $f : \mathbb{R}^N \to \mathbb{R}^N$. Show that the game $G' = (N, S, f \circ u)$ has Nash equilibria s^* and s^{**} such that $s^* \geq s \geq s^{**}$ for every rationalizable strategy profile s in game G'. What is the relation between s^* and s^{**} and the extremal equilibria of G.

Answer: Let $\bar{s}$ and $\underline{s}$ be the extremal equilibria of G. Then, $\bar{s}$ and $\underline{s}$ are Nash equilibria of G' and $\bar{s} \geq s \geq \underline{s}$ for every rationalizable strategy s in G'. That is, $s^* = \bar{s}$ and $s^{**} = \underline{s}$. Let me first show that $\bar{s}$ is a Nash equilibrium of G'. For any $i \in N$ and $s_i \in S_i$, since $\bar{s}$ is a Nash equilibrium of G,

$$u_i(s_i, \bar{s}_{-i}) \leq u_i(\bar{s}_i, \bar{s}_{-i}),$$

and since f is increasing,

$$f_i(u_i(s_i, \bar{s}_{-i})) \leq f_i(u_i(\bar{s}_i, \bar{s}_{-i})),$$

showing that $\bar{s}$ is a Nash equilibrium of G'. Similarly, $\underline{s}$ is a Nash equilibrium of G'. I will now show that $\bar{s} \geq s \geq \underline{s}$ for every rationalizable strategy s in G', following Milrom and Roberts. For this it suffices to generalize their lemma for G to G': if $x_i \not\geq b_i(\underline{x})$, then x_i is strictly dominated by $x_i \vee b_i(\underline{x})$ in G, i.e.,

$$u_i(x_i \vee b_i(\underline{x}), s_{-i}) > u_i(x_i, s_{-i}) \qquad \forall s_{-i} \in S_{-i}, \tag{6}$$

where $\underline{x}$ is the smallest element of S and b_i is the smallest best reply. But since f_i is increasing, (6) immediately implies that

$$f_i(u_i(x_i \vee b_i(\underline{x}), s_{-i})) > f_i(u_i(x_i, s_{-i})) \qquad \forall s_{-i} \in S_{-i}, \tag{7}$$

i.e., x_i is strictly dominated by $x_i \vee b_i(\underline{x})$ in G'. That is $b(\underline{x})$ is the smallest strategy profile among the profiles that survives the first round of elimination of doiminated strategies in G'. Applying (7) inductively as in Milgrom and Roberts, one concludes that if $s_i \not\geq \underline{s}_i$, then s_i is eliminated at some round on iterated dominance for G'. (Recall that $\underline{s} = \sup_k b^k(\underline{x})$.)

(c) (8 points) Give an example of supermodular Bayesian game (N, Θ, T, A, u, p), where $(A, \geq)$ is a complete lattice and $u(a, \theta, t)$ is continuous and supermodular in a for each θ and t, and a strictly increasing function $f : \mathbb{R}^N \to \mathbb{R}^N$ such that the largest rationalizable strategy profile in game $G' = (N, \Theta, T, A, f \circ u, p)$ is not a Bayesian Nash equilibrium of game G'. [Hint: You can take $|T| = 1$, so that there is only payoff uncertainty.]

Answer: Many of you found the following solution, which is simpler than my original solution. Take $|T| = 1$, so that there is no asymmetric information, and suppress the trivial type profile in the notation. Take $\theta \in \{0, 1\}$ with $\Pr(0) = 1/2$. Take the payoff matrix and the actions as follows:

$\theta = 1$	a	b
a	α, α	$0, 0$
b	$0, 0$	$-1, -1$

$\theta = 0$	a	b
a	$-1, -1$	$0, 0$
b	$0, 0$	α, α

Take $\alpha > 1$ so that the game is supermodular with the order $a > b$. Take $f(\alpha) \in (0,1)$, $f(0) = 0$ and $f(-1) = -1$, a monotone transformation. Then, the normal form game is

	a	b
a	$(f(\alpha)-1)/2, (f(\alpha)-1)/2$	$0, 0$
b	$0, 0$	$(f(\alpha)-1)/2, (f(\alpha)-1)/2$

where (a, a) and (b, b) are not Nash equilibria.

4. Consider the following variation on Morris-Shin currency attack model. There are two speculators and a government. The competitive exchange rate is θ, but the exchange rate is pegged at $e^* > 1$. The fundamental, θ, is uniformly distributed on $[0, L]$ where $L > e^*$ and is privately known by the government. Each speculator observes a signal $x_i = \theta + \eta_i$ where η_i is uniformly distributed on $[-\varepsilon, \varepsilon]$ where ε is very small but positive and θ, η_1, and η_2 are independent. The value of the peg for the government is v, which is privately known by the government and uniformly distributed on $[\bar{v} - 1, \bar{v}]$ for some $\bar{v} \in (0,1)$, close to 1. Each speculator i shortsells $a_i \in [0,1]$ units of the currency, and observing a_1 and a_2 the government decides whether to defend the currency. The value of the peg for the government is v, and buying each unit of the currency costs c, where $c \in (1/(2e^*), \bar{v}/2)$, so that the payoff of the government is $v - c \cdot (a_1 + a_2)$ if it defends and 0 otherwise. The payoff of speculator i is $(e^* - \theta) a_i - a_i^\gamma/\gamma - t a_i$ if the government does not defend the currency and $-a_i^\gamma/\gamma - t a_i$ if it defends, where $\gamma \geq 2$ and $0 < t < (1 - \bar{v}) e^*$.

   (a) (10 points) Show that there exists a pure strategy s^* such that all perfect Bayesian Nash equilibrium strategies converge to s^* as $\varepsilon \to 0$. (Formalize this statement.)
   Answer: By sequential rationality, the government defends the currency iff

$$v - c \cdot (a_1 + a_2) > 0;$$

   the case of indifference is irrelevant. Call this strategy s_G^*. Given s_G^*, the probability that the peg will be abandoned given a_1 and a_2 is

$$P(a_1, a_2) = \Pr(v \leq c \cdot (a_1 + a_2)) = c \cdot (a_1 + a_2) + (1 - \bar{v}).$$

   Hence, conditional on θ, a_1 and a_2, the expected payoff of speculator i is

$$u_i(a_1, a_2, \theta) = P(a_1, a_2)(e^* - \theta) a_i - t a_i - a_i^\gamma/\gamma.$$

   For a PBNE, we need to have a Bayesian Nash equilibrium in this reduced form game between the speculators. If the reduced form game satisfies the assumptions of Frankel, Morris, and Pauzner, then all rationalizable strategies and in particular all Bayesian Nash equilibria will converge to some (s_1^*, s_2^*), and all PBNE in the original game will converge to $s^* = (s_1^*, s_2^*, s_G^*)$. I now check their conditions:

   1. Dominance for $a_i = 0$: At $\theta = e^*$,

$$\partial u_i(a_1, a_2, \theta = e^*)/\partial a_i = -t - a_i^{\gamma-1} \ll 0.$$

      Since u_i is continuously differentiable there is $\bar{\theta} < e^*$ such that u_i is strictly decreasing in a_i for any $\theta \in (\bar{\theta}, e^*]$.

2. Dominance for $a_i = 1$: At $\theta = 0$, for any $a_1, a_2 \in [0,1]$,

$$\begin{aligned}\partial u_i\left(a_1, a_2, \theta = 0\right)/\partial a_i &= \left(2ca_i + ca_j + 1 - \bar{v}\right)e^* - t - a_i^{\gamma-1} \\ &> \left(2ce^* a_i - a_i^{\gamma-1}\right) + \left(1 - \bar{v} - t\right) > 0.\end{aligned}$$

Since u_i is continuously differentiable, u_i is strictly increasing in a_i on a neighborhood of $\theta = 0$.

3. Strategic complementarity:

$$\frac{\partial^2 u_i}{\partial a_i \partial a_j} = c\left(e^* - \theta\right) \geq 0.$$

4. Monotonicity:

$$\frac{\partial^2 u_i}{\partial a_i \partial \theta} = -\left(2ca_i + ca_j + 1 - \bar{v}\right) < -\left(1 - \bar{v}\right) < 0.$$

Since $\frac{\partial^2 u_i}{\partial a_i \partial \theta}$ is negative and uniformly bounded away from 0, their uniform bound is satisfied (with the reverse order on θ).

(b) (10 points) Taking $1 - \bar{v} \cong t \cong 0$ (i.e. ignoring these variables), compute s^*.

Answer: Recall that $\left(s_1^*\left(x\right), s_2^*\left(x\right)\right)$ is a symmetric Nash equilibrium of the complete information (reduced-form) game with $\theta = x$; s_G^* is already fixed. Hence, we first need to compute the symmetric Nash equilibria of the complete information game for speculators. By continuity of u_i, NE are continuous with respect to $\bar{v}$ and t; hence, we will take $1 - \bar{v} = t = 0$. For an interior equilibrium $a_1 = a_2 = a \in (0,1)$, we need

$$\frac{\partial u_i}{\partial a_i} = \left(2ca + ca\right)\left(e^* - x\right) - a^{\gamma-1} = 0,$$

i.e.,

$$a = \left(3c\left(e^* - x\right)\right)^{1/(\gamma-2)}.$$

We have $a \in (0,1)$, when $3c\left(e^* - x\right) \in (0,1)$, i.e. when $x \in \left(e^* - \frac{1}{3c}, e^*\right)$. Check that the second order condition $\partial^2 u_i / \partial a_i^2 = 2c\left(e^* - x\right) - \left(\gamma - 1\right)a^{\gamma-2} < 0$ is satisfied. The only other possible symmetric pure strategy equilibria are the corner solution. For $a_1 = a_2 = 0$ to be an equilibrium, we need

$$0 \geq \frac{\partial u_i}{\partial a_i}|_{a_j=0} = 2ca_i\left(e^* - x\right) - a_i^{\gamma-1} = a_i\left(2c\left(e^* - x\right) - a_i^{\gamma-2}\right).$$

The condition will not be satisfied around 0 whenever $x < e^*$. Hence we have this equilibrium only at $x = e^*$.[2] In order for $a_1 = a_2 = 1$ to be an equilibrium, we need

$$\frac{\partial u_i}{\partial a_i}|_{a_i=a_j=1} = \left(2ca_i + c\right)\left(e^* - x\right) - a_i^{\gamma-1} = 3c\left(e^* - x\right) - 1 \geq 0,$$

[2] There was a region where $a_1 = a_2 = 0$ is the only equilibrium, but this dominance region shrinks to a single point as $t \to 0$.

which is satisfied only when $x \leq e^* - \frac{1}{3c}$. It turns out that there is a unique symmetric pure strategy Nash equilibrium for each value of x. Since s_i^* is a selection from these equilibria, it is simply the unique solution:

$$s_i^*(x) = \begin{cases} (3c(e^* - x))^{1/(\gamma-2)} & \text{if } x > e^* - \frac{1}{3c} \\ 1 & \text{otherwise.} \end{cases}$$

(c) (5 points) Briefly describe what rationality assumptions you have made for your solution.

Notice that we only used sequential rationality for the government and then applied rationalizability (by Frankel, Pauzner, and Morris) in the reduced game. Hence, the assumptions are

1. the government is sequentially rational,
2. speculators are rational, and
3. (i) and (ii) are common knowledge among the speculators.

14.126 Game Theory
Spring 2010