Tutorial #3

Exercise 1

Compute the maximin strategies and security levels for both players for this game:

	L	R
T	6, 8	2, 0
B	1, 4	5, 8

Exercise 2

In class we saw that we can reduce a player's uncertainty about the actions available to her opponents to her uncertainty about the utility functions involved. Construct a similar example that shows that it is also possible to reduce a player's uncertainty about *how many* opponents she is playing against to her uncertainty about utility functions.

Do not attempt to construct a general mathematical argument. Instead, only provide a simple example that illustrates the phenomenon alluded to above. The purpose of this exercise is to demonstrate that the model of Bayesian games introduced in class is very general and covers all sorts of different kinds of uncertainty in games.

Exercise 3

Consider the following two-player Bayesian game. If the row player is *lucky* (which is the case with prior probability $p = \frac{1}{2}$), she has a guaranteed payoff of 1 (and the column player gets 0). But if she is *unlucky*, the two players play a variant of the Prisoner's Dilemma.

	L	R
T	1, 0	1, 0
B	1, 0	1, 0

row player is lucky

	L	R
T	3, 3	0, 5
B	5, 0	1, 1

row player is unlucky

Answer the following questions:

(a) What are the pure strategies available to the row player?

(b) What are the pure strategies available to the column player?

(c) How many pure-strategy profiles are there?

(d) Which of theses pure-strategy profiles are Bayes-Nash equilibria? Why?

Exercise 4

Recall the definition of Bayesian games given in class, which we had formulated in terms of *epistemic types*. Here is an alternative definition, formulated in terms of *information sets:*

> A *Bayesian game* is a tuple $\langle N, \boldsymbol{A}, \mathcal{G}, p, \boldsymbol{\sim} \rangle$, where $N = \{1, \ldots, n\}$ is a finite set of players; $\boldsymbol{A} = A_1 \times \cdots \times A_n$ is a finite set of action profiles; $\mathcal{G}$ is a finite set of games of the form $\langle N, \boldsymbol{A}, \boldsymbol{u} \rangle$ that only differ in $\boldsymbol{u}$; $p \in \Pi(\mathcal{G})$ is a common prior over that set of games; and $\boldsymbol{\sim} = (\sim_1, \ldots, \sim_n)$ is a vector of equivalence relations on $\mathcal{G}$, each of which is partitioning $\mathcal{G}$ into information sets for one of the players.

The idea is that p is used to draw a game G from the set of all games $\mathcal{G}$ to be played in a given situation. While every player i knows the common prior p, she is uncertain about which game ends up being played and only can tell that it is some game G' such that $G' \sim_i G$. The set of all such games G' is her information set.

Discuss these two alternative approaches to modelling incomplete information and try to understand how you would translate a given game from one formalism to the other.