GAME THEORY: ANALYSIS OF STRATEGIC THINKING
Exercises on Multistage Games with Chance Moves, Randomized Strategies and Asymmetric Information

Pierpaolo Battigalli
Bocconi University

A.Y. 2006-2007

Abstract

The number of stars denotes the difficulty of the exrercise:
(*) easy
(**) medium
(***) difficult

Exercise 1

Pat and Andy play the following *principal-agent game*:

Andy is a manager. If he works for Pat he can either put effort in a project ($e = 1$) or shirk ($e = 0$). The project can either be successful ($y = 1$) or unsuccessful ($y = 0$). The probability of success depends on effort as follows:

$p(y\|e)$	$e = 0$	$e = 1$
$y = 0$	p	$1 - p$
$y = 1$	$1 - p$	p

where $p > \frac{1}{2}$. We assume that y is verifiable but e is not. Thus, an enforceable contract can only specify a reward to Andy that depends on y, not on e.

Pat offers Andy an enforceable contract of the form (w, B) where w is a baseline wage and B is bonus to be awarded if and only $y = 1$. Andy can accept (Yes) or reject (No). If he rejects the game is over, if he accepts he then chooses the effort level.

If Andy says *No*, the players get their outside-option monetary payments, which is 0 for both. If Andy accepts payoffs are as follows:

$$\begin{aligned} u_P((w,B), Yes, e, y) &= (\Pi - B)y - w, \\ u_A((w,B), Yes, e, y) &= \sqrt{w + By} - e \end{aligned}$$

(Andy likes money, but he is risk-averse).

Note that *if* e were verifiable, Pat could offer Andy a constant wage w just high enough to make him accept a committment to exert effort. Solving the participation constraint $\sqrt{w} \geq 1$ as an equality we obtain $w = 1$. We *assume* that Pat would find this more profitable than hiring Andy with zero wage to exert zero effort:

$$p\Pi - 1 > (1-p)\Pi.$$

(1) (*) Draw a graph representing this game.

(2) (**) Find the least expensive contract inducing Andy to exert effort (assuming that Andy breaks indifferences in favor of Pat). [**Hint:** write the incentive constraint stating that Andy weakly prefers to exert effort; note that the difference $\sqrt{w+B} - \sqrt{w}$ is decreasing in the baseline wage w; use this observation to obtain $w = 0$, then derive B).

(3) (**) Derive the subgame perfect equilibrium (SPE) as a function of the parameters p and Π. Determine the region of the parameter space where the equilibrium contract induces effort.

(4) (**) Suppose that Pat and Andy play this game *infinitely* many times with discount factor δ, assuming that e is *observable, although not verifiable.* (Note that in the one-shot game the contract (w, B) is enforceable, thus Pat cannot renege on either w, or B if $y = 1$. In the solution of the one-shot game it does make any difference whether Pat is able to pay an additional discretionary bonus after she has observed e, because subgame perfection dictates that she would not pay it. However, this possibility would make a difference in the repeated game. To simplify the exercise, *we assume away the possibility of paying a discretionary bonus after the observation of* e. Any "reward" for the current effort must necessarily take place with a possibly higher wage in future periods.) Assume that p and Π are such that in the SPE of the one-shot game Pat does not offer an effort-inducing contract because it is too costly (see point 3). Show that if $\delta \geq 1/\sqrt{(2p-1)\Pi}$ there is a SPE of the repeated game in which no bonus is offered and nonetheless Andy exerts effort on the equilibrium path. [**Hint:** Consider trigger strategies reverting to the SPE of the one-shot game and note that Andy must have no incentive to shirk, and Pat must prefer the trigger-strategy payoff to the repetition of the one-shot-game SPE.)

(5) (***, optional) Under the same assumptions of point 4, show that for $\delta \geq \max\left(\frac{\sqrt{(1-p)\Pi}}{\sqrt{(1-p)\Pi}+\sqrt{p\Pi}-1}, \frac{1}{\sqrt{p\Pi}}\right)$ there is a SPE where (on the equilibrium path) Andy exerts effort and Pat offers $B = 0$ and $w = p\Pi$, thus getting zero profit in expectation. [**Hint:** The usual trigger strategies do not work. Andy may be punished with a reversion to the one-shot SPE (where he gets zero), but if Pat is punished with a reversion to the one-shot SPE she gets a positive payoff under the punishment. Thus, Andy must reject any offer that leaves Pat with a positive expected payoff in the current period, and he must by willing to do so under the assumption that if he accepts today he will get zero from tomorrow onward. This yields the first threshold. The second threshold comes from Andy's incentive to exert effort.)

Exercise 2 (*)

This exercise is taken from Osborne Rubinstein page 216.

Consider the following game with *imperfect* observation of past actions: player 1 can terminate the game (B) or give the move to 2, who can terminate the game (L) or give the move back to 1 in one of two ways (M or R). If 1 gets to move in the third stage, she does not observe whether 2 chose M or R. In the game-tree representation the latter assumption is represented by the dashed line connecting history/node (A, M) to history/node (A, R).

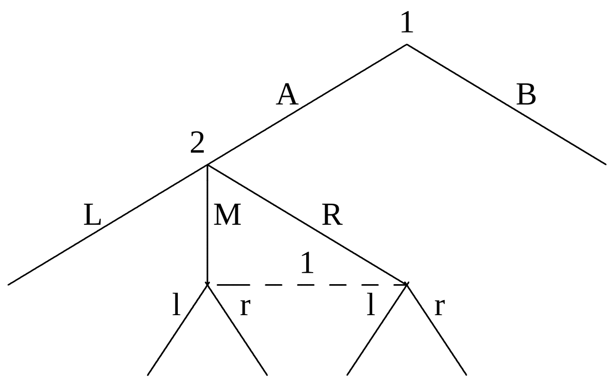

A pure strategy of 1 specifies a feasible action for each situation in which she can find herself playing: the beginning of the game, and the *information set* $I = \{(A, M), (A, R)\}$ (since 1 cannot distinguish between the two histories (A, M) and (A, R)). For example, (A, r) is the pure strategy that first selects A and prescribes to play r *if* I were reached.

A behavioral strategy of 1 is given by two probabilities: the probability of A at h^0 and the probability of r at I.

Find the behavioral strategy of player 1 that is equivalent to the mixed strategy in which she plays (B, r) with probability 0.4, (B, l) with probability 0.1 and (A, l) with probability 0.5

Exercise 3
Consider the following one-person game tree with perfect information:

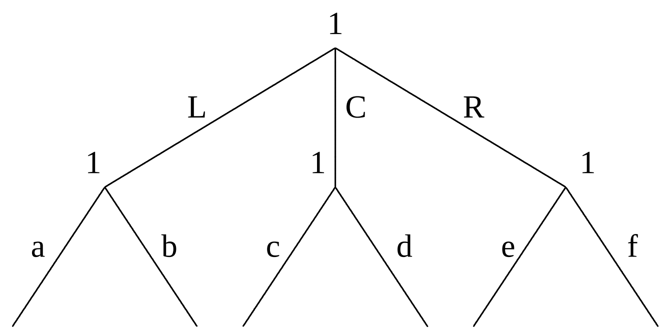

(1) (*) How many strategies does player 1 have? And how many "plans of action" (strategies of the reduced normal form)?

(2) (**) Consider the following mixed strategy:

$$\hat{\sigma}_1(s_1) = \begin{cases} \frac{2}{12} & \text{if } s_1 = (Lace) \\ \frac{1}{12} & \text{if } s_1 = (Lbcf) \\ \frac{1}{8} & \text{if } s_1 = (Cbcf) \\ \frac{1}{8} & \text{if } s_1 = (Cacf) \\ \frac{1}{4} & \text{if } s_1 = (Rbcf) \\ \frac{1}{4} & \text{if } s_1 = (Radf) \\ 0 & \text{otherwise} \end{cases}$$

Compute the distribution on the set of the terminal histories induced by $\hat{\sigma}_1$. Is $\hat{\sigma}_1$ the only mixed strategy that induces such distribution? If not, provide an example of another mixed strategy that generates the same distribution on Z.

(3) (**) Compute the behavioral strategy realization-equivalent to mixed strategy $\hat{\sigma}_1$.

(4) (**) Consider the behavioral strategy of point (3). Compute the realization-equivalent mixed strategy derived from this behavioral strategy under the assumption of "independence across agents". Is it equal to $\hat{\sigma}_1$? Explain. What is the distribution on the set of terminal histories induced by the mixed strategy you have just found? Explain.

Exercise 4

Consider the following *Minipoker game*:

There are two players (Ann and Bob), and a deck that contains three cards: a King (K), a Queen (Q) and a Jack (J). The King wins over the Queen which wins over the Jack (of course the King wins over the Jack too).

At the beginning of the game, each player puts one Euro on the pot.

The cards are shuffled and the resulting order is random. There are 6 possible orders (we write only the first two cards because the third is residually determined): KQ, KJ, QK, QJ, JK, JQ. Each order is equally likely. The first card is given to Ann, and the second is given to Bob. We denote with θ_i the card of player i.

Ann is the first player to move. She can put an additional euro on the pot (action $B =$"Bet") or leave (L). If Ann leaves, the game ends and Bob wins the pot.

If Ann bids, Bob can put a euro on the pot (action $c =$"call"), or leave the pot to Ann without running the risk to lose an additional euro (action $f =$"fold"). If Bob calls, cards are shown and the player with the higher card wins (if, for example, Bob has a Queen and Ann has a Jack , Bob wins).

This Minipoker can be described as an extensive form game with imperfect information on the initial random move (that is, the order of the cards in the deck). To answer the following question, you do not have to represent the whole game-tree, but it could be helpful.

TERMINOLOGY. In this context (games with asymmetric information regarding an initial random move), the term *history* can be interpreted in two different ways. We can either regard the random move as the first element of a history, or we can interpret it as a profile of "types" and include in histories h only the players' actions. Note that, in the first case, histories correspond to the nodes of the tree. Both terminologies are allowed and used in the literature. *Following the notation we have used for games with incomplete information, we adopt the second one.* Then, each *node* of the "arborescence" representing the game corresponds to a pair (θ, h) where $\theta = (\theta_i)_{i\in N}$ is a profile of types and h is a sequence of players' actions (or, more generally, a sequence of action profiles). An *information set* of player i is then a set of nodes defined in the following way:

$$I(\overline{\theta}_i, \overline{h}) = \{(\theta_i, \theta_{-i}, h) : \theta_i = \overline{\theta}_i, \theta_{-i} \in \Theta_{-i}, h = \overline{h}\}.$$

The game we described has a *unique Bayesian perfect equilibrium* (actually, even a unique Bayes-Nash equilibrium) in partially mixed behavioral strategies. You can determine it, answering to the following questions.

(1) (*) To each order of the cards in the deck (e.g. KQ, KJ, etc.) we can associate a sub-tree that describes the possible moves of Ann and Bob and their payoffs (payoff of player i=money received - money put by i on the pot) at terminal nodes. (Of course the different sub-trees are connected through the information sets, but disregard this aspect for the moment). Draw the sub-trees corresponding to KQ and JK with the related payoffs.

(2) (*) How many information sets does each player have? Write down one example of information set for each player.

(3) (*) Find the best responses of Bob if Ann bets, in the two cases $\theta_{Bob} = K$ (Bob has the King) and $\theta_{Bob} = J$ (Bob has the Jack).

(4) (*) *Given the answer to the previous point,* find the equilibrium actions of Ann in the two cases $\theta_{Ann} = K$ and $\theta_{Ann} = Q$ (respectively, Ann has the King and Ann has the Queen).

(5) (*) Denote with $\mu = \mu(\theta_{Ann} = K|B, \theta_{Bob} = Q)$ the probability that Bob assigns to the event "Ann has the King " given that Bob has the Queen and that Ann bets. For what values of μ is it rational for Bob to randomize between c e f when he has the Queen?

(6) (**) Denote with $\beta = \beta(a_{Ann} = B|\theta_{Ann} = J)$ the probability with which Ann bets when she has a *Jack* (this is the *probability Ann "bluffs"*). *Taking into account the answer you have given to the previous questions* (in particular given the answer to point (4)), find μ as a function of β using Bayes' rule. For which values of β do we obtain the value of μ calculated at point (5)?

(7) (**) Denote with $\gamma = \gamma(a_{Bob} = c|B, \theta_{Bob} = Q)$ the probability that Bob calls when he has the Queen. For what values of γ, is it rational for Ann to randomize between L and B when she observes a Jack?

(8) (**) By construction, the actions determined at point (2) and (3) and the mixed actions determined at point (5) and (6) (together with the belief system derived from the behavior strategy of Ann) determine an equilibrium of the game (can you understand why?). Try to show there are not other equilibria.

Exercise 5

Consider the following signaling game between an informed player (player 1) and an uninformed one (player 2). In particular, 1 can be of two types: θ' or θ'', and he chooses between left (l) and right (r). If 1 chooses right, 2 observes the move (message) of 1 and then chooses between up (u), middle (m) and down (d). If 1 chooses left, the game ends. There is common knowledge that, *a priori*, 2 assigns equal probability to the two types. The game is represented below:

$$
\begin{array}{ccccccccccc}
 & & & & & & & & & & (0,3) \\
 & & l & & \{\frac{1}{2}\} & & r & & & \nearrow & \\
(1,1) & \longleftarrow & - & - & 1 & - & - & \longrightarrow & 2 & \rightarrow & (2,1) \\
 & & & & \theta' & & & & : & \searrow & \\
 & & & & & & & & : & & (0,0) \\
 & & & & & & & & : & & \\
 & & & & & & & & : & & \\
 & & & & & & & & : & & (0,0) \\
 & & & & \theta'' & & & & : & \nearrow & \\
(1,1) & \longleftarrow & - & - & 1 & - & - & \longrightarrow & 2 & \rightarrow & (0,1) \\
 & & l & & \{\frac{1}{2}\} & & r & & & \searrow & \\
 & & & & & & & & & & (2,3)
\end{array}
$$

(1) (*) Denote with μ the probability of θ' given the message d, that is $\mu = \mu(\theta'|d)$. Find the expected payoff of each action of player 2 as a function of μ.

(2) (**) Find the set of separating equilibria.

(3) (**) Find the set of pooling equilibria. As an easier alternative, try to find at least one pooling equilibrium.

(4) (**) Show that in each equilibrium, no type of player 1 randomizes.

(5) (**) Does any of the equilibria satisfies a *forward-induction* condition, that is, is any of them consistent with the assumption that player 2 strongly believes in the rationality of player 1?

Exercise 6

Consider the following signaling game in which the payoff of the receiver depends on the parameter ν (payoffs of the sender are in **bold**):

$$
\begin{array}{ccccccccccc}
(\mathbf{1},4) & & & & & & & & & & (\mathbf{2},1) \\
 & \nwarrow u & & & l & \{\frac{1}{2}\} & r & & & u \nearrow & \\
 & & 2 & \longleftarrow & - \;\; - & 1 & - \;\; - & \longrightarrow & 2 & & \\
 & \swarrow d & : & & & \theta' & & & : & d \searrow & \\
(\mathbf{1},0) & & : & & & & & & : & & (\mathbf{0},0) \\
 & & : & & & & & & : & & \\
 & & : & & & & & & : & & \\
(\mathbf{1},0) & & : & & & & & & : & & (\mathbf{0},v) \\
 & \nwarrow u & : & & & \theta'' & & & : & u \nearrow & \\
 & & 2 & \longleftarrow & - \;\; - & 1 & - \;\; - & \longrightarrow & 2 & & \\
 & \swarrow d & & & l & \{\frac{1}{2}\} & r & & & d \searrow & \\
(\mathbf{1},1) & & & & & & & & & & (\mathbf{0},1)
\end{array}
$$

(1) (*) Suppose first $v = 2$. Determine the unique perfect Bayesian equilibrium.

(2) (**) Assume $v = 0$. Find two pure perfect Bayesian equilibria.

(3) (***) Which of these equilibria corresponds to a Bayes-Nash equilibrium (not necessarily perfect) also if $\nu = 2$? (Neither?, Only one? Both?)

(4) (**) Which of these equilibria satisfy the forward induction criterion of Dominated Messages? (Neither?, Only one? Both?)

Exercise 7 (***)

Consider the following signaling game (we assume that there is common knowledge of the probability that 2 assigns to the three types of player 1):

$$
\begin{array}{ccccccccc}
 & (1,1) & & & (1,1) & & & (1,1) & \\
 & U \uparrow & & & U \uparrow & & & U \uparrow & \\
 & \{\frac{1}{2}\}\ \mathbf{a} & & & \{\frac{1}{3}\}\ \mathbf{b} & & & \{\frac{1}{6}\}\ \mathbf{c} & \\
 & D \downarrow & & & D \downarrow & & & D \downarrow & \\
 & 2\ - & - \;\; - & - \;\; - \;\; - & -\ 2\ - & - \;\; - & - \;\; - \;\; - & -\ 2 & \\
a \swarrow & b \downarrow & \searrow c & a \swarrow & b \downarrow & \searrow c & a \swarrow & b \downarrow & \searrow c \\
(0,1) & (0,0) & (0,0) & (0,0) & (2,1) & (2,0) & (0,0) & (2,0) & (2,1)
\end{array}
$$

Find the PBE of this game.

Exercise 8

Consider the following signaling game with revealing messages, also called *"disclosure game"*: player 1 (Ann) is the informed player and plays at the beginning. Player 2 (Bob) is not informed and plays after player 1. Ann can be of two different types: $\theta^0 = 0$ or $\theta^1 = 1$. She has the opportunity to reveal (disclose) her type to Bob with a *verifiable* statement "my type is θ". False statements are severely punished, and we assume them away. This simplify the analysis without changing the results.

To be more formal, the set of signals is: $A_1 = \{R^0, R^1, N\}$ ($R =$"Reveal" and $N =$ "Not reveal"), and the set of feasible signals depends on θ: $A_1(\theta^0) = \{R^0, N\}$, $A_1(\theta^1) = \{R^1, N\}$. Thus, if Ann chooses R^k ($k = 0, 1$) she discloses that her type is θ^k. Payoffs are not directly affected by the action-message chosen by player 1.

Bob's goal is to correctly guess the type of Ann: his utility function is

$$u_2(\theta, a_2) = -(\theta - a_2)^2,\ a_2 \in \mathbb{R}.$$

Ann wants Bob to choose a high action a_2: her utility function is

$$u_1(\theta, a_2) = a_2.$$

(1) (*) Represent this game using trees and information sets (in the graphical representation you may pretend that Bob has only two actions, or you may use "fans").

(2) (**) Denote with $\mu = \mu(\theta^1|N)$ the conditional probability of θ^1 given the non-revealing signal $N \in A_1$. Find the best response of Bob as a function of μ. Find the best response of Bob to the revealing signals R^1 and R^0.

(3) (***) Show that pooling equilibria do not exist.

(4) (***) Show that μ is a uniquely determined in (perfect Bayesian) equilibrium. Determine the set of perfect Bayesian equilibrium assessments.