EconS 424 - Strategy and Game Theory

Why do we need Perfect Bayesian equilibrium?

*Asking for sequential rationality in sequential-move games with incomplete information**

March 28, 2013

1 Motivating example

Let us consider the following sequential-move game where player 1 decides to make a gift (G) or not make a gift (N) to player 2. Player 1 is privately informed about whether he is a "Friendly-type", or an "Enemy-type". Player 2, however, does not observe such information, and must decide whether to accept or reject player 1's gift.

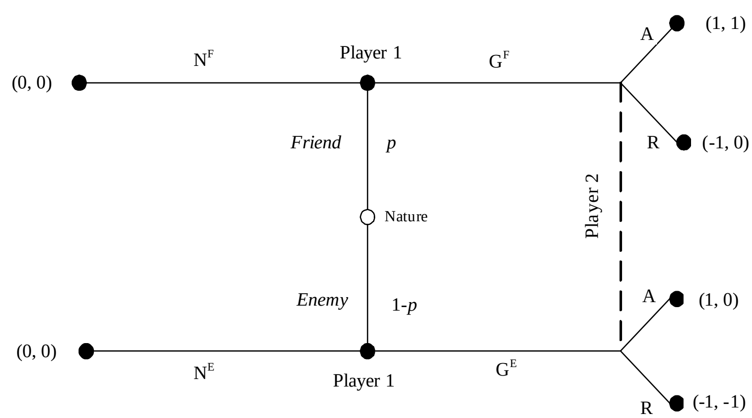

For additional practice, let us briefly analyze the set of BNE of this game. With that goal, let us first represent the above game tree in its Bayesian normal form representation, as depicted in the matrix below. (Recall that this matrix includes expected payoffs for each player. In addition, the uninformed player 2 has only two available strategies (two columns), whereas the informed player 1 has four different strategies (four rows))

*Félix Muñoz-García, School of Economic Sciences, Washington State University, 103G Hulbert Hall, Pullman, WA, 99163, fmunoz@wsu.edu.

		Player 2 A	R
	G^FG^E	$\underline{1}$,$\underline{p}$	-1, $p-1$
Player 1	G^FN^E	p,$\underline{p}$	$-p, 0$
	N^FG^E	$1-p$, $\underline{0}$	$p-1$, $p-1$
	N^FN^E	0,$\underline{0}$	$\underline{0}$,$\underline{0}$

Bayesian normal form representation

Underlining expected payoffs in order to identify each player's best response, we can conclude that there are two BNE in this game: (G^FG^E,A) and (N^FN^E,R).

But, one second, is the BNE (N^FN^E,R) sequentially rational for player 2? No!!

In this strategy profile, no type of sender makes a gift in equilibrium. Hence, if a gift is ever observed (a surprising event, that only happens off-the-equilibrium path), the receiver will compare the expected utility of accepting and rejecting the gift, based on the off-the-equilibrium belief μ, which denotes the probability that such a surprising gift originates from a friendly type, i.e., probability of being in the top right-hand node of the game tree. In particular, the receiver finds that

$$EU_2(A) = 1\mu + 0(1-\mu) = \mu, \quad \text{and}$$

$$EU_2(R) = 0\mu + (-1)(1-\mu) = \mu - 1$$

Hence, player 2 accepts the gift since $\mu > \mu - 1 \iff 0 > -1$. Importantly, this acceptance holds for any arbitrary off-the-equilibrium beliefs, μ, that player 2 might sustain upon observing a gift. Therefore, the gift rejection that the BNE (N^FN^E,R) prescribes cannot be sequentially rational.

2 Demanding sequential rationality to the BNEs

Can we find some equilibrium concept that selects equilibria predicting actions that are sequentially rational for all players, at any information set they are called to move? Yes, the Perfect Bayesian Equilibrium (PBE), which we analyze next. First, however, we must precisely define three concepts that must be satisfied in any PBE.

2.1 *Conditional beliefs about types*

First, note that player 2's initial beliefs about player 1's type coincide with nature's probability distribution p (which we refer as prior probabilities).

But when player 2 observes player 1's action, player 2 might learn something about player 1's type through his decisions. We say that player 2 **updates** his beliefs about player 1's type.

Example:

If P2 thinks that P1's optimal actions are N^F and G^E, and he observes that P1 is sending a gift, it must be that such gift only comes from the Enemy type (see figure below, with shaded branches for N^F and G^E). In terms of notation, we say that P2's beliefs are $\mu(Friend|G) = 0$, where $\mu(Friend|G)$ denotes the probability that player 2 believes to be in the top right-hand node (receving a gift from a Friend) *conditional* on that information set being reached (i.e., conditional on receiving a gift).

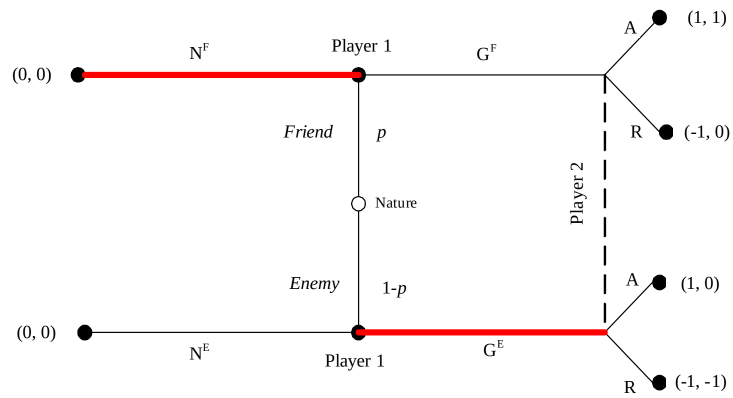

2.2 *Sequential rationality in an incomplete information environment*

We already encountered sequential rationality in SPNE, but now we apply it to games with incomplete information:

> at every information set at which every player is called to move, every player chooses the strategy that maximizes his expected utility level, given that all other players will do the same, and given his own beliefs about the other players' types.

Example: Player 2 accepts a gift if and only if $EU_2(A) > EU_2(R)$. That is, if

$$\begin{aligned} 1\mu + 0(1-\mu) &> 0\mu + (-1)(1-\mu) \\ \mu &> \mu - 1 \end{aligned}$$

which holds for all μ, since $\mu \in [0,1]$ by definition. [Note that in other exercises, we could be finding a cuotoff rule, i.e., player 2 prefering to accept if μ is relatively high, e.g., $\mu > \frac{2}{3}$, but reject otherwise.]

2.3 *Consistency of beliefs*

Let us denote by α^F the probability that player 1 plays G^F, and by α^E the probability that player 1 plays G^E. Then, player 2's belief that the gift coming from player 1 is in fact coming from a Friendly type, μ, can be expressed as

$$\mu = \frac{p\alpha^F}{p\alpha^F + (1-p)\alpha^E}$$

That is, player 2's beliefs are defined by the probability that a player 1 is a friendly type *and* he makes a gift, $p\alpha^F$, over the probability that player 1 is a friendly type and makes a gift in addition to the probability that player 1 is an enemy type and makes also a gift. In other words, we divide the probability that player 1 is a friendly type and makes a gift, $p\alpha^F$, over the probability that any type player 1 (friend or enemy) makes a gift. Hence, the consistency requirement that we imposse on beliefs is that beliefs must be found by using Bayes' rule. You may have seen this rule about conditional probabilities in you Stats class. Don't worry, we will just use it in the way that it is specified above. Next, I present one example.

Example: Let us assume that $p = \frac{1}{2}$, $\alpha^F = \frac{1}{8}$ and $\alpha^E = \frac{1}{16.}$. Then, player 2's posterior beliefs about P1 being a Friend after receiving a gift, μ, are

$$\mu = \frac{p\alpha^F}{p\alpha^F + (1-p)\alpha^E} = \frac{\frac{1}{2}\frac{1}{8}}{\frac{1}{2}\frac{1}{8} + \frac{1}{2}\frac{1}{16}} = \frac{2}{3}$$

A few notes about off-of-equilibrium beliefs:

a) If player 2's information set is not reached (what can only happen in this example if both types of player 1 choose to make no gifts), then the denominator in the above formula for Bayes' rule is zero (i.e., the probability of receiving a gift from any type of player 1 is zero, since $\alpha^F = 0$ and $\alpha^E = 0$). This makes μ to be indeterminate, since it especifically becomes $\frac{0}{0}$. In these cases, we are allowed to arbitrarily specify the value of μ (any value between zero and one, $\mu \in [0, 1]$). We will describe how to do it as generally as possible in worked-out exercises next.

b) In this case, we refer to μ as "off-of-equilibrium" beliefs, since it specifies beliefs about the probability of being in a node belonging to an information set that is actually unreached in equilibrium. In the above example, when both types of player 1 choose not to make a gift, then player 2's information set is never reached (i.e., it is an off-of-equilibrium event). Hence, in this case μ would specify off-of-equilibrium beliefs, and it can be arbitrarily specified, i.e., $\mu \in [0, 1]$.

c) Why do we care about off-of-equilibrium beliefs? Because they determine what P2 does in the event of receiving a gift. This can induce P1 to make gifts (or deter him from doing so), thereby affecting our equilibrium results.

3 Perfect Bayesian Equilibrium

We are now ready to combine the above 3 requirements for a PBE into its definition.

Definition of PBE. A strategy profile for all players $(s_1, s_2, ..., s_N)$ and beliefs μ over the nodes at all information sets are a PBE if:

a) Each player's strategies specify optimal actions, given the strategies of the other players, and given his beliefs.

b) The beliefs are consistent with Bayes' rule, whenever possible.

Note that the first bit of condition (a): "Each player's strategies specify optimal actions, given the strategies of the other players" resembles the condition for players' best responses in the definition of NE, whereas the last bit of this condition "...given his beliefs" resembles the definition of BNE for incomplete information games. Finally, note that condition (b) states that beliefs must be consistent with Bayes' rule "whenever possible". In particular, this last element of condition (b) is related with the previous note about off-of-equilibrium beliefs. Indeed, it is possible to specify beliefs which are consistent with Bayes' rule only when we are dealing with information sets that are reached in equilibrium. However, when we are in information sets that are unreached in equilibrium, Bayes' rule cannot be applied; and beliefs must be arbitrarily determined.

We will encounter different types of PBE. In a *separating* PBE different types of the privately informed player (e.g., player 1) behave differently. For instance, the friendly type makes gifts whereas the enemy type does not. In contrast, in the *pooling* PBE all types of the privately informed player behave similarly, e.g., all types of P1 make gifts.

Procedure to find PBE:[1]

1. Specify a profile of actions for the player with types (informed player 1), either separating or pooling.

   - In our above example, there can only be four possible profiles of actions for player 1: two separating strategy profiles, $N^F G^E$ and $G^F F^E$, and two pooling strategy profiles, $G^F G^E$ and $N^F N^E$.

2. Calculate the receiver's beliefs μ (the beliefs of the uninformed agent) using Bayes' rule at all information sets, both in-equilibrium and out-of-equilibrium.

   - In our above example, we only need to specify beliefs in one information set (that arising after player 2 receives a gift). Note however that this information set can either be reached in equilibrium (implying that μ specifies equilibrium beliefs) or not be reached in equilibrium (implying that μ specifies off-the-equilibrium beliefs).

3. Given μ from the previous step, find the optimal action (the optimal action of the uninformed player (player 2), i.e., player 2's optimal response given his beliefs about the type of P1 he faces.

4. Given the optimal action of the uninformed player, find the optimal action for the informed player.

5. Then check if this action profile for P1 coincides with the profile you suggested on step 1. If this is the case, then this strategy profile and beliefs can be supported as a PBE of the game. Otherwise, we say that this strategy profile cannot be sustained as a PBE of the game.

[1] You can find a more detailed description of this producedure on the short paper posted on the course website "A systematic procedure for finding Perfect Bayesian Equilibria in Incomplete Information Games." Here is the link: faculty.ses.wsu.edu/Munoz/Teaching/EconS491_Spring2011/Slides/Procedure_to_find_PBEs_June_2012.pdf

4 Worked-out Example

4.1 Separating equilibrium with $N^F G^E$

1. We first specify the profile of actions $N^F G^E$ for the informed player 1, i.e., player 1 makes gifts as an enemy but does not make gifts as a friend. For easy reference, the next figure shades the path $N^F G^E$ that this separating strategy profile describes.

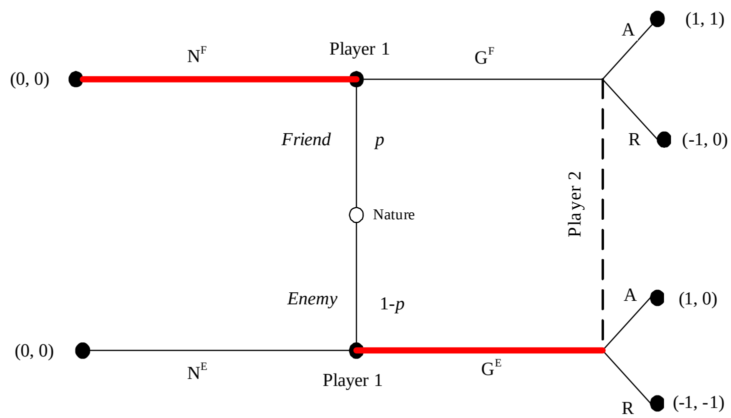

2. We can now update player 2's beliefs by using Bayes' rule: Since a gift can only come from an Enemy in this separating equilibrium,

$$\mu = \frac{p\alpha^F}{p\alpha^F + (1-p)\alpha^E} = \frac{p0}{p0 + (1-p)1} = \frac{0}{1-p} = 0$$

3. Optimal action for the responder (player 2) is to choose A, since he puts full probability to the event that he is called to move at the lower node of the information set, i.e., $0 > -1$.

   - We can now shade the branch labelled with A for player 2 (do it!!). Note that, since player 2 cannot distinguish player 1's type, he accepts any gift made to him, which graphically implies that we must shade branch A both in the upper and lower node.

4. Given player's optimal action (accept), we can now move to the informed player. If player 2 accepts the gift, then player 1's optimal action is to make a gift when he is a Friendly type, i.e., G^F, since $1 > 0$.

5. Therefore, the separating strategy profile $N^F G^E$, where player 1 makes gifts only when he is an enemy type, *cannot* be sustained as a PBE of this game.

4.2 Separating equilibrium with $G^F N^E$

1. We first specify the profile of actions $G^F N^E$ for the informed player 1, i.e., player 1 makes gifts as a friend but does not make gifts as an enemy. For easy reference, the next figure shades the path $G^F N^E$ that this separating strategy profile describes.

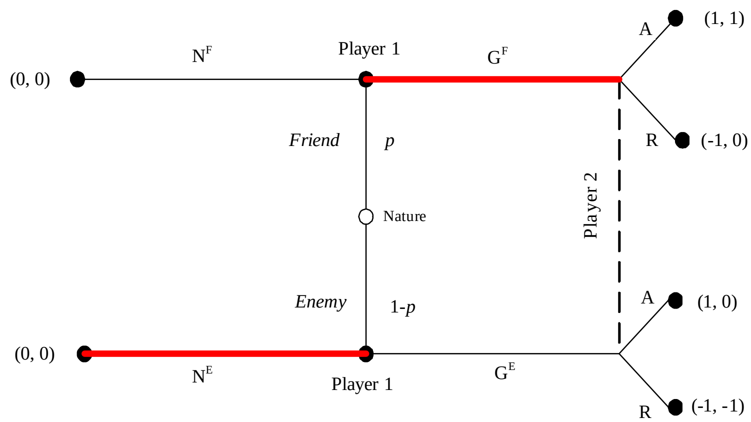

2. We can now update player 2's beliefs by using Bayes' rule: Since a gift can only come from an Friend in this separating equilibrium,

$$\mu = \frac{p\alpha^F}{p\alpha^F + (1-p)\alpha^E} = \frac{p1}{p1 + (1-p)0} = \frac{p}{p} = 1$$

3. Optimal action for the responder (player 2) is to choose A, since he puts full probability to the even that he is called to move at the upper node of the information set, i.e., $1 > 0$.

   - We can now shade the branch labelled with A for player 2 (do it!!). Note that, since player 2 cannot distinguish player 1's type, he accepts any gift made to him, which graphically implies that we must shade branch A both in the upper and lower node.

4. Given player's optimal action (accept), we can now move to the informed player. If player 2 accepts the gift, then player 1's optimal action is to send a gift both when he is a Friend (1>0) and when he is an Enemy (1>0), i.e., $G^F G^E$.

5. Then, the separating strategy profile $G^F N^E$, where player 1 makes gifts only when he is a friendly type, *cannot* be sustained as a PBE of this game either.

4.3 Pooling equilibrium with $G^F G^E$

1. We first specify the pooling profile of actions $G^F G^E$ for the informed player 1, i.e., player 1 makes gifts both as a friend and as an enemy. For easy reference, the next figure shades the path $G^F G^E$ that this pooling strategy profile describes.

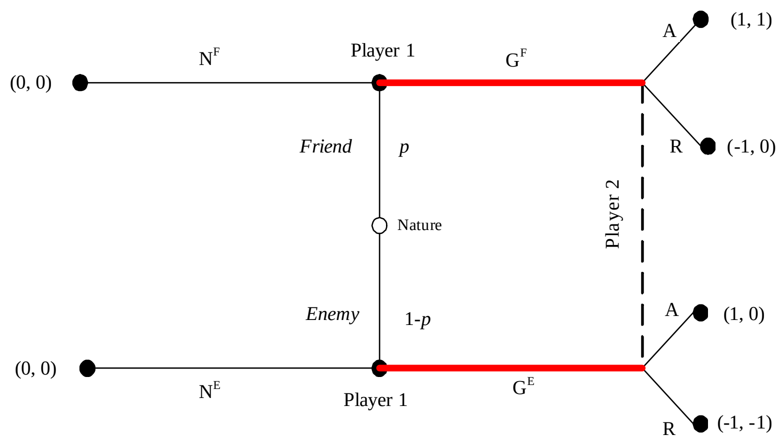

2. Update player 2's beliefs by using Bayes' rule: Since a gift can come from any type of player in this pooling equilibrium, observing a gift does not give player 2 any additional information about the actual type of player 1 he faces. Indeed, using Bayes' rule it is easy to check that player 2's posterior probability (his beliefs) coincide with the prior probability that player 1 is a Friendly type.

$$\mu = \frac{p\alpha^F}{p\alpha^F + (1-p)\alpha^E} = \frac{p1}{p1 + (1-p)1} = \frac{p}{p + 1 - p} = p$$

3. Optimal action for the responder (player 2) must be found by calculating player 2's expected utility from accepting and rejecting

$$EU_2(A) = 1p + 0(1-p) = p, \quad \text{and}$$

$$EU_2(R) = 0p + (-1)(1-p) = p - 1$$

   Hence, player 2 accepts the gift since $p > p - 1$ for any value of p.

   - We can now shade the branch labelled with A for player 2 (do it!!). Note that, since player 2 cannot distinguish player 1's type, he accepts any gift made to him, which graphically implies that we must shade branch A both in the upper and lower node.

4. Given player's optimal action (accept), we can now move to the informed player. Since player 2 accepts the gift, player 1's optimal action is to make a gift both when he is a Friend (1>0) and when he is an Enemy (1>0), i.e., $G^F G^E$.

5. Therefore, the pooling strategy profile $G^F G^E$, where both types of player 1 make gifts, *can* be supported as a PBE of this game.

4.4 Pooling equilibrium with $N^F N^E$

1. We first specify the pooling profile of actions $N^F N^E$ for the informed player 1, i.e., player 1 makes no gifts regardless of his type. For easy reference, the next figure shades the path $N^F N^E$ that this pooling strategy profile describes.

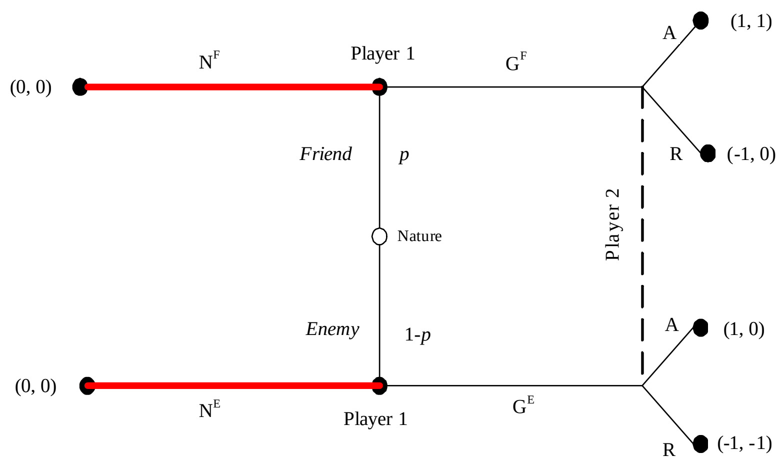

2. Update player 2's beliefs by using Bayes' rule: Since a gift is never observed in this pooling strategy profile, Bayes' rule is undetermined in this case (denominator is zero).

$$\mu = \frac{p\alpha^F}{p\alpha^F + (1-p)\alpha^E} = \frac{p0}{p0 + (1-p)0} = \frac{0}{0} \implies \mu \in [0, 1]$$

Receiving a gift in this case is considered and out-of-equilbrium event. As a result, μ can be arbitrarily specified in $\mu \in [0, 1]$.

3. Optimal action for the responder (player 2) must be found by calculating player 2's expected utility from accepting and rejecting (given a general value of μ, which specifies his out-of-equilibrium beliefs)

$$EU_2(A) = 1\mu + 0(1-\mu) = \mu, \quad \text{and}$$

$$EU_2(R) = 0\mu + (-1)(1-\mu) = \mu - 1$$

Hence, player 2 accepts the gift since $\mu > \mu - 1$ is satisfied for any value of μ.

- We can now shade the branch labelled with A for player 2 (do it!!). Note that, since player 2 cannot distinguish player 1's type, he accepts any gift made to him, which graphically implies that we must shade branch A both in the upper and lower node.

3. Given player's optimal action (accept), we can now move to the informed player. Since player 2 accepts the gift, then player 1's optimal action is to send a gift both when he is a Friend, G^F, and when he is an Enemy, G^E.

4. Therefore, the pooling strategy profile $N^F N^E$, where no type of player 1 makes gifts, *cannot* be supported as a PBE of this game.