Microeconomics II

Lecture 5: Equilibrium Refinements

Karl Wärneryd
Stockholm School of Economics

November 2016

• PROBLEM 5

Extending backward induction

Consider the following extensive form game.

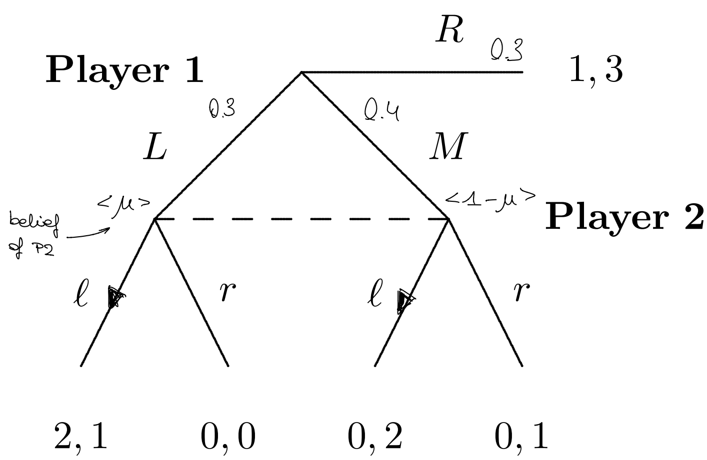

The game has the following normal form.

		Player 2	
		ℓ	r
Player 1	L	$2, 1$	$0, 0$
	M	$0, 2$	$0, 1$
	R	$1, 3$	$1, 3$

There are two equilibria in pure strategies, (L, ℓ) and (R, r).

Since the only subgame is the game itself, both equilibria are subgame perfect. Nevertheless, the (R, r) equilibrium is based on a non-credible threat. If it ever got to be Player 2's turn, he would never actually play r.

No subgame starts at Player 2's information set because he does not know at which of the two nodes he is, and a subgame has to start at a single node. But we could make something like a subgame out of it by specifying a *belief* held by Player 2 at his information set, and then continuing in the spirit of backward induction. A belief will be a probability distribution over the nodes in the information set. Clearly, admissible beliefs must depend on the strategies played, since in the (L, ℓ) equilibrium Player 2 will know for certain where he is.

That is, the Nash equilibrium concept induces beliefs at information sets reached on the equilibrium path, but not beliefs after zero-probability events. ?

So what should Player 2 believe if the (R, r) equilibrium is played and he nevertheless find he gets the move? This is exactly the kind of counterfactual or off-the-equilibrium-path reasoning that is central to the concept of subgame perfection. We shall investigate some different answers to this question. In the example, it turns out not to matter.

In the example, a belief for Player 2 is simply the probability μ that Player 2 assigns to Player 1 having played L.

To extend the notion of backward induction, it now seems reasonable to require that each player's action is rational given his beliefs, i.e., that it maximizes his expected payoff. In the example, Player 2's expectation from playing ℓ if it becomes his turn is $\mu \cdot 1 + (1 - \mu)2 = 2 - \mu$, while his expectation from playing r is $\mu \cdot 0 + (1 - \mu) \cdot 1 = 1 - \mu$. Hence it is strictly better to play ℓ for *all* possible beliefs.

Bayes' rule

In the example, Player 2 is unconstrained in the beliefs he may have at his information set. More generally, it seems reasonable to require that beliefs at different information sets are consistent with each other and with the strategies being played, and that they relate according to Bayes' rule whenever applicable.

Recall that if A and B are two possible events, we have that

$$\boxed{P(A|B) = \frac{P(A \text{ and } B)}{P(B)}.}$$

← share of cases where B happens and A also happens

In the example, suppose Player 1 plays L with probability .3, M with probability .4, and R with probability .3. Then Player 2's belief given that he gets to act must be

$$\mu = \frac{.3}{.3 + .4} = 3/7.$$

Example. *A patient has just been told by his physician that his result was positive on a test that may reveal whether or not an individual has a deadly disease. The precision of the test is as follows. If the subject has the disease, the test is positive in 98% of cases. If the subject does not have the disease, the test is negative in 98% of cases. In the population as a whole, one in 200 people has the disease. How concerned should the patient be?*

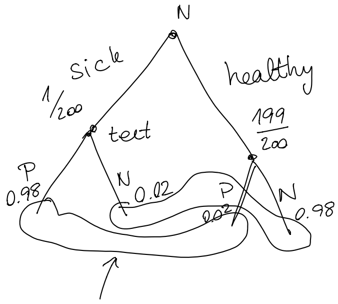

$$\frac{\frac{1}{200}\cdot 0.98}{\frac{1}{200}\, 0.98 + \frac{199}{200}\cdot 0.02} \approx 0.20$$ ← only 1/5 who test positive have the disease

The game show “paradox” (Monty Hall problem)

You are a contestant on a game show (similar to the US classic *Let’s Make a Deal*). You are asked to select one of three doors, behind one of which is \$1,000,000. After you have made your choice, the host, who knows what is behind each door, opens one of the other doors, behind which is nothing. You are then asked if you would like to stick with your original choice, or switch to the remaining unopened door.

In fact, you are always better off switching, since the probability that the million dollars is behind the remaining door is 2/3. This is a classic puzzle, with a solution many people (including eminent mathematicians) apparently find wildly counterintuitive or even wrong. The revelation that there is nothing behind one of the doors you did not choose seems to contain no relevant new information.

Equilibrium refinements

Ruling out dynamic irrationality of the kind that may appear in our first example is of course especially important in games of incomplete information, since they start with Nature selecting the types of the players and hence have no other subgames than the game itself.

Today, we study some equilibrium refinements that apply the type of reasoning we used in the example, i.e., that have in common that they require

- that players have beliefs, and that the beliefs are consistent (in ways to be specified), and
- that players are *sequentially rational* in the sense that at any point in the game they maximize their expected payoff for the continuation game, given their beliefs and the strategies being played.

Trembles (perturbed game)

We start by considering Selten's (1975) notion of *trembling-hand perfect*, or simply *perfect*, equilibrium. Consider the following normal form game.

		ε	$1-\varepsilon$
		Player 2	
		L	R
Player 1	T	$1,1$	$0,0$
	B	$0,0$	$0,0$

This game has two equilibria, (T, L) and (B, R). Consider the equilibrium (B, R). Suppose Player 1 thinks that Player 2 fully intends to play R, but with some small probability he makes a mistake (his hand trembles) and instead plays L. No matter how small this probability is, for Player 1 to play T is now a strict best reply. Hence the (B, R) equilibrium is not stable against perturbations. (This also coincides with its involving weakly dominated strategies, which is always the case in two-player finite normal form games.)

$$P1: \begin{cases} T: 1\varepsilon + 0(1-\varepsilon) \\ B: 0\varepsilon + 0(1-\varepsilon) \end{cases} \Rightarrow T \succ B \Rightarrow \varepsilon > 0$$

The more general notion of perfection in normal form games involves the idea of games where the players are required to play all their pure strategies with positive probability. A strategy σ_i is *completely mixed* if $\sigma_i(s_i) > 0$ for all s_i.

Definition. *A strategy profile σ of a normal form game is a* perfect equilibrium *if there is a sequence of completely mixed strategy profiles $\sigma^n \to \sigma$ such that for all i, $u_i(\sigma_i, \sigma^n_{\sim i}) \geq u_i(s_i, \sigma^n_{\sim i})$ for all s_i.*

In the example, if Player 2 plays L with positive probability, playing T with probability one is the best reply of Player 1, and conversely. Hence there is a sequence of strategy profiles where more and more weight is put on L and T, respectively, such that L and T are best responses, and hence (L, T) is perfect.

Considering normal-form trembles is not enough, however, to rule out all behavior that seems suspicious in extensive form games. A perfect equilibrium of the normal form is not necessarily subgame perfect.

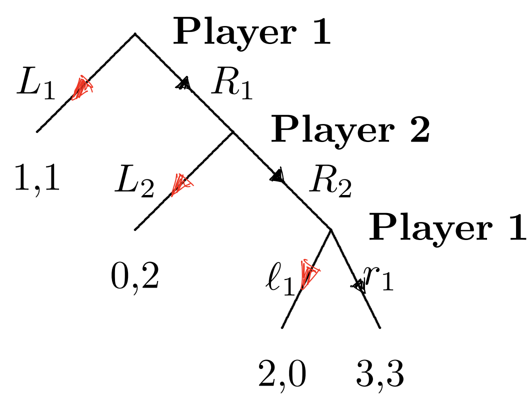

The unique subgame perfect equilibrium of this game is $(R_1 r_1, R_2)$. But consider the normal form of the game.

		Player 2	
		L_2	R_2
Player 1	$L_1\ell_1$	$1, 1$	$1, 1$
	$L_1 r_1$	$1, 1$	$1, 1$
	$R_1\ell_1$	$0, 2$	$2, 0$
	$R_1 r_1$	$0, 2$	$3, 3$

Consider the equilibrium $(L_1\ell_1, L_2)$. The strategy $L_1\ell_1$ is a best reply to any strategy that puts weight sufficiently close to 1 on L_2. The strategy L_2 is a best response to any strategy that puts weight sufficiently close to 1 on $L_1\ell_1$ and weight sufficiently higher on $R_1\ell_1$ than on R_1r_1. Hence $(L_1\ell_1, L_2)$ is a perfect equilibrium.

To avoid this, Selten proposes that normal-form perfection for extensive form games should be replaced by perfection in the *agent-strategic form*, in which each information set is considered to belong to a separate agent, who gets the same payoffs in the end as the original player whose information set it is.

In the example, we would then have three agents, Player 1's first agent, who controls the action at his first information set, Player 2, and Player 1's second agent, who controls the action at his second information set. Now $(L_1\ell_1, L_2)$ is not perfect, because it can never be optimal for Player 1's second agent to play ℓ_1, given that there is some positive probability of his information set being reached.

Note that the issue of beliefs never arises under perfection, since all the relevant conditional probabilities are induced by the completely mixed strategies, i.e., Bayes' rule is always applicable in the perturbed games.

Sequential equilibrium

Kreps and Wilson (1982) suggest a weakening of perfection for extensive form games. As under perfection, they require that beliefs be derivable from some trembles on the part of the players, but these trembles need not correspond to optimal strategies.

A player is now required to have a belief μ_i that consists of a probability distribution over the nodes in the information set for every one of his information sets. A profile (σ, μ) is called an *assessment*. Let Ψ be the set of all possible assessments.

Definition. *An assessment (σ, μ) is* sequentially rational *if for any information set h and strategy $\sigma'_{i(h)}$ we have that*

$$u_{i(h)}(\sigma \mid h, \mu(h)) \geq u_{i(h)}((\sigma'_{i(h)}, \sigma_{\sim i(h)}) \mid h, \mu(h)).$$

Let Σ^0 be the set of all completely mixed behavioral strategies. Given $\sigma \in \Sigma^0$, the conditional probabilities of all nodes follow from Bayes' rule. Let Ψ^0 be the set of assessments such that $\sigma \in \Sigma^0$ and μ is induced by σ.

Definition. *An assessment (σ, μ) is* consistent *if $(\sigma, \mu) = \lim_{n\to\infty}(\sigma^n, \mu^n)$ for some sequence (σ^n, μ^n) in Ψ^0.*

Definition. *An assessment* (σ, μ) *is a* sequential equilibrium *if it is sequentially rational and consistent.*

The point of sequential equilibrium is that it requires that beliefs at off-the-equilibrium-path information sets be the limits of the conditional probabilities induced by some sequence of strategies with trembles that converges to the actual strategies, and that strategies at every information set be optimal given those beliefs.

Every perfect equilibrium is a sequential equilibrium, but the converse is not true. Hence sequential equilibrium is a weaking of perfect equilibrium. For generic games, the concepts coincide.

Perfect Bayesian equilibrium

Fudenberg and Tirole (1991) suggest a weakening of sequential equilibrium that they call *perfect Bayesian equilibrium.* Unlike sequential equilibrium, PBE is not defined for general extensive form games.

We give an informal definition of a PBE as follows.

Definition. *An assessment is a* perfect Bayesian equilibrium *if*

- *each player's strategy maximizes his expected payoff at each information set given his beliefs and the strategies of all players for the continuation game, and*
- *each player's beliefs follow from Bayes' rule, whenever applicable, given the strategies of all players.*

That is, a PBE essentially requires that strategies at information sets reached with zero probability be such that there are *some* beliefs that respect the actual conditional probabilities and make the strategies optimal.

$$\frac{\varepsilon}{\varepsilon + (1-\varepsilon)2\varepsilon} = \frac{1}{3-2\varepsilon} \xrightarrow[\varepsilon\to 0]{} \frac{1}{3}$$

$$\frac{\varepsilon}{\varepsilon + (1-\varepsilon)\frac{2\varepsilon}{1-\varepsilon}} = \frac{1}{3}$$

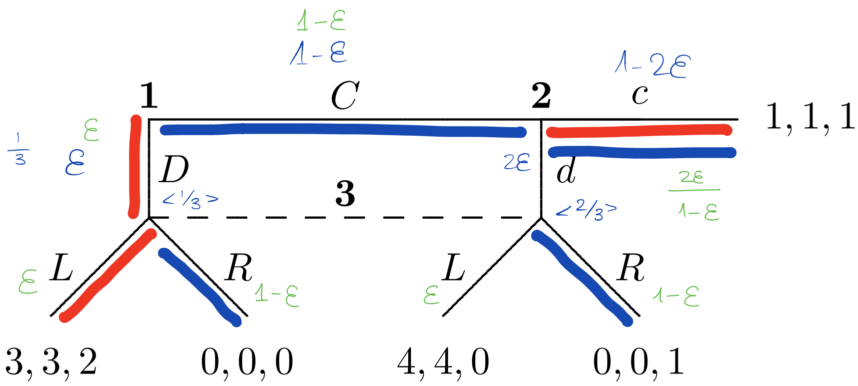

To relate these different concepts, consider the example known as "Selten's horse." It has two pure-strategy equilibria, one where $\sigma_1(D) = 1$, $\sigma_2(c) = 1$, and $\sigma_3(L) = 1$, and another where $\sigma_1(C) = 1$, $\sigma_2(c) = 1$, and $\sigma_3(R) = 1$.

The first equilibrium is not a sequential, perfect Bayesian, or perfect equilibrium, since it fails sequential rationality. Player 2's action is dominated if he gets the move.

The second equilibrium is sequential together with a belief on the part of Player 3 with $\mu_3(D) = 1/3$. To see that this is a consistent belief, consider a sequence of strategy profiles with $\sigma_1(C) = 1-\varepsilon$ and $\sigma_2(c) = 1-2\varepsilon$. Under any such profile the conditional probability that D has been played, given that Player 3 gets the move, is $\varepsilon/(\varepsilon + (1-\varepsilon)2\varepsilon) = 1/(3-2\varepsilon)$, which converges to $1/3$ as $\varepsilon \to 0$. It follows that this is also a PBE.

To see that the second equilibrium is also perfect, consider trembles such that $\sigma_1(C) = 1-\varepsilon$, $\sigma_2(d) = 2\varepsilon/(1-\varepsilon)$, and $\sigma_3(R) = 1-\varepsilon$.

Signaling games

We may define a class of finite signaling games as follows. There are two players. Player 1 is of type θ and takes an action a_1. Player 2 observes Player 1's action, but not his type, and takes an action a_2, whereupon both get payoffs.

An equilibrium of a signaling game is a strategy profile (σ_1, σ_2) such that for all θ, we have that

$$\sigma_1(\cdot \mid \theta) \in \arg\max_{\alpha_1} u_1(\alpha_1, \alpha_2, \theta),$$

(i.e., that Player 1 plays a best reply) and $\sigma_2(\cdot \mid a_1)$ maximizes

$$\sum_\theta p(\theta) \left(\sum_{a_1} \sum_{a_2} \sigma_1(a_1 \mid \theta) \sigma_2(a_2 \mid a_1) u_2(a_1, a_2, \theta) \right),$$

Player 2's *ex ante* expected payoff. The latter condition is equivalent to Player 2's action maximizing his expected payoff at every information set that occurs with positive probability under Player 1's strategy, since otherwise he could increase his *ex ante* expectation by changing his choice at some information set. That is, we need that for all a_1 such that $\sigma_1(a_1 \mid \theta) > 0$ for some θ, we have that $\sigma_2(\cdot \mid a_1)$ maximizes

$$\sum_\theta p(\theta \mid a_1) u_2(a_1, \alpha_2, \theta) =$$

$$\sum_\theta \frac{p(\theta)\sigma_1(a_1 \mid \theta)}{\sum_{\theta'} p(\theta')\sigma_1(a_1 \mid \theta')} u_2(a_1, \alpha_2, \theta).$$

A PBE of such a signaling game is an assessment $(\sigma_1, \sigma_2, \mu)$ such that (σ_1, σ_2) is a Nash equilibrium and Player 2's belief μ is such that $\mu(\theta) = p(\theta \mid a_1)$ for every a_1 that occurs with positive probability under σ_1 (and p). That is, a PBE in this setting does not constrain beliefs at all at unreached information sets, it just specifies that there must be *some* belief that supports Player 2's actions there. Hence only dominated actions are ruled out.

Consider the following signaling game.

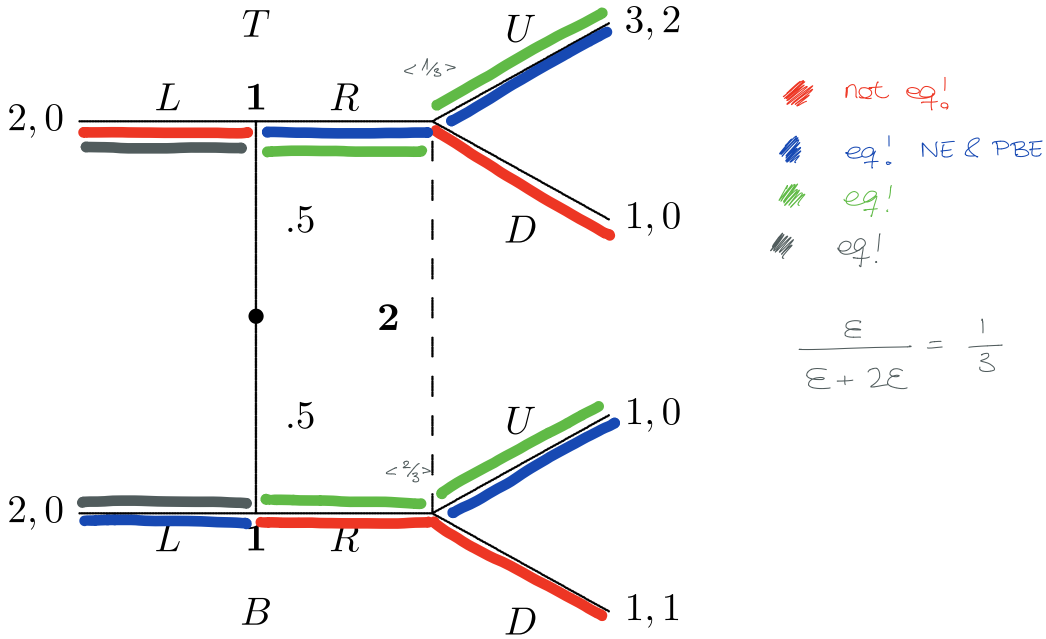

Let μ be the probability Player 2 assigns to Player 1 being of type B when his information set is reached.

Does this game have a *separating* PBE?

There are two possibilities. Suppose Player 1 plays L when of type T and R when of type B. Then we have $\mu = 1$, and Player 2 must play D. But then Player 1 would prefer to play L also when of type B, so this is not an equilibrium.

prob P2 assigns to P1 being B

Suppose Player 1 plays R when of type T and L when of type B. Then we have $\mu = 0$, and Player 2 must play U. Since given this, both types of Player 1 strictly prefer their choices, we do indeed have a PBE.

Is there a *pooling* PBE?

Suppose Player 1 always plays R. Then we have $\mu = .5$, and Player 2 must play U. But then Player 1 prefers L when he is of type B. (Of course, Player 1 always prefers L when he is of type B.) not an eq!

Suppose Player 1 always plays L. This is a best reply for Player 1 if Player 2 plays D with probability at least .5. It is rational for Player 2 to play D if we have $\mu \geq 2/3$. Hence there are PBEs where Player 1 always plays L, Player 2 plays D, and Player 2 has a belief $\mu > 2/3$.

Any PBE in this class of games is also a sequential equilibrium, since *any* belief of Player 2 at an unreached information set can be made consistent.

Spence's (1974) job-market signaling game

A worker is of ability θ_L or $\theta_H > \theta_L$, with prior probabilities p_L and p_H, respectively. The worker makes an investment e in education. An employer observes e, but not the type of the worker, and, because the labor market is competitive, pays a wage w equal to the expectation of the worker's type. To model the employer's payoff explicitly, assume it is $-(w-\theta)^2$. The worker's payoff is $w - (e/\theta)$.

We now look for PBEs in pure strategies. Let e_L and e_H be the equilibrium choices of a low and high type. There are now two types of PBE.

Pooling PBE: In a pooling equilibrium, we have $e_L = e_H = e^\star$. Since observing the education choice of the worker then reveals no new information, we must have $w^\star = p_L\theta_L + p_H\theta_H$. In order for it to be rational for both types of the worker to invest $e^\star$, the wage $w(e)$ paid at any deviation must be such that

$$w^\star - \frac{e^\star}{\theta_L} \geq \overbrace{w(e) - \frac{e}{\theta_L}}^{\text{deviation: invest some } e \neq e^\star \text{ (more or less)}}$$

prior of employer: if $e \neq e^\star$, worker is $\theta_L \Rightarrow w(e) = \theta_L$.

and

$$w^\star - \frac{e^\star}{\theta_H} \geq w(e) - \frac{e}{\theta_H}.$$

↳ then θ_L's optimal response is $e = 0$!

The simplest way to get this to hold in a PBE is for the employer to believe that the worker's type is θ_L for any $e \neq e^\star$. Then $w(e) = \theta_L$ for $e \neq e^\star$, and the optimal deviation for the worker is then to set $e = 0$ when of low type, so the incentive conditions reduce to

$$w^\star - \frac{e^\star}{\theta_L} \geq \theta_L,$$

which is equivalent to

$$e^\star \leq \theta_L p_H(\theta_H - \theta_L).$$

Separating PBE: In a separating equilibrium, the two types of worker make different investments, and hence the education choice perfectly reveals the worker's type. We must have $e_L = 0$, since the wage paid to the low type is then θ_L regardless of e_L. Assuming the employer believes the worker is of low type for any education level not e_L or e_H, the incentive conditions are then

$$\theta_L \geq \theta_H - \frac{e_H}{\theta_L} \quad \longleftarrow \text{ } \theta_L \text{ should not want to "act" as a } \theta_H$$

and

$$\theta_H - \frac{e_H}{\theta_H} \geq \theta_L, \quad \longleftarrow \text{ } \theta_H \text{ should not want to "act" as } \theta_L$$

which implies

$$\theta_L \underbrace{(\theta_H - \theta_L)}_{\theta_d} \leq e_H \leq \theta_H(\theta_H - \theta_L). \implies e_H \in [\theta_L\theta_d, \theta_H\theta_d]$$

Problem. *Find all equilibria, subgame perfect equilibria, and perfect Bayesian equilibria in pure strategies in the following two games. (Note that beliefs matter most at information sets that are not reached with positive probability in equilibrium.)*

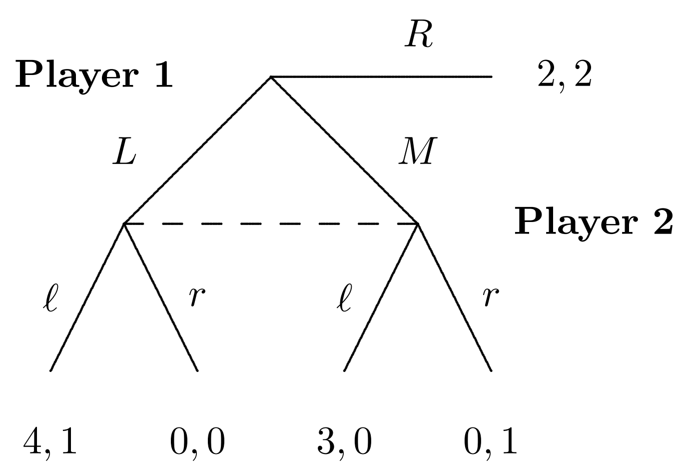

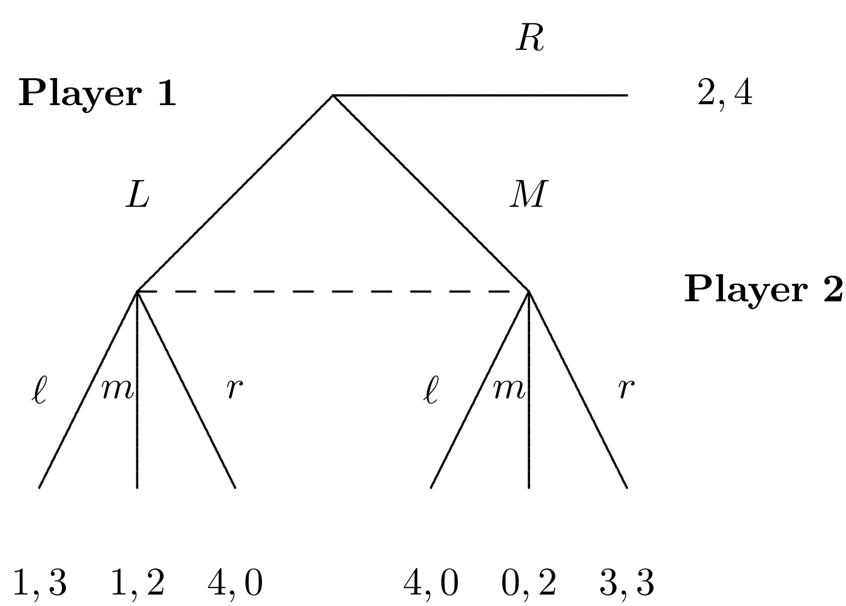

Problem. *Show that the following game has no perfect Bayesian equilibrium in pure strategies. Then find a PBE in mixed strategies.*

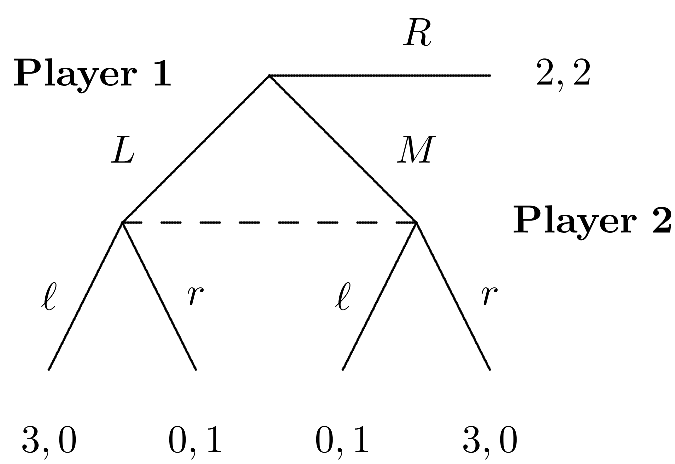

Problem. *Consider the following signaling game (due to Cho and Kreps, QJE 1987). Player 1's type may be either Surly or Wimpy, with probabilities 0.9 and 0.1, respectively. Knowing his type, Player 1 enters a tavern and must order either beer or quiche. He prefers beer if he is surly, but he prefers quiche if he is wimpy. Player 2 does not know 1's type but, after seeing what player 1 has ordered, player 2 must decide whether to fight with player 1 or just leave him alone and go away. Find the sequential and perfect Bayesian equilibria of this game.*

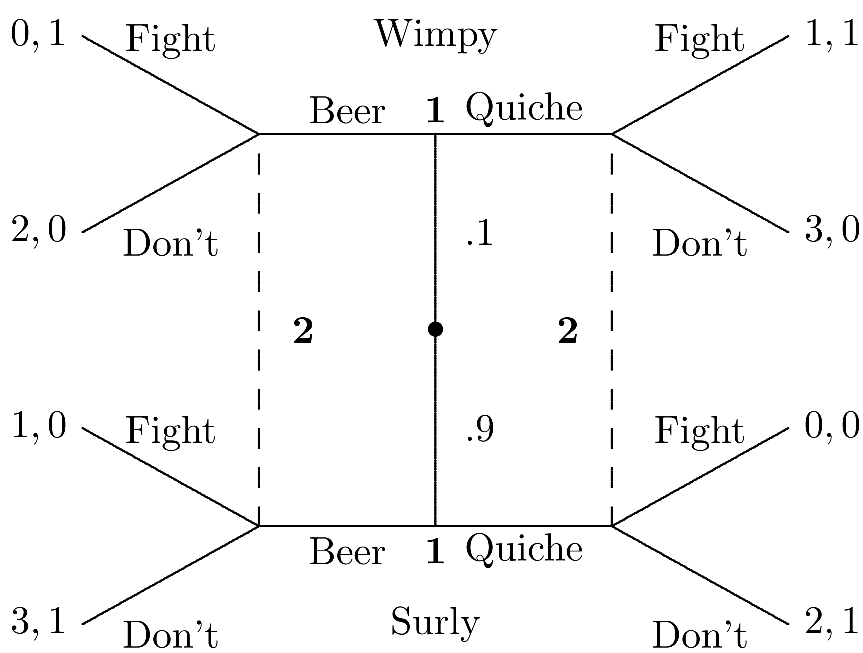

Problem. *Player 1 is in a car accident with Player 2. Player 1 knows whether he is to blame or not, but Player 2 does not. If the case goes to court the judge will learn the truth with certainty. Player 1 now makes a take-it-or-leave it pre-trial offer to Player 2, which gives Player 2 either 3 or 5. Player 2 either accepts or rejects the offer. If he rejects the offer, the case goes to court, both players have to pay legal expenses of 6, and Player 1 has to pay 5 to Player 2 if he is guilty and 0 otherwise. Formulate this as a signaling game and find its sequential equilibria.*

Problem. *Consider a sequence of firms,* $i = 1, 2, \ldots$, *who each in turn have to decide whether to switch to a new production technology or not. Each firm observes the choices of those preceding it in the sequence. Switching yields a benefit of* 1 *with probability* $.5$ *and* 0 *with probability* $.5$, *and costs* $.5$. *Each firm observes the choice of those preceding it in the sequence, but also gets an independent, private signal* $x_i \in \{L, H\}$. *If the benefit of switching is* 1, *then* $x_i = H$ *with probability* $p > .5$. *If the benefit of switching is* 0, *the probability of* $x_i = H$ *is* $1 - p$. *Assume a firm that is indifferent switches with probability* $.5$. *An informational cascade (Bikhchandani, Hirshleifer, and Welch 1992) is said to start when it is rational for a firm to ignore its own signal, given the observed choices of previous firms. Find the probability that an informational cascade starts with Firm 3.*

PROBLEM 1

a)

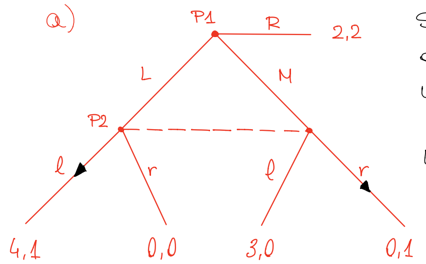

Since P2 does not know at which node he is, and there is no ex-ante better-off option, we cannot use backward induction.

Let me write the game in normal form,

		P2	
		ℓ	r
P1	L	4, 1	0, 0
	M	3, 0	0, 1
	R	2, 2	2, 2

$NE = \{(L, \ell), (R, r)\}$

SPE = NE (only whole subgame)

Let me now check for PBE:

- (L, ℓ): assign probability p for P1 playing L (trembling hand). Since $p=1$ is the correct belief given P1's strategy (P1 is playing L) and P2 maximises his payoff by playing ℓ, $(L, \ell; p=1)$ is PBE

- (R, r): notice that P2's information set is never reached, so this case is more involved. If P2's information set was reached, P2's expected payoff from playing r would be

$$E_r = p0 + (1-p)1 = 1-p$$

Similarly for ℓ,

$$E_\ell = p1 + (1-p)0 = p$$

In order to sustain (R, r) a PBE, it must be rational to play r! Hence

$$r \succeq \ell \implies 1-p \geq p \implies p \leq 1/2$$

Hence, $(R, r; p \in [0, 1/2])$ is a PBE

b)

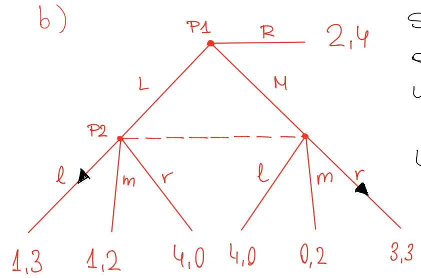

Since P2 does not know at which node he is, and there is no ex-ante better-off option, we cannot use backward induction.

Let me write the game in normal form,

		P2		
		ℓ	m	r
P1	L	1,3	1,2	4,0
	M	4,0	0,2	3,3
	R	2,4	2,4	2,4

$NE = \{(R, m)\}$

SPE = NE (only whole subgame)

Let me now check for PBE. Notice that P2's information set is never reached, so this case is more involved. If P2's information set was reached, P2's expected payoff would be

$$E_\ell = 3p + 0(1-p) = 3p$$

$$E_m = 2p + 2(1-p) = 2$$

$$E_r = 0p + 3(1-p) = 3-3p$$

In order to sustain (R, m) as PBE, it must be rational to play m! Hence

$$m \geq \ell \Rightarrow 2 \geq 3p \Rightarrow p \leq 2/3$$

$$m \geq r \Rightarrow 2 \geq 3-3p \Rightarrow p \geq 1/3$$

Hence, $(R, r; p \in [1/3, 2/3])$ is a PBE

PROBLEM 2

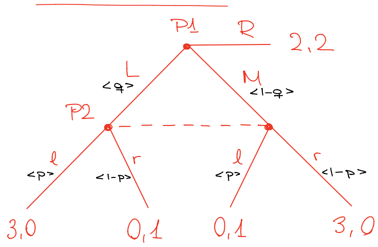

Since P2 does not know at which node he is, and there is no ex-ante better-off option, we cannot use backward induction.

Let me write the game in normal form,

P2

		ℓ	r
P1	L	3, 0	0, 1
	M	0, 1	3, 0
	R	2,2	2,2

Notice there are no pure NE. Let me find the mixed NE. Suppose P2 plays ℓ with prob p (and r with prob 1-p). If P1 plays, he obtains

$$E_L = 3p + 0(1-p) = 3p$$

$$E_M = 0p + 3(1-p) = 3(1-p)$$

$$E_R = 2$$

Hence, P1 plays R whenever

$$R \geq L \Rightarrow 2 \geq 3p \Rightarrow p \leq 2/3$$

$$R \geq M \Rightarrow 2 \geq 3-3p \Rightarrow p \geq 1/3$$

As long as $p \in [1/3, 2/3]$, P1's best reply is R. Let me now move to P2. If P1 plays R, then P2 is indifferent between ℓ and r (check the normal form game). Hence, $q = 1/2$ would be the required belief. Since this equilibrium is only satisfied when $p \in [1/3, 2/3]$,

$$PBE = \{(R, p\ell + (1-p)r ; p \in [1/3, 2/3], q = 1/2)\}$$

There are many other PBE's. Notice from the normal form game that if P2's information set is ever reached (P1 plays L or M), P2 will not be indifferent between ℓ and r.

PROBLEM 3

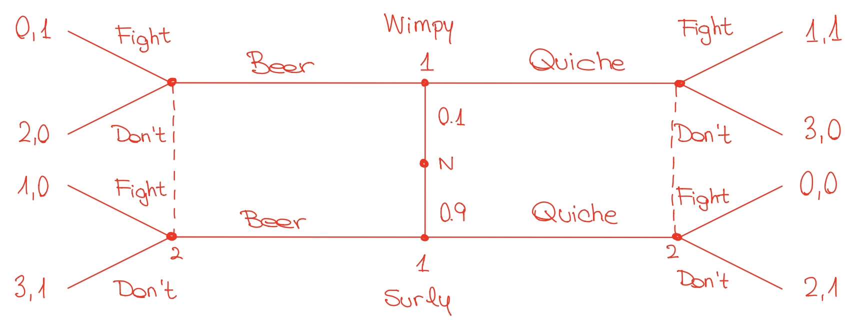

Let me first consider the separating equilibria (SE)

- (B, Q) : in this case P2 can form priors about P1. P2 knows that when P1 is Wimpy will play Beer. Hence P2 will choose Fight. On the other hand, when P1 is Surly will play Quiche. In that case P2 will play Don't. Therefore, this equilibrium is

  "(Beer if Wimpy, Quiche if Surly ; Fight if Beer, Don't if Quiche)"

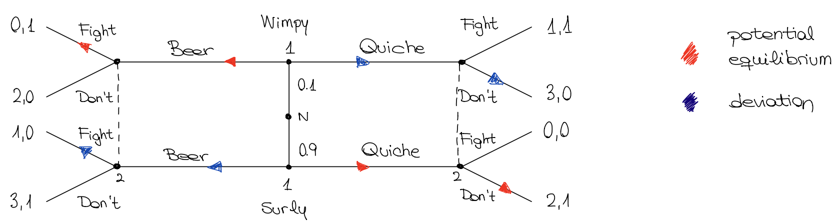

  Is there any incentive to deviate for P1? If P1 as Wimpy plays Quiche instead, he will be mistakenly thought as a Surly, and P2 will play Don't. By this, P1 as Wimpy would get a higher payoff ($0 < 3$), thus it is not an equilibrium!

  Note: when P1 is Surly there would be no incentive to deviate ($2 > 1$)

- (Q, B) : in this case P2 can form priors about P1. P2 knows that when P1 is Wimpy will play Quiche. Hence P2 will choose Fight. On the other hand, when P1 is Surly will play beer. In that case P2 will play Don't. Therefore, the equilibrium is

  "(Quiche if Wimpy, Beer if Surly ; Fight if Quiche, Don't if Beer)"

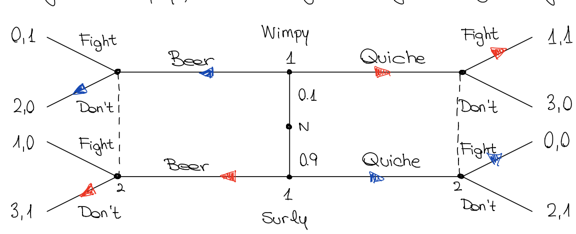

  Is there any incentive to deviate for P1? If P1 as Wimpy plays Beer instead, he will be mistakenly thought as a Surly, and P2 will play Fight. By this, P1 as Wimpy would get a higher payoff ($1 < 2$), thus it is not an equilibrium!

Note: when P1 is Surly there would be no incentive to deviate $(3>0)$.

Let me now consider the pooling equilibria (PE)

- (Q, Q): in this case P1 cannot form any prior on P1. P2's expected payoffs are

$$\mathbb{E}_F = 0.1 \cdot 1 + 0.9 \cdot 0 = 0.1$$

$$\mathbb{E}_D = 0.1 \cdot 0 + 0.9 \cdot 1 = 0.9$$

Hence, P2 will play Don't. But what will P2 do if he observes Beer? Consider δ his prior on P1 playing Beer on Wimpy (and $1-\delta$ on P1 playing Beer while being Surly). Payoffs for P2 would be

$$\mathbb{E}_F = 1\delta + 0(1-\delta) = \delta$$

$$\mathbb{E}_D = 0\delta + 1(1-\delta) = 1-\delta$$

Hence, P2 plays Fight if $\delta \geq 1/2$, and Don't otherwise. The goal now is to define δ such that there are no possible deviations (i.e., find an equilibrium). P1 is now obtaining payoffs of 3 (Wimpy) and 2 (Surly). Notice that if P2 plays Don't, P1 would like to deviate when Surly. However, if P2 plays Fight no version of P1 would deviate ($3>0$ if Wimpy, $2>1$ if Surly). Hence, we need $\delta \geq 1/2$

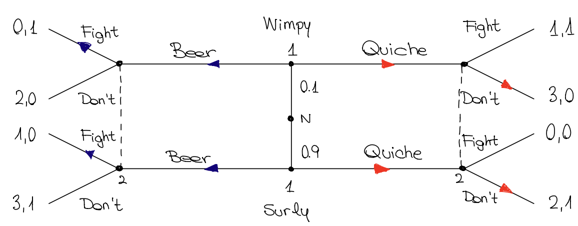

$PE = \left\{ \text{"}\left(\text{Quiche always; Don't if Quiche, Fight if Beer; } \delta \in [1/2, 1]\right)\text{"} \right\}$

- (B, B): in this case P1 cannot form any prior on P1. P2's expected payoffs are

$$\mathbb{E}_F = 0.1 \cdot 1 + 0.9 \cdot 0 = 0.1$$

$$\mathbb{E}_D = 0.1 \cdot 0 + 0.9 \cdot 1 = 0.9$$

Hence, P2 will play Don't. But what will P2 do if he observes Quiche? Consider δ his prior on P1 playing Quiche on Wimpy (and $1-\delta$ on P1 playing Quiche while being Surly). Payoffs for P2 would be

$$\mathbb{E}_F = 1\delta + 0(1-\delta) = \delta$$

$$\mathbb{E}_D = 0\delta + 1(1-\delta) = 1-\delta$$

Hence, P2 plays Fight if $\delta \geq 1/2$, and Don't otherwise. The goal now is to define δ such that there are no possible deviations (i.e., find an equilibrium). P1 is now obtaining payoffs of 2 (Wimpy) and 3 (Surly). Notice that if P2 plays Don't, P1 would like to deviate when Wimpy. However, if P2 plays Fight no version of P1 would deviate ($2>1$ if Wimpy, $3>0$ if Surly). Hence, we need $\delta \geq 1/2$

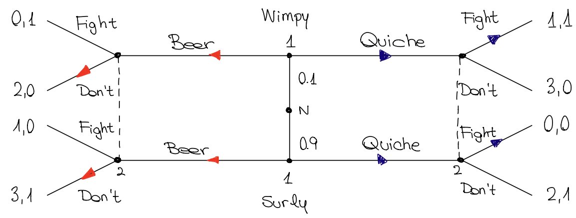

PE = { " (Beer always ; Don't if Beer, Fight if Quiche ; $\delta \in [1/2, 1]$) " }

PROBLEM 4

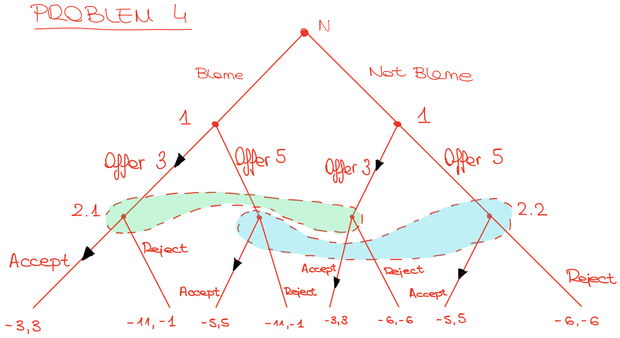

Notice that, even though it is a game of asymmetric information, we can apply backward induction. It is always optimal for P2 to accept. Given this, it is P1's best response to offer the lower amount. Hence, the equilibria are

NE = {(Blame ; Offer 3, AcceptAccept), (Not Blame ; Offer 3, AcceptAccept)}

PROBLEM 5 ← ask to do!

The expected value of changing technology for F1 depends only on the signal (since it is the first firm, there is no previous experience). If F1 receives signal $x_1 = H$, the expected value of changing technology is

$$\mathbb{E}(V|H) = \frac{\mathbb{P}(V=1,\ x_1=H)}{\underbrace{\mathbb{P}(V=1,\ x_1=H)}_{\text{"If the benefit is 1, then } x_i=H \text{ with prob } p>0.5\text{"}} + \underbrace{\mathbb{P}(V=0,\ x_1=H)}_{\text{"If the benefit is 0, the prob. of } x_i=H \text{ is } 1-p\text{"}}} = \frac{0.5p}{0.5p + 0.5(1-p)} = p$$

And the expected value of changing technology after observing signal $x_1 = L$ is

$$\mathbb{E}(V|L) = \frac{\mathbb{P}(V=1,\ x_1=L)}{\mathbb{P}(V=1,\ x_1=L) + \mathbb{P}(V=0,\ x_1=L)} = \frac{0.5(1-p)}{0.5(1-p) + 0.5p} = 1-p$$

Given that $\mathbb{E}(V|H) = p$ and $\mathbb{E}(V|L) = 1-p$, F2 can conclude that if he observes F1 changing technology then the signal F1 received was H, and if not changing, L.

If F2 receives signal H, the expected value of changing technology is

$$\mathbb{E}(V \mid H_1 \wedge H_2) = \mathbb{P}(V=1 \mid x_1 = H \ \wedge \ x_2 = H) = \frac{\mathbb{P}(V=1,\ x_1 = H,\ x_2 = H)}{\mathbb{P}(V=1,\ x_1 = H,\ x_2 = H) + \mathbb{P}(V=0,\ x_1 = H,\ x_2 = H)} =$$

$$= \frac{0.5p^2}{0.5p^2 + 0.5(1-p)^2} = \frac{p^2}{p^2 + (1-p)^2} > 0.5$$

Hence, F2 changes technology if F1 already changed it and F2 obtains signal H. What about if F2 observes signal L?

$$\mathbb{E}(V \mid H_1 \wedge L_2) = \mathbb{P}(V=1 \mid x_1 = H \ \wedge \ x_2 = L) = \frac{\mathbb{P}(V=1,\ x_1 = H,\ x_2 = L)}{\mathbb{P}(V=1,\ x_1 = H,\ x_2 = L) + \mathbb{P}(V=0,\ x_1 = H,\ x_2 = L)} =$$

$$= \frac{0.5p(1-p)}{0.5p(1-p) + 0.5(1-p)p} = 0.5$$

(identical result if F1 didn't change technology (i.e. F1 observed L) and F2 observed H, $\mathbb{E}(V \mid L_1 \wedge H_2) = 0.5$). Now consider F3. If F1 and F2 have both changed technology, and F3 has received signal L. The expected value for changing technology is

$$\mathbb{E}(V \mid H_1 \wedge H_2 \wedge L_3) = \mathbb{P}(V=1 \mid x_1 = H,\ x_2 = H,\ x_3 = L) =$$

$$= \frac{\mathbb{P}(V=1,\ x_1 = H,\ x_2 = H,\ x_3 = L)}{\mathbb{P}(V=1,\ x_1 = H,\ x_2 = H,\ x_3 = L) + \mathbb{P}(V=0,\ x_1 = H,\ x_2 = H,\ x_3 = L)} =$$

$$= \frac{0.5p^2(1-p)}{0.5p^2(1-p) + 0.5p(1-p)^2} = p > 0.5$$

Hence, it is rational for F3 to change technology even after observing signal $x_3 = L$.