FDPE Microeconomics 3: Game Theory
Spring 2018
Lecture notes 4: Games of incomplete information
Pauli Murto

1 Bayesian games

- The last part of the course considers games of incomplete information. Additional material:
  - MWG 8.E, 9
  - Fudenberg and Tirole Ch. 6, Ch. 8
  - Osborne and Rubinstein Ch. 2.6, 11, 12
  - Myerson Ch. 3.9 - 3.11
  - Maschler, Solan and Zamir Ch. 9 - 10

1.1 Modeling incomplete information

- So far we have assumed that players know each other's characteristics, in particular, their preferences
- It is clear that this is often unrealistic
- How other players behave, depends on their payoffs
- Therefore, player i's payoffs depends on his *beliefs* of others' payoffs, and so i's *actions* will depend on his beliefs
- But then, other players' optimal actions depend on what they believe about i's beliefs
- ... and so i's optimal action depends on his beliefs about other players' beliefs about his beliefs, and so on
- A full model of incomplete information should specify the *belief hierarchy* containing beliefs of all orders (belief over payoffs, beliefs over beliefs, beliefs over beliefs over beliefs, and so on)
- Harsanyi (1967-68) showed how such a model can be specified in a tractable way
- His insight is to assume that a random variable, "nature's move", specifies the "type" for each player, where the "type" of player i contains the information about i's payoffs as well as i's beliefs of all orders
- The probability distribution of this random variable is assumed to be common among the players (*common prior*) and the players then use Bayesian rule to reason about probabilities

- In applications, resulting belief hierarchies are typically very simple
- Mertens and Zamir (1985) showed that in principle one can construct a rich enough type space to model any situation of incomplete information in such a manner
- It should be mentioned that the assumption of "common prior" underlying Harsanyi model is critical
- Note that Harsanyi model in effect just models incomplete information as a standard extensive form game with imperfect information (nature takes the first move, and players are asymmetrically informed about this move)
- However, it is more practical to treat such models as a separate class of games that are called Bayesian games

1.2 Bayesian game

- A game of incomplete information, or a Bayesian game, is defined by:

1. Set of players, $\mathcal{I} = \{1, 2, ..., I\}$.
2. Set of possible types for each player: Θ_i, $i \in \mathcal{I}$. Let $\theta_i \in \Theta_i$ denote a typical type of player i. We adopt the following notation $\theta = (\theta_1, ..., \theta_I)$, $\theta_{-i} = (\theta_1, ..., \theta_{i-1}, \theta_{i+1}, ..., \theta_I)$ etc.
3. Natures move: θ is drawn from a joint probability distribution F on $\Theta = \Theta_1 \times \cdots \times \Theta_I$ (common prior)
4. Set of actions available to each player: A_i, $i \in \mathcal{I}$. Let $a_i \in A_i$ denote a typical action taken by i.
5. Strategies, $s_i : \Theta_i \to A_i$, for $i \in \{1, 2, ..., I\}$. The action that type θ_i takes is then given by $s_i(\theta_i) \in A_i$. Denote the strategy space of i by S_i.
6. Payoffs, $u_i(a_1, ..., a_N; \theta_1, ..., \theta_N)$.

- The game proceeds as follows:
  - First, nature chooses θ according to F.
  - Then, each player i observes the realized type $\widehat{\theta}_i$ and updates her beliefs of other players types based on F. Denote the distribution on θ_{-i} conditional on $\widehat{\theta}_i$ by $F_i\left(\theta_{-i} \mid \widehat{\theta}_i\right)$.
  - Players then choose their actions simultaneously (we may also interprete those as representing full action plans in extensive form games)
- If Θ_i are finite, then F is just a discrete probability distribution on Θ, and we may denote by $p(\theta_1, ..., \theta_I)$ the probability of realization $\theta = (\theta_1, ..., \theta_I)$.

- Then the expected payoff of i given $\widehat{\theta}_i$ and profile s is simply

$$
\begin{aligned}
& \mathbb{E}_{\theta_{-i}}\left[u_i\left(s_i\left(\widehat{\theta}_i\right), s_{-i}\left(\theta_{-i}\right), \left(\widehat{\theta}_i, \theta_{-i}\right) \middle| \widehat{\theta}_i\right)\right] \\
= \; & \Sigma_{\theta:\theta_i=\widehat{\theta}_i} u_i\left(s_i\left(\widehat{\theta}_i\right), s_{-i}\left(\theta_{-i}\right), \theta\right) p_i\left(\theta_{-i} \mid \widehat{\theta}_i\right)
\end{aligned}
$$

- In many applications type space is continuous, and then the expected payoff is defined analogously using intergral instead of summation

- Some classifications of models:

  - If for all i, u_i is independent of θ_j, $j \neq i$, then we have private values. Otherwise, the model has interdependent values.
  - If θ_i, $i = 1, ..., I$, are distributed independently, we have a model with indepenent types. Otherwise, types are correlated.
  - The simplest case is independently distributed, private values

1.3 Bayesian Nash equilibrium

- Bayesian Nash Equilibrium is just the Nash Equilibrium in the current context:

- That is, each player chooses a strategy s_i that is a best response to other players strategies s_{-i}:

Definition 1 *A strategy profile* $(s_1, ..., s_N)$ *is a Bayesian Nash Equilibrium if*

$$
\mathbb{E}_\theta\left[u_i\left(s_i\left(\theta_i\right), s_{-i}\left(\theta_{-i}\right); \theta\right)\right] \geq \mathbb{E}_\theta\left[u_i\left(s_i'\left(\theta_i\right), s_{-i}\left(\theta_{-i}\right); \theta\right)\right]
$$

for all $s_i' \in S_i$.

- Another way to think about a Bayesian equilibrium is to note that if s is a best response to s_{-i}, then each possible type θ_i must be playing a best reponse to the conditional distribution of the other player's types.

- Hence, the above definition can be restated as follows: a profile s is a Bayesian Nash equilibrium if and only if, for all i and all $\widehat{\theta}_i \in \Theta_i$ occurring with positive probability,

$$
\begin{aligned}
& \mathbb{E}_{\theta_{-i}}\left[u_i\left(s_i\left(\widehat{\theta}_i\right), s_{-i}\left(\theta_{-i}\right), \left(\widehat{\theta}_i, \theta_{-i}\right) \middle| \widehat{\theta}_i\right)\right] \\
\geq \; & \mathbb{E}_{\theta_{-i}}\left[u_i\left(s_i'\left(\widehat{\theta}_i\right), s_{-i}\left(\theta_{-i}\right), \left(\widehat{\theta}_i, \theta_{-i}\right) \middle| \widehat{\theta}_i\right)\right]
\end{aligned}
$$

  for all $s_i' \in S_i$.

- We can also allow mixed strategies as before

1.3.1 Cutoff strategies

- With many types and actions, strategy spaces are large
- In many applications, best-response strategies are monotonic in types
- This often leads to tractable models that exploit cutoff strategies
- Consider a binary game with a high action a^H and a low action a^L
- If $u_i(a^H, a_{-i}, \theta_i, \theta_{-i}) - u_i(a^L, a_{-i}, \theta_i, \theta_{-i})$ is monotonic in θ_i for all a_{-i}, then all best responses take the form of a cutoff strategy:

$$\theta_i < \theta_i^* \Rightarrow a_i = a^L,$$

$$\theta_i \geq \theta_i^* \Rightarrow a_i = a^H.$$

- You have many examples based on this technique in problem set 4.

1.3.2 Example: Interpreting mixed strategies through incomplete information (purification)

- We have often relied on mixed strategies in our analysis of games
- Sometimes mixed strategies raise objections. Why would people randomize?
- Harsanyi suggested that mixed strategies may be interpreted as pure strategies in an incomplete information game, in the limit where incomplete information vanishes
- To see how this works, consider the following game

	L	R
U	$0, 0$	$0, -1$
D	$1, 0$	$-1, 3$

- This has a unique equilibrium in mixed strategies: $\sigma_1(U) = 3/4$, $\sigma_1(D) = 1/4$, $\sigma_2(L) = \sigma_2(R) = 1/2$
- Let us change the game so that payoffs are given by

	L	R
U	$\varepsilon\theta_1, \varepsilon\theta_2$	$\varepsilon\theta_1, -1$
D	$1, \varepsilon\theta_2$	$-1, 3$

where $\varepsilon > 0$ is a small number and θ_1 and θ_2 are independent random variables uniformly distributed over $[0, 1]$

- θ_1 is private information to player 1 and θ_2 is private information to player 2

- This is a Bayesian game with independent private values

- Note that here $u_1((U, a_2), \theta_1) - u_1((D, a_2), \theta_1)$ is increasing in θ_1 for all a_2, and similarly $u_2((a_1, L), \theta_2) - u_2((a_1, R), \theta_2)$ is increasing in θ_2 for all a_1, so best responses are cutoff strategies

- In particular, there should be some cutoff levels θ_1^* and θ_2^* such that player 1 chooses U iff $\theta \geq \theta_1^*$ and player 2 chooses L iff $\theta \geq \theta_2^*$

- For such cutoff strategies to be an equilibrium, the cutoff types should be indifferent between the two actions, which holds if:

$$\begin{aligned} \varepsilon\theta_1^* &= (1-\theta_2^*)\cdot 1 + \theta_2^* \cdot (-1) \text{ and} \\ \varepsilon\theta_2^* &= (1-\theta_2^*)\cdot(-1) + \theta_2^* \cdot 3 \end{aligned}$$

- Solving these gives equilibrium cutoffs:

$$\begin{aligned} \theta_1^* &= (2+\varepsilon)/\left(8+\varepsilon^2\right) \\ \theta_2^* &= (4-\varepsilon)/\left(8+\varepsilon^2\right) \end{aligned}$$

- Check that then all other types are choosing strictly optimal action

- Ex ante, player 1 chooses U with probability $1-(2+\varepsilon)/\left(8+\varepsilon^2\right)$ and D with probability $(2+\varepsilon)/\left(8+\varepsilon^2\right)$, whereas player 2 chooses L with probability $1-(4-\varepsilon)/\left(8+\varepsilon^2\right)$ and R with probability $(4-\varepsilon)/\left(8+\varepsilon^2\right)$

- Let $\varepsilon \to 0$, and note that these probabilities converge to the unique mixed strategy equilibrium probabilities of the original game

1.3.3 Electronic mail game

- Recall that the *type* may entail not only what the payoffs of a player are, but also what player's believe about each others payoffs, what other players believe about what other players believe about payoffs, and so on

- In other words, a type should capture a belief hierarchy

- In the example considered above (as in much of the literature), belief hierarchies are very simple: if types are independently distributed, then all beliefs of order two or higher are degenerate (i *knows* j's belief for sure, because prior is common, and j's private information does not affect his belief over i's payoffs)

- To demonstrate a more complicated case, consider the following game analyzed by Rubinstein (1989)

- Each of the two players has to choose action A or B
- With probability $p < 1/2$ the payoffs are given by game G_b and with probability $1-p$ payoffs are given by game G_a:

Game G_a:

	A	B
A	M, M	$1, -L$
B	$-L, 1$	$0, 0$

Game G_b:

	A	B
A	$0, 0$	$1, -L$
B	$-L, 1$	M, M

where $L > M > 1$.

- In both games, it is mutually beneficial to choose the same action, but the best action depends on the game
- Note that B is the more risky action: even if the true game is G_b, choosing B is bad if the other player chooses A
- Assume first that player 1 knows the true game, but player 2 does not
- Then this is a very simple game of incomplete information, where both player choose A in the unique Bayesian equilibrium (see why?)
- On the other hand, if both players know the game, then there is an equilibrium where (A, A) is played in game G_a and (B, B) is played in game G_b
- The interesting case is the following: suppose that only player 1 knows the true state, and players communicate through a special protocol as follows
- If the game is G_a, then there is no communication
- If the game is G_b, then player 1's computer sends an automatic message to player 2
- If player 2's computer receives a message, then it sends automatically a confirmation to player 1
- If player 1's computer receives a confirmation, then it sends automatically a further confirmation to player 2, and so on
- The confirmations are sent automatically, but in each transmission there is a small probability ε that the message does not get through
- If a message does not get through, then communication ends
- At the end of the communication phase each player sees on her computer screen exactly how many messages her computer has sent
- To model this as a Bayesian game, define a type Q_1 of player one to be the number of messages her computer has sent, and type Q_2 of player two to be the number of messages her computer has sent

- Then, if $Q_1 = 0$, the game is G_a, otherwise it is G_b
- Note that both players know the true game, except type $Q_2 = 0$ of player 2 (she is quite convinced that game is G_a is ε is small)
- This is a Bayesian game with payoffs

$$u_1\left(\left(Q_1, Q_2\right), (A, A)\right) = \begin{cases} M \text{ if } Q_1 = 0 \\ 0 \text{ if } Q_1 > 0 \end{cases}$$

.... and so on

- The common prior is:

$$\Pr\left(\left(Q_1, Q_2\right) = \left(q_1, q_2\right)\right) = \begin{cases} & 1-p \text{ if } q_1 = q_2 = 0 \\ p\varepsilon\left(1-\varepsilon\right)^{q_1+q_2-1} & \text{if } q_1 \geq 1 \text{ and } \left(q_2 = q_1 - 1 \text{ or } q_2 = q_1\right) \\ & 0 \text{ otherwise} \end{cases}$$

- How to compute players' beliefs?
- Consider player 1. She knows her own type Q_1, and she knows that player 2's type is either $Q_2 = Q_1 - 1$ or $Q_2 = Q_2$ (depending on whether her own message, or the next message by player 2, failed to go through). Therefore, his belief of player 2's types are $\Pr\left(Q_2 = Q_1 - 1\right) = \varepsilon / \left(\varepsilon + \left(1-\varepsilon\right)\varepsilon\right) > 1/2$ and $\Pr\left(Q_2 = Q_1\right) = \left(\left(1-\varepsilon\right)\varepsilon\right) / \left(\varepsilon + \left(1-\varepsilon\right)\varepsilon\right) < 1/2$ (and $\Pr\left(Q_2 = q\right) = 0$ for all other q)
- Similarly, player 2 knows her own type Q_2 and she knows that player 1's type is either $Q_1 = Q_2$ or $Q_1 = Q_2 + 1$
- Hence, her beliefs of player 2's types are $\Pr\left(Q_1 = Q_2\right) = \varepsilon / \left(\varepsilon + \left(1-\varepsilon\right)\varepsilon\right) > 1/2$ and $\Pr\left(Q_1 = Q_2 + 1\right) = \left(\left(1-\varepsilon\right)\varepsilon\right) / \left(\varepsilon + \left(1-\varepsilon\right)\varepsilon\right) < 1/2$ (and $\Pr\left(Q_1 = q\right) = 0$ for all other q)
- What are the higher order beliefs?
- Claim: there is a unique Bayesian equilibrium, where all types play (A, A)
- To prove the result: show first that for type $Q_1 = 0$, it is a dominant action to choose A
- Next, show that then it is optimal for $Q_2 = 0$ to choose A, then also for type $Q_1 = 1$, then $Q_2 = 1$, and so on... (check)
- So, even when 256 messages have gone through, players play (A, A) (no matter how large number M is)
- The problem is: even if both players know for sure that game is G_b, it is not *common knowledge*

- An event is common knowledge among the players if all the players know the event, all the players know that all the players know the event, all the players know that all the players know that all the players know the event, and so on
- Can you see that in the electronic mail game event "game is G_b" is not common knowledge even when $Q_1 = 256$ and $Q_2 = 256$ (think first about e.g. case $Q_1 = 2$ and $Q_2 = 1$)

2 Dynamic games of incomplete information

- The concepts that we defined above apply to extensive form games: one may think of the game as a strategic form representation of extensive form
- But as before, the idea of sequential rationality requires modeling explicitly the dynamic structure of the game
- In dynamic games of incomplete information, nature first draws types for the players, each player observes her own type, and then the players play some extensive form game
- With incomplete information, sub-game perfectness refinement has no bite, so the relevant refinements used here are the perfect Bayesian equilibrium and sequential equilibrium
- The games we consider fall in the class of "Multi-stage games with observed actions and incomplete information", or "Bayesian extensive game with observed actions"
- This is a Bayesian variant of the multi-stage games with observed actions; the only uncertainty is about the types of the players
- In such a model:
  - First, nature chooses a private type for each player
  - Then, players play a multi-stage game, where at the beginning of each stage, all previous actions are observed (except the initial move by nature)
  - We assume here that the type distribution is more restricted than in general Bayesian games: types are independentely distributed across the players
- More precicely: the game starts by the nature choosing a type $\theta_i \in \Theta_i$ for all players and types are independent:

  $$p(\theta) = p_1(\theta_1) \cdot ... \cdot p_I(\theta_I)$$

  and this distribution is common knowledge.

- Now we may summarize information at stage t as a list of previous actions by all players:

$$h^t = (a_1, ..., a_{t-1}).$$

- Without loss of generality, we may assume that all the players have a move in each period (if not, then we may let i choose from a one-element set $A_i(h^t) = \{a\}$)

- A behavior strategy profile σ now assigns for each h^t a distribution over actions that depends on type: $\sigma_i(a_i | h^t, \theta_i)$ is the probability with which i chooses action $a_i \in A_i(h^t)$ given h^t and her type θ_i

- At the beginning of t, players know exactly all the actions taken in the past, so uncertainty concerns just the types of the other players. The belief system can therefore be summarized as

$$\mu(h^t) = \left\{\mu_i\left(\theta_{-i} \left| h^t\right.\right)\right\}_{i \in \mathcal{I}},$$

defined for all h^t, where $\mu_i(\theta_{-i} | h^t)$ is i's probability assessment of the other players' types

- Following Fudenberg-Tirole, it is reasonable to define Perfect Bayesian Equilibrium in this context by imposing several natural extra requirements in addition to weak consistency for the belief system:

C1: At all h^t, players share a common belief on each others' types. That is, players i and j have identical belief, denoted $\mu(\theta_k | h^t)$, on k's type:

$$\mu_i\left(\theta_k \left| h^t\right.\right) = \mu_j\left(\theta_k \left| h^t\right.\right) := \mu\left(\theta_k \left| h^t\right.\right) \text{ for all } h^t,\ \theta_k \in \Theta_k, \text{ and } i \neq j \neq k.$$

Moreover, these assesments remain independent across players throughout the game:

$$\mu\left(\theta \left| h^t\right.\right) = \mu\left(\theta_1 \left| h^t\right.\right) \cdot ... \cdot \mu\left(\theta_I \left| h^t\right.\right).$$

C2: Other players' belief of player i's type do not depend on actions by $j \neq i$ (even if i takes an unexpected action):

$$\mu\left(\theta_i \left| \left(h^t, a^t\right)\right.\right) = \mu\left(\theta_i \left| \left(h^t, \widehat{a}^t\right)\right.\right) \text{ whenever } a_i^t = \widehat{a}_i^t.$$

C3: Bayes rule is applied whenever possible. That is, for all i, h^t, and $a_i^t \in A_i(h^t)$, if there exists θ_i with $\mu(\theta_i | h^t) > 0$ and $\sigma_i(a_i^t | h^t, \theta_i) > 0$, then

$$\mu\left(\theta_i \left| \left(h^t, a_i^t\right)\right.\right) = \frac{\mu(\theta_i | h^t)\, \sigma_i(a_i^t | h^t, \theta_i)}{\sum_{\theta_i' \in \Theta_i} \left(\mu\left(\theta_i' | h^t\right) \sigma_i\left(a_i^t \left| h^t, \theta_i'\right.\right)\right)}.$$

- Note that this applies also when history h^t is reached with probability 0.

- In particular: players update their beliefs about player i using Bayes' rule until her behavior contradicts her strategy, at which point they form a new common belief about i's type. From then on, this new belief will serve as the basis of future Bayesian updating
- Bayes rule also applies to beliefs about player i if some other player j takes an unexpected action (i.e. if $\sigma_k\left(a_j^t\left|h^t,\theta_j\right.\right)=0$ for all $\theta_j\in\Theta_j$)
- To summarize: conditions C1 - C3 basically say that after each h^t, all players $-i$ have a common probability distribution on i's type, denoted by $\mu\left(\theta_i\left|h^t\right.\right)$, and this is derived from $\mu\left(\theta_i\left|h^{t-1}\right.\right)$ by Bayesian rule whenever that is applicable.
- Moreoever, if $\mu\left(\theta_i\left|h^t\right.\right)$ cannot be derived by Bayesian rule from $\mu\left(\theta_i\left|h^{t-1}\right.\right)$, it is independent on actions a_j^t, $j\neq i$
- Sequential rationality is as before: σ is sequentially rational if for all i and all h^t,

$$u_i(\sigma|h^t,\theta_i,\mu\left(\theta_{-i}\left|h^t\right.\right))\geq u_i(\sigma|h^t,\theta_i,\mu\left(\theta_{-i}\left|h^t\right.\right))\text{ for all }\sigma_i'\in\Sigma_i.$$

- A perfect Bayesian equilibrium (PBE) is a pair (σ,μ) such that σ is sequentially rational (given μ), and μ satisfies C1-C3.
- Literature often uses the concept of PBE rather than sequential equilibrium in this class of games, because it is simpler and more easily checked
- But PBE as defined here is closely related to sequential equilibrium:
  - Every sequential equilibrium is PBE
  - By a result by Fudenberg-Tirole (1991), if either each player has at most two possible types, or if there are two periods, then the set of perfect Bayesian equilibrium coincides with the set of sequential equilibria.
  - These conditions apply to many applications, so in those cases there is no difference between sequential equilibrium and PBE.

2.1 Example: Spence's signalling model

- A worker's talent (= value to employer) is either low or high: $\theta\in\left\{\theta^L,\theta^H\right\}$, $\theta^L<\theta^H$
- $\Pr\left(\theta=\theta^H\right)=p>0$.
- The worker knows her talent but the employer does not
- Employer offers a wage w

- Suppose that employer minimizes $(w-\theta)^2$, so that given her belief, optimal wage is $w = E(\theta)$.
- This is just a short-cut way to model a labor market, where competition drives wage to the expectation of talent
- Before seeking a job, the worker chooses the level of education $e \geq 0$
- We assume that education does not affect productivity, but has cost e/θ
- The payoff for worker is $w - e/\theta$ if she accepts the job with wage w
- The game proceeds as follows:
  - Worker observes her type θ
  - Worker chooses education level e
  - Employer offers wage w
  - Employer accepts or rejects, and the game ends
- There are many equilibria. We will next look separately at pooling equilibria and separating equilibria.

Pooling equilibrium

- In a pooling equilibrium, both type of workers choose the same education level $e^L = e^H = e^*$
- Since the employer learns nothing from education level, she offers wage $w^* = (1-p)\theta^L + p\theta^H$
- For this to be an equilibrium, it can not be optimal for worker to choose some $e \neq e^*$
- The easiest way to satisfy this is to consider a belief system for the employer, where she believes that any deviation from education level e^* originates from a worker type θ^L
- Hence consider employer's strategy $w(e^*) = w^*$, $w(e) = \theta^L$ for $e \neq e^*$
- When is this an equilibrium?

Separating equilibrium

- In a separating equilibrium both types choose different education levels and hence employer can tell them apart
- Clearly, type θ^L should choose $e^L = 0$

- To ensure that no type wants to mimic the other, we must have

$$\theta^L \geq \theta^H - e^H/\theta^L \text{ and } \theta^H - e^H/\theta^H \geq \theta^L,$$

or

$$\theta^L \left(\theta^H - \theta^L\right) \leq e^H \leq \theta^H \left(\theta^H - \theta^L\right).$$

- Since $\theta^H > \theta^L$, a continuum of feasible values of e^H exists
- Check that you can complete those to a PBE
- This model has a lot of PBE
- Note: this is a multi-stage game with observable actions with at most two types, so sequential equilibria are the same as PBE
- There is a large literature that considers refinements to sequential equilibrium to narrow down the plausible predictions
- In this case, the so called *intuitive criterium* by Cho and Kreps (1987) selects the best separating equilibrium (see MWG Chapter 13, Appendix A)

2.2 Reputation effects: chain-store game with incomplete information

- We consider here a simple model of reputation following seminal papers by Kreps and Wilson (1982) and Milgrom and Roberts (1982)
- Consider the following variant of the chain-store game
- A single long-run incumbent firm faces potential entry by a series of short-run firms
- Each entrant plays only once but observes all previous play
- Payoff matrix

	Fight if entry	Accommodate if entry
Enter	$-1, -1$	$b, 0$
Stay out	$0, a$	$0, a$

- Note that an entrant enters if she considers probability of incumbent fighting to be less than $b/(b+1)$
- As we observed earlier, this game with a finite number of entrants has a unique subgame perfect equilibrium, where every entrant enters and the incumbent accommodates every time

- With an infinite horizon, there are in fact many equilibria, including one where every entrant enters, and one where entry is deterred (if discount factor is high enough). Which one should we expecte to be played?
- Introduce incomplete information: with a probability $p^0 > 0$ the incumbent is tough and prefers to fight (with complementary probability the incumbent is weak and gets payoffs as shown in the table)
- Assume also that each entrant is tough with probability $q^0 > 0$ and prefer to enter no matter how incumbent responds. Weak entrant gets the payoffs shown in the table.
- Clearly, a one period game has a sequential equilibrium, where weak entrant enters if $p^0 < b/(b+1)$ and stays out if $p^0 > b/(b+1)$
- (and tough entrant enters with probability 1, tough incumbent fights with probability 1, and weak incumbent accommodates with probability 1)
- Now consider a game with a finite number of periods, where incumbent maximizes the sum of payoffs over periods
- If the incumbent could credibly commit to fighting every entrant, then it would be optimal to do so if $a\left(1-q^0\right) > q^0$, that is
$$q^0 < a/(a+1).$$
- Assume that this inequality holds. We will see that under that condition, the incumbent gets close to that behavior in PBE
- The reason why a weak incumbent might fight is that this could make it look more likely to the entrants that the incumbent is tough
- To see this, suppose that entrants' belief that incumbent is tough is p, and suppose that a weak incumbent fights with probability π
- Then, if entry takes place and incumbent fights, p jumps to
$$p'(p,\pi) = \frac{p}{p+\pi(1-p)} \geq p$$
by Bayesian rule.
- Clearly, $p'(p,\pi)$ is decreasing in π with $p'(p,1) = p$ and $p'(p,0) = 1$
- That is, if weak incumbent fights with probability 1, then entrants learn nothing. And if weak incumbent fights with probability 0, then entrants learn perfectly incumbent's type
- Consider the second last period of the game
- If $a\left(1-q^0\right) > 1$, then the incumbent finds it worthwhile to fight in the current period if this deters entry (of weak entrant) in the final period

- For simplicity, suppose that this is the case, that is:

$$q^0 < (a-1)/a.$$

- The analysis would be qualitatively similar if this does not hold, but we would need more backward induction steps from the last period to get repuation effects work
- What does a weak entrant do? This depends on the probability with which incumbent fights
- Let p denote the current period belief of entrants about incumbent's type
- If $p > b/(1+b)$, then weak entrant should not enter
- This also means that a weak incumbent should fight (because this deters entry for the next period)
- What if $p < b/(1+b)$?
- Suppose $p < b(1-b)$, and consider incumbents equilibrium fighting probability π that leads to final period belief $p'(p,\pi)$
- Could we have π such that $p' < b/(1+b)$? No, because then next period weak entrant enters, so it would be better to accommodate in the current period than fight
- Could we have π such that $p' > b/(1+b)$? No, because then next period weak entrant does not enter, so it is better to fight than accommodate
- The only candidate for equilibrium is π that leads to $p' = b/(1+b)$:

$$\pi = p/((1-p)b)$$

- This also requires that in the final period weak entrant (who will then be indifferent between entry and not entering) enters with a probability that makes the incumbent indifferent
- The total probability that entry is fought in second last period is therefore

$$p \cdot 1 + (1-p) \cdot p/((1-p)b) = p(b+1)/b$$

  and therefore entrant stays out if $p > [b/(1+b)]^2$.
- Now, what happens in the third last period?
- Continuing with the same backward induction logic, if $p > [b/(1+b)]^2$, then weak entrant stays out and weak incumbent fights (to deter entry of next period weak incumbent)

- If $[b/(1+b)]^3 < p < [b/(1+b)]^2$, then weak entrant stays out and weak incumbent randomizes
- If $p < [b/(1+b)]^3$, then weak entrant enters and weak incumbent randomizes
- More generally, for the k:th period from the end, weak entrant stays out and weak incumbent fights if $p > [b/(1+b)]^{k-1}$
- Note that $[b/(1+b)]^{k-1}$ goes geometrically to zero as k increases
- Therefore, for a fixed prior probability p^0, there is some $\overline{k}$ such that the incumbent fights with probability one for the first $N - \overline{k}$ periods
- Hence, for a fixed p^0, as the total number of periods N increases, the total payoff of the incumbent converges to the payoff that it would obtain by committing to always fighting
- Posterior p stays constant at p^0 for the $N - \overline{k}$ first periods, so the incumbent does not make the entrants believe that it is a tough type
- Rather, reputational concerns make the incumbent *behave as if it was tough*
- This is the unique sequential equilibrium of the game (= PBE in this case)