10 Dynamic Games of Incomple Information

DRAW SELTEN'S HORSE

- In game above there are no subgames. Hence, the set of subgame perfect equilibria and the set of Nash equilibria coincide.
- However, the sort of logic that led us to consider subgame perfection also suggests that player 2 should play "down" if he/she gets the chance to play (2>1).
- That would break the Nash equilibrium indicated above.

10.1 Beliefs and Sequential Rationality

Definition 1 *Given an extensive form game K, a system of beliefs is a map $\mu : T \backslash Z \rightarrow [0,1]$ such that $\sum_{x \in h} \mu(x) = 1$ for every information set $h \in H$.*

Definition 2 *(σ, μ) is sequentially rational if for every i and every information et h that belongs to i, the continuation strategy σ is a best reply to the continuation startegies of the other players given beliefs μ. That is*

$$Eu_i(s_i|h, \sigma|h; \mu) \geq Eu_i(s_i|h, \widehat{\sigma}_i|h, \sigma_{-i}|h; \mu)$$

holds for every i, h and every alternative continuation strategy $\widehat{\sigma}_i|h$.

Definition 3 *A strategy profile σ is said to be sequentially rational if there exists μ such that the pair (σ, μ) is sequentially rational.*

Remark 1 *In games with perfect information σ is sequentially rational if and only if σ is a backwards induction equilibrium. This can be seen by noting that all information sets are singletons, so*

$$Eu_i\left(s_i|h,\sigma|h;\mu\right) \geq Eu_i\left(s_i|h,\widehat{\sigma}_i|h,\sigma_{-i}|h;\mu\right)$$

simply asks the player to optimize given the history in nodes preceeding the endnodes....and so on, meaning that σ will be a backwards induction equilirbium by construction.

Definition 4 (σ,μ) *is said to be a (qualifier missing) perfect Bayesian equilibrium if it is sequentially rational and if μ is obtained from σ by Bayes rule "whenever Bayes rule is applicable". [In MWG the qualifier in "weak" and "whenever Bayes rule is applicable" is taken to mean "on the equilibrium path".*

10.2 Example

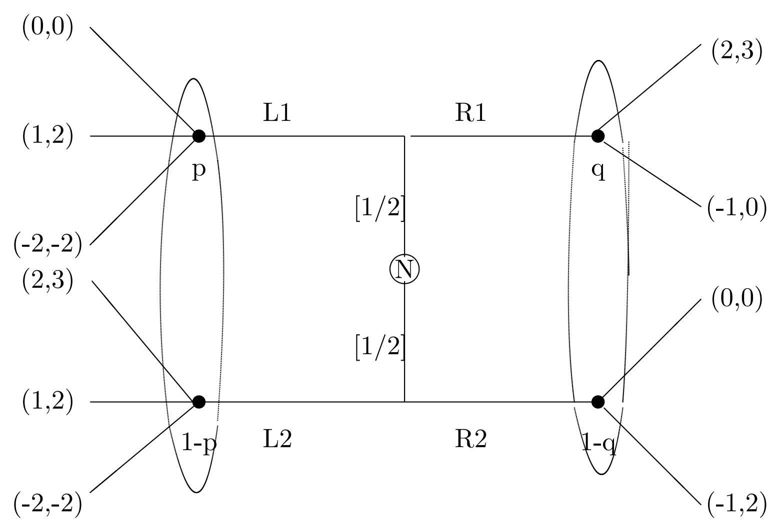

Figure 1: A Simple Signaling Game

- $(R_1L_2, aT), (p,q) = (0,1)$ is Perfect Bayesian

Bayes rule

$$\begin{aligned} p &= \Pr[t_1|L] = \frac{\Pr[L|t_1]\Pr[t_1]}{\Pr[L|t_1]\Pr[t_1] + \Pr[L|t_2]\Pr[t_2]} \\ &= \frac{0\frac{1}{2}}{0\frac{1}{2} + 1\frac{1}{2}} = 0 \\ q &= \Pr[t_1|R] = \frac{\Pr[R|t_1]\Pr[t_1]}{\Pr[R|t_1]\Pr[t_1] + \Pr[R|t_2]\Pr[t_2]} \\ &= \frac{1\frac{1}{2}}{1\frac{1}{2} + 0\frac{1}{2}} = 1 \end{aligned}$$

- CHECK! $(L_1L_2, bB); (p,q) = (\frac{1}{2}, q)$ is a PBE if $q \geq \frac{2}{5}$ (free to specify beliefs at unreached information sets, but must be done so that they support equilibrium strategies).
- CHECK! (R_1R_2, cT) is a Nash equilibrium, but not a PBE.

10.3 Example: How to Play Poker

Consider the following simplification of poker.

- There are two players, a professor and a student.
- First, each player first puts a dollar down (the ante). As we wont consider this to be a choice we will not make this part of the extensive form of the game (the reason for describing the setup this way is that it is a natural way of generating the payoffs).
- Then the professor draws a single card from a deck consisting of an equal number of queens and kings. Only the professor draws a card.
- After observing the card, the professor may either "bet" or "fold". If the professor folds, the student wins the pot containing the ante (implying that the student wins 1 and the professor loses 1). If the professor "bets", he places another dollar in the pot.
- The student then has the option between "fold" and "call". If the student folds, then the professor takes the pot (implying a net transfer of 1 from the student to the professor). If the student "calls" he adds another dollar to the pot. In this case the

professor wins if he has a king (in which case there is a net transfer of 2 from student to professor) and the student wins if the professor has a queen (in which case there is a net transfer of 2 from professor to student.

The extensive form is depicted in Figure **??**.

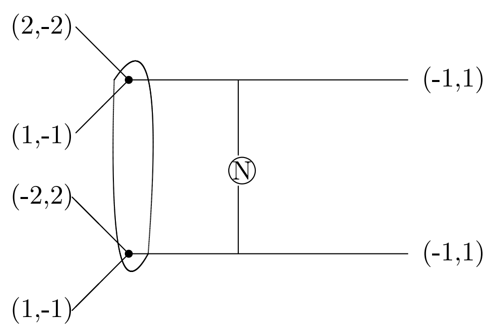

Figure 2: Simple Form of Poker

This game has a unique equilibrium, which is in mixed strategies. *Let* $q = \Pr[\text{bet}|\text{queen}]$ be the randomization probability for the professor. The student also needs to randomize and let p be the probability that the student calls. The probability that the professor has a king when betting is

$$\begin{aligned}\Pr[\text{king}|\text{bet}] &= \frac{\Pr[\text{bet}|\text{king}]\Pr[\text{king}]}{\Pr[\text{bet}|\text{king}]\Pr[\text{king}] + \Pr[\text{bet}|\text{queen}]\Pr[\text{queen}]} \\ &= \frac{1 \times \frac{1}{2}}{1 \times \frac{1}{2} + q \times \frac{1}{2}} = \frac{1}{1+q}\end{aligned}$$

The student is indifferent if

$$\begin{aligned}\frac{1}{1+q}(-2) + \left(1 - \frac{1}{1+q}\right)2 &= \frac{1}{1+q}(-1) + \left(1 - \frac{1}{1+q}\right)(-1) \\ &= -1\end{aligned}$$

Solving this we get $q = \frac{1}{3}$. To make the professor with a Queen indifferent between folding and calling it must be that

$$-1 = (1-p)\,1 - 2p,$$

which can be solved for $p = \frac{2}{3}$. I will not show it, but by succesive elimination of all remaining possibilities you'll find that this is the unique Nash equilibrium (as well as the unique perfect Bayesian equilibrium).

Notice that the game puts the player who can observe the card at an advantage: If the professor gets a king, the expected payoffs are

$$\begin{aligned} p2 + (1-p) &= \frac{2}{3}2 + \frac{1}{3} = \frac{5}{3} \text{ for the professor} \\ &\quad -\frac{5}{3} \text{ for the the student} \end{aligned}$$

If the professor gets a queen, the expected payoffs are

$$q(-1) + (1-q)(p(-2) + (1-p)1) = -\frac{2}{3} + \frac{1}{3}\left[\frac{2}{3}(-2) + \frac{1}{3}\right] = -1$$

for professor and 1 for student. Hence the expected payoff for the professor is

$$\frac{1}{2}\frac{5}{3} + \frac{1}{2}(-1) = \frac{1}{3}.$$

10.4 Defining "Signaling"

The concept of "signaling" refers to strategic models where one or more informed agents take some observable actions *before* one or more uninformed agents make their strategic decisions. This leads to situations where the uninformed agent care about the actions taken by the informed agent not only because the actions affect payoffs directly, but also because the action taken say something about the type of the player. This in turn creates incentives to select actions to send the right signal about type.

Formally,

- There are two players. We refer to player 1 as the *sender* and to player 2 as the *receiver.*

- Player 1 has private information about his type. We denote the type space by Θ and write θ for a generic element. The set of available actions for the sender is M, so a (pure) strategy is function s_1 with $s_1(\theta) \in M$ for every type θ. We let p denote the (prior) probability distribution over the senders type space.

- Player 2, the *receiver*, observes the action chosen by the sender and then take some action in A. A pure strategy for player is a function s_2, where $s_2(m) \in A$ for every $m \in M$

- Utility functions,

$$\begin{aligned} &u_1(a, m; \theta) \text{ for the sender} \\ &u_2(a, m, \theta) \text{ for the receiver,} \end{aligned}$$

10.5 Forward Induction in Signalling Games

Example

We note that (uB, b) is subgame perfect (and perfect Bayesian). However, many game theorists would argue that the outcome is implausible. The argument is that if player 2 would see that her information set is reached, then she couldn't possibly believe that player 1 is aiming at the $(1,5)$ outcome. Then, it would be better for player 1 to play "up".

The reduced normal form is

	t	b
u	$2,0$	$2,0$
AB	$0,0$	$1,5$
AT	$5,1$	$0,0$

and here we see that, indeed, AB is strictly dominated by up. Many "forward induction" ideas are based on the normal form.

Refinements by forward induction-like criteria have been used extensively in signalling and many concepts are developed specifically for signalling games. One of the simplest is:

Definition 5 *Consider a signalling game for which there exists a perfect Bayesian equilibrium with equilibrium payoff* $u_1^*(\theta)$ *for a sender of type* θ*. Then, we say that message* m *is equilibrium dominated for type* $\theta \in \Theta$ *if*

$$u_1^*(\theta) > \max_{a \in A} u_2(a, m, \theta)$$

Definition 6 *A Perfect Bayesian equilibrium passes the test of equilibrium domination if for all out of equilibrium messages* $m \in M$*, the reciever assigns probability 0 to* m *being sent by* θ *if* m *is equilibrium dominated for* θ *and if there exists some other type* θ' *such that* m *is not equilibrium dominated for* θ'.

EXAMPLE-TRUTH GAME

10.6 Sequental Equilibrium

10.6.1 Example-Weak Perfect Bayesian not "structurally consistent"

9.C.4 in MGW

10.6.2 Example–Weak Perfect Bayesian not subgame perfect

9.C.5 in MGW

10.6.3 Sequental Equilibrium

There are several attempts to "fix" the notion of Perfect Bayesian equilibrium. However, it turns out that a relatively satisfactory solution solves some of the issues by taking a quite different approach.

Definition 7 *(σ, μ) is said to be a sequential equilibrium if:*

1. *(σ, μ) is sequentially rational*

2. *There exists a sequence of mixed strategies with full support $\left\{\sigma^k\right\}_{k=1}^{\infty}$ with $\sigma^k \to \sigma$ such that $\mu = \lim_{k\to\infty} \mu_k$, where μ_k are the beliefs derived from σ^k using Bayes rule.*

   $9.C.4$

   $9.C.5$

10.7 A Simple Version of the Spence Model

Suppose:

- Types are given by $\theta \in \{1, 2\}$

- μ_0 is the probability that $\theta = 2$

- A Worker may choose (i.e. commit to) any education of length $t \geq 0$

- The utility is given by

$$u(w, t, \theta) = w - \frac{t}{\theta}$$

- Game: 1) Worker choose education. 2) Firms compete Bertrand for workers and get a profit

$$\pi = \theta - w$$

  if it manages to attract the worker.

Let $\mu(t)$ denote the probability that firms' asses that $\theta = 2$. In a sequential equilibrium beliefs must be consistent (and the same for both firms) and firms must both optimize given beliefs, implying that

$$w(t) = 2\mu(t) + (1 - \mu(t)) = 1 + \mu(t)$$

10.7.1 Separating Equilibria

Note that:

- $t(1) = 0$ in any separating sequential equilibrium (since $\mu(t(1)) = 0 \Rightarrow w(t(1)) = 1$ and $w(0) \geq 1$)

- Neither type can have an incentive to mimic the other

$$w(t(1)) - t(1) \geq w(t(2)) - t(2)$$
$$1 \geq 2 - t(2) \Leftrightarrow t(2) \geq 1$$

  and

$$w(t(2)) - \frac{t(2)}{2} \geq w(t(1)) - t(1)$$
$$2 - \frac{t(2)}{2} \geq 1 \Leftrightarrow t(2) \leq 2$$

By setting beliefs

$$\mu(t) = \begin{cases} 1 & \text{if } t = t(2) \\ 0 & \text{if } t \neq t(2) \end{cases} \quad \text{or } \mu(t) = \begin{cases} 1 & \text{if } t \geq t(2) \\ 0 & \text{if } t < t(2) \end{cases}$$

any $t(2) \in [1,2]$ can be supported as a sequential equilibrium.

10.7.2 Pooling Equilibria

To support as large a set of education levels as pooling equilibria, suppose that $t(1) = t(2) = t^*$ is a pooling equilibrium and let beliefs be

$$\mu(t) = \begin{cases} \mu_0 & \text{if } t = t^* \\ 0 & \text{if } t \neq t^* \end{cases},$$

which obviously is consistent with Bayes rule where relevant. The best deviation for both types is then to $t = 0$ and this is not profitable for the low productivity/high cost of education type if

$$\mu_0 - t^* \geq 1.$$

Clearly, the high productivity type has no incentive to deviate if the low [productivity type has no incentive to deviate, so any

$$t^* \in [0, 1 - \mu_0]$$

can be supported as a pooling equilibrium.