# Economics and Computation

ECON 425/563 and CPSC 455/555

Professor Dirk Bergemann and Professor Joan Feigenbaum

Lecture II

In case of any questions and/or remarks on these lecture notes, please contact Oliver Bunn at oliver.bunn(at)yale.edu.

# 1 Games with Complete Information

Last lecture we have been considering the following general setup:

- A set of players $\mathcal{I} = \{1, ..., I\}$, where one particular player was denoted by $i \in \mathcal{I}$.
- A set of actions $A_i$ for each player $i \in \mathcal{I}$, where $A_i = \{a_i^1, ..., a_i^K\}$. One particular action for player $i$ is denoted by $a_i \in A_i$.
- A payoff function $u_i$ for each player $i \in \mathcal{I}$ that maps a tuple of actions, one by each player, into the real numbers, i.e. $u_i(a_1, ..., a_I) \in \mathbb{R}$.

## 1.1 Iterated Elimination of Weakly Dominated Strategies

Now, the general setting will be extended to incorporate uncountable and infinite action-sets, in particular $\mathbb{R}_+$. In the following, consider the **Cournot duopoly**-game:
Two firms need to make a decision about the non-negative quantity of an identical good that they want to supply to the market.[1] Denote the firm's quantities by $q_i$, for $i \in \{1, 2\}$. The firms make their decision simultaneously and, according to their decision, the price of the good is determined as

$$p = 1 - q_1 - q_2.$$

Clearly, this pricing-function exhibits the intuitive property that the price decreases in the quantity supplied. As an example for this situation, one may think of both firms as steel-producers who supply into one steel-market.
The firm's cost of producing are set to zero, i.e. one firms profit function is given by

$$\begin{aligned} \Pi_i(q_i, q_j) &= pq_i \\ &= (1 - q_i - q_j)q_i. \end{aligned}$$

Here, $i \in \{1, 2\}$ and $j \neq i$.
In order to determine the optimal quantity that a firm is supposed to supply, one can optimize $\Pi_i$ with respect to $q_i$. This yields the following FOC:

$$\begin{aligned} & \frac{\partial \Pi_i(q_i, q_j)}{\partial q_i} \overset{!}{=} 0 \\ \Leftrightarrow \quad & 1 - 2q_i - q_j = 0 \\ \Leftrightarrow \quad & q_i = \frac{1 - q_j}{2}. \end{aligned}$$

Explicitly, one has computed the following result:

[1] The name "duopoly" is derived from the situation that two firms are in a "kind of monopolistic situation".

- In response to firm 2 setting a quantity $q_2$, firm 1 optimally sets

$$q_1^*(q_2) = \frac{1 - q_2}{2}. \quad (1)$$

- In response to firm 1 setting a quantity $q_1$, firm 2 optimally sets

$$q_2^*(q_1) = \frac{1 - q_1}{2}. \quad (2)$$

Now, the Nash-equilibrium of the game can be determined as the solution to the system of equations given by (1) and (2). Hence

$$q_i^* = \frac{1}{3}, q_j^* = \frac{1}{3}. \quad (3)$$

Graphically, the Nash-equilibrium can be determined as the intersection of the two lines that are described by (1) and (2). As can be expected, the intersection is located at the point $(\frac{1}{3}, \frac{1}{3})$.

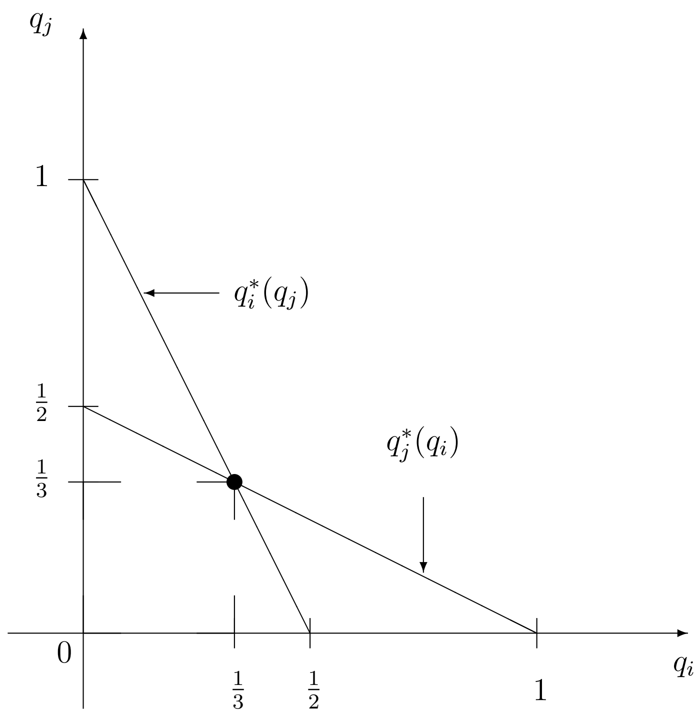

In contrast to many situations in Computer Science, the solution of this game does not try to minimize costs in any way. Quite the contrary, the firms' sole objective is to maximize their profits.

In the following, an alternative derivation of the solution to the *Cournot Duopoly* will be presented. This derivation is based on an iterative procedure that subsequently eliminates weakly dominated strategies. Heavy use will be made of the fact that any player's optimal response is monotonely decreasing in the quantity that the other player supplies. Given this property, one obtains boundaries on reasonable quantity-choices given that the opponent will only choose from a specific interval.
First of all, the fact that the optimal response functions $q_1^*(q_2) = \frac{1-q_2}{2}$ and $q_2^*(q_1) = \frac{1-q_1}{2}$ depend on the quantity supplied by the other firm rules out the existence of a dominant strategy. So, the question arises whether one can possible rule out any weakly dominated strategies. This will be done in the following iterative manner:

1. Define the sets
$$S_1^0 := \mathbb{R}_+, S_2^0 := \mathbb{R}_+.$$
That is, both $S_1^0$ as well as $S_2^0$ are set to the player's full action-set.

2. For $q_2 = 0$, the best-response functions (1) prescribes that player 1 will optimally never choose to supply more than $\frac{1}{2}$. Therefore, all quantities in the set $S_1^0$ that are located above $\frac{1}{2}$ are weakly dominated for player 1. The same argument applies to player 2 facing of player 1 choosing $q_1 = 0$. Hence, one can construct sets $S_1^1$ and $S_2^1$ from $S_1^1$ as well as $S_2^1$ by eliminating weakly dominated strategies:
$$S_1^1 := \left[0, \frac{1}{2}\right], S_2^1 := \left[0, \frac{1}{2}\right].$$

3. Now, by the methodology of the iteration, player 2's actions are now restricted to $S_2^1$. Because player 2 can never choose a quantity bigger than $\frac{1}{2}$, it follows from player 1's optimal response function via
$$q_1^*\left(\frac{1}{2}\right) = \frac{1}{4},$$
that player 1 will never choose a quantity smaller than $\frac{1}{4}$. That is, the interval $[0, \frac{1}{4})$ contains weakly dominated actions for player 1. The identical argument applies to player 2 facing player 1 choosing from the set $S_1^1$. Hence, the elimination of weakly dominated strategies implies
$$S_1^2 := \left[\frac{1}{4}, \frac{1}{2}\right], S_2^2 := \left[\frac{1}{4}, \frac{1}{2}\right].$$

4. Player 2's actions are now restricted to $S_2^2$. Because player 2 can never choose a quantity smaller than $\frac{1}{4}$, it follows from player 1's optimal response function via

$$q_1^*\left(\frac{1}{4}\right) = \frac{3}{8},$$

that player 1 will never choose a quantity bigger than $\frac{3}{8}$. That is, the interval $(\frac{3}{8}, \frac{1}{2}]$ contains weakly dominated actions for player 1. The identical argument applies to player 2 facing player 1 choosing from the set $S_1^2$. Hence, the elimination of weakly dominated strategies implies

$$S_1^3 := \left[\frac{1}{4}, \frac{3}{8}\right], S_2^3 := \left[\frac{1}{4}, \frac{3}{8}\right].$$

5. $$S_1^4 := \left[\frac{5}{16}, \frac{3}{8}\right], S_2^4 := \left[\frac{5}{16}, \frac{3}{8}\right].$$

6. $$S_1^5 := \left[\frac{5}{16}, \frac{11}{32}\right], S_2^5 := \left[\frac{5}{16}, \frac{11}{32}\right].$$

7. etc.

The described procedure has an algorithmic feature that appears easily implementable. The only assumption that one needs to impose in order to get the iteration going deals with the rationality of the players. Since the iteration involves repetitive reasoning it is required that

- Both players are rational.
- Both players know that their counterpart is rational.
- Both players know that their counterpart knows that they are rational.
- etc.

Outcomes that survive this iterated elimination of weakly dominated strategies are called **rationalizable**. In the particular situation of the *Cournot duopoly*, the rationalizable outcome coincides with the Nash-equilibrium.
In general, the set of rationalizable outcomes is a superset of the set of Nash-equilibria. For example, in the case of the *Battle of the Sexes* there is no outcome that can be eliminated by the iterative procedure outlined above. That is, all outcomes of the game are rationalizable, whereas there are only three Nash-equilibria. The iterated removal of weakly dominated strategies imposes very strong rationality-assumptions on the players

which make the elimination of outcomes more difficult than under the procedure to find Nash-equilibria.

The reason why the iterated elimination of weakly dominated strategies has worked so well for the situation of the *Cournot Duopoly* has to do with the mutual relationship between the quantities in the best-response functions (1) and (2). Explicitly, $q_1$ and $q_2$ are strategic substitutes because one player's optimal response function is decreasing in the other player's quantity.

# 2 Games with Incomplete Information

Instead of games with complete information, the focus will now be shifted towards games of incomplete information, i.e. games where some players are uncertain about the payoffs (own or others), the strategies or the players in the game. This will be made explicit by the consideration of so-called **types**. Instead of the previous specification of a utility function as a mapping $u_i(a_1, ..., a_I)$ for every player $i \in \mathcal{I}$, we will specify

$$u_i(a_1, ..., a_I; \underbrace{t_1, ..., t_I}_{\text{types}}).$$

Hereby, the type of player $i$ is denoted by $t_i$. $t_i$ captures private information that is only available to player $i$. As an example, $t_i$ may capture the valuation that a certain bidder in an auction has for a specific painting. Before any action is taken in a game, no player besides player $i$ knows the specific value of $t_i$, but there will be a general assessment about the distribution of types in the population. But in the course of the game it may well be possible that players reveal (part of) their private information to other players by the way in which they act. To introduce private information of players appears to be highly reasonable if one thinks about situations like the Internet as a place with highly dispersed information-structure.

The new notion of a **Bayesian game** that captures incomplete information consists of the following elements:

- A set of players $\mathcal{I} = \{1, ..., I\}$.
- A set of actions $A_i$ for every $i \in \mathcal{I}$. A particular action for player $i$ is denoted by $a_i \in A_i$.
  The set of all tuples of actions that may arise in the game is denoted by $A = A_1 \times ... \times A_I$.

- A set of types $T_i$ for every $i \in \mathcal{I}$. A particular type for player $i$ is denoted by $t_i \in T_i$. The particular type for player $i$ is this player's private information.
  The set of all tuples of actions that may arise in the game is denoted by $T = T_1 \times ... \times T_I$.
- A common prior, also called common prior distribution, denoted by

$$\begin{aligned} p : T \quad &\rightarrow \quad [0, 1], \\ \text{such that} \quad &\sum_{t \in T} p(t) = 1. \end{aligned}$$

- Payoff-functions

**Definition 1** *A* ***Bayesian game*** $\Gamma_B$ *is given by*

$$\Gamma_B := \left\{ I, \{A_i\}_{i=1}^I, \{T_i\}_{i=1}^I, p, \{u_i\}_{i=1}^I \right\}. \tag{4}$$

**Remark 1** *The two additional elements that distinguish a Bayesian game from the setting previously considered are the type-spaces* $\{T_i\}_{i=1}^I$ *and the common prior* $p$*.*

Starting from the common prior, every player can make use of his private information $t_i \in T_i$ to update $p$ via Bayes' rule to obtain

$$p(t_i | t_i),$$

i.e. player $i$ can form a conditional belief[2] on the distribution of the other players' type given what he knows about himself. As an example for this reasoning from the prior to the conditional belief consider the following distribution of $p$:

**Player b**

| | | $t_b^l$ | $t_b^h$ |
|---|---|---|---|
| **Player a** | $t_a^l$ | $\frac{1}{4}$ | $\frac{1}{4}$ |
| | $t_a^h$ | $\frac{1}{4}$ | $\frac{1}{4}$ |

Here, the game consists of two players, player a and player b. Each of the players has two possible types, a low type and a high type, i.e.

$$T \in \{t_a^l, t_a^h\} \times \{t_b^l, t_b^h\}.$$

[2]This is sometimes referred to as the posterior belief.

The numbers in the fields of the matrix stand for the probability by which the corresponding pair of types appears according to the common prior. In the example above, each of the pair of types $(t_a^l, t_b^l)$, $(t_a^l, t_b^h)$, $(t_a^h, t_b^l)$ and $(t_a^h, t_b^h)$ comes up with probability $\frac{1}{4}$. Unfortunately, this structure of the probability distribution does not allow any inference at all, because the type-distributions of the two players are independent from each other. This can be seen from the fact that the probability distribution in the matrix can be represented as the product of two marginal distributions. Explicitly, if one assumes that each player's distribution over just his types is given by $(\frac{1}{2}, \frac{1}{2})$ and both type-distributions are independent, then one would obtain the above matrix-representation via

$$\begin{aligned}
p((t_a^l, t_b^l)) &\overset{\text{Indep.}}{=} \text{Prob(type a is } t_a^l) \cdot \text{Prob(type b is } t_b^l) = \frac{1}{2} \cdot \frac{1}{2} = \frac{1}{4}. \\
p((t_a^l, t_b^h)) &\overset{\text{Indep.}}{=} \text{Prob(type a is } t_a^l) \cdot \text{Prob(type b is } t_b^h) = \frac{1}{2} \cdot \frac{1}{2} = \frac{1}{4}. \\
&\textit{etc.}
\end{aligned}$$

An example where the conditioning role of private information manifests itself is given by the following common prior distribution:

**Player b**

| **Player a** | | $t_b^l$ | $t_b^h$ |
|---|---|---|---|
| | $t_a^l$ | $\frac{1}{3}$ | $\frac{1}{6}$ |
| | $t_a^h$ | $\frac{1}{6}$ | $\frac{1}{3}$ |

Obviously, the previous argument about not being able to draw any inference does not apply anymore. So, it makes sense to compute

$$p(t_b^l | t_a^l).$$

This is done by Bayes' rule as follows:

$$\begin{aligned}
p(t_b^l | t_a^l) &= \frac{p((t_a^l, t_b^l))}{p((t_a^l, t_b^l)) + p((t_a^l, t_b^h))} \\
&= \frac{\frac{1}{3}}{\frac{1}{3} + \frac{1}{6}} \\
&= \frac{2}{3}.
\end{aligned}$$

Similarly, one can compute the following conditional probability

$$\begin{aligned} p(t_a^l|t_b^h) &= \frac{p((t_a^l, t_b^h))}{p((t_a^l, t_b^h)) + p((t_a^h, t_b^h))} \\ &= \frac{\frac{1}{6}}{\frac{1}{6} + \frac{1}{3}} \\ &= \frac{1}{3}. \end{aligned}$$

A simple example of a Bayesian game is given by:

**Type $t_2^1$ of Player 2**

| Player 1 | | Left | Right |
|---|---|---|---|
| | Up | 1,1 | 0,0 |
| | Down | 0,1 | 1,0 |

**Type $t_2^2$ of Player 2**

| Player 1 | | Left | Right |
|---|---|---|---|
| | Up | 1,0 | 0,1 |
| | Down | 0,0 | 1,1 |

Player 1 is choosing between "Up" and "Down" without knowing whether the left or the right matrix is played. Before making any decision, player 2 observes whether she has type $t_2^1$ or $t_2^2$. According to her type, player 2 knows whether the left or the right matrix is played. Then, she makes a decision between "Left" and "Right". Player 2's type is the uncertain variable in this Bayesian game. Player 2 knows whether he has type $t_2^1$ or type $t_2^2$, but from the point of view of player 1 this type is uncertain. So, there is a common prior on player 2's type given by[3]

$$p(t_2^1) = \alpha, p(t_2^2) = 1 - \alpha.$$

Observe the following characteristics of the game:

- Player 1 obtains the same payoff in both matrices.
- For player 2, action $L$ is a dominant strategy if she has type $t_2^1$, whereas $R$ is a dominant strategy if she has type $t_2^2$.

Before we can solve for the equilibrium of the above game, we need to formally define the notion of a **strategy** in a Bayesian game:

[3]Because there is no uncertainty about player 1's type, his type will be omitted as an argument of the common prior $p$ for notational convenience.

**Definition 2** *A **pure strategy** for player $i \in \mathcal{I}$ in a Bayesian game is given by*

$$s_i : T_i \to A_i.$$

*A **mixed strategy** for player $i \in \mathcal{I}$ in a Bayesian game is given by*

$$s_i : T_i \to \Delta(A_i).$$

So, a strategy is always a mapping from the player's type-set into the action-set. That is, for player $i$, the mapping $s_i$ assigns an action $a_i$ (or a probability distribution over actions) to any type-realization $t_i$. This definition incorporates the idea of a contingent plan. It can be seen as a list of instructions of how to behave in response to any possible type-realization.
If there is only one type - a situation that we have faced in the games of complete information analyzed before - then the above notion of a strategy simply prescribes one action or one probability distribution over actions. This is consistent with the way in which we have used the notion of a strategy before.

**Remark 2** *The question may arise why one needs to specify the whole contingent plan of an player. If player $i$ already knows his type, then she can also determine her action according to this one type and she may not see any need to determine the full strategy-mapping beyond this one type-action-pair. But notice that, as the other players are still uncertain about the specific player's type, they will need to determine a probability distribution over player $i$'s actions, depending on her type. So, the specification of a whole function is important for the definition of every player's strategy in a consistent way.*

Now, we want to come back to the initially studied Bayesian game. The solution-concept to be used is that of a Bayesian Nash equilibrium, which is a simple generalization of the idea of mutual best responses to the situation of incomplete information. For the analysis, the probability $\alpha \in [0, 1]$ will be fixed. As already outlined in Remark 4 of Lecture I, an Expected-Utility model will be used in this uncertain environment:

Player 2's best-response-strategy is given by

$$s_2^*(t_2) = \begin{cases} L & \text{if } t_2 = t_2^1 \\ R & \text{if } t_2 = t_2^2 \end{cases} .$$

Player 2's best response is irrespective of any action of player 1 because player 2's action is a dominant strategy for any of his type. In the left matrix, player 2 always wants to choose "Left" and in the right matrix it is "Right".
Player 1's best response to player 2's action will now depend on the probability $\alpha$. Player 1 only has one type, so her strategy will simply be one action (or one probability distribution

to choose). She already knows that she will face the choice "Left" in the left matrix and "Right" in the right matrix. Now, the decision for any alternative depends on which matrix will actually be played, i.e. on the probability $\alpha$. Player 1 compares the following two payoffs:

$$\text{for "Up"} \quad : \quad \alpha \cdot 1 + (1-\alpha) \cdot 0, \tag{5}$$

$$\text{for "Down"} \quad : \quad \alpha \cdot 0 + (1-\alpha) \cdot 1. \tag{6}$$

The logic of the computation is as follows:
From the point of view of player 1, the left matrix is played with probability $\alpha$. Here, player 2 will choose "Left". So, "Up" yields a payoff of 1 and "Down" yields 0. The right matrix is chosen with probability $1-\alpha$. Here, player 2 will choose "Right". So, "Up" yields a payoff of 0 and "Down" yields 1. As one can see from (6) and (5), the choice "Up" is favorable for $\alpha > \frac{1}{2}$ and the choice "Down" is favorable for $\alpha < \frac{1}{2}$. For $\alpha = \frac{1}{2}$, both choices yields the same payoff, so any convex combination between "Up" and "Down" yields an identical payoff. Summarizing, player 1's best response is given by

$$s_1^* = \begin{cases} U & \text{if } \alpha > \frac{1}{2}, \\ \lambda U + (1-\lambda) D \text{ for any } \lambda \in [0,1] & \text{if } \alpha = \frac{1}{2}, \\ D & \text{if } \alpha < \frac{1}{2}. \end{cases}$$

A Bayesian Nash-equilibrium of the above game is now given by the pair $(s_1^*, s_2^*(s_1^*))$.