Games and Information

MEPS Course

Summer Semester 2024

Karl-Josef Koch

School of Economic Disciplines,
University of Siegen

Contents

1 Motivation

Glorious days of information economics



George A. Akerlof, A. Michael Spence, Joseph E. Stiglitz,Nobel-prize winners 2001

Glorious days of contract theory


Oliver Hart, Bengt Holmström,Nobel-prize winners 2016

2 Background

The world is not always what it seems at first glance.

There are hidden secrets.

In this course you will learn to understand the difference between

Moral Hazard

Adverse Selection

Signaling

The Value of Information

2.1 Some Examples

- The goal of this course is the provision of economic analysis of such problems.
- To accomplish this goal we start with the analysis of some examples.

Robinson Crusoe's problem to sell a goat

- Robinson Crusoe owns a goat yielding s liters of milk every day. He can imagine to sell the goat at a reasonable price.
- Friday thinks of buying the goat. He does not know whether the goat is milking well or not.
- $$U_R = \begin{cases} 2s, \text{ if he keeps the goat} \\ p, \text{ if he sells the goat} \end{cases}$$
- $$U_F = \begin{cases} 0, \text{ if does not buy the goat} \\ 3s - p, \text{ if he buys the goat} \end{cases}$$

- Friday knows that some goats milk well and others don't.
- He knows the distribution of the amount of milk s a goat usually yields. s is equally distributed on the interval $[0, 10]$.
- Friday is risk neutral. He uses the expected value of s^e as an estimator of the true yield of milk.
- A priori the estimator is $s^e = 5$.
- Robinson Crusoe is willing to offer the goat at a fair price p_{fair}.
- Friday concludes, that $p_{\text{fair}} \geq 2s$, in other words $0 \leq s \leq p_{\text{fair}}/2$.
- Hence Friday alines his estimator $s^e(p_{\text{fair}}) = p_{\text{fair}}/4$.
- He will accept Crusoe's offer, if $3s^e(p_{\text{fair}}) \geq p_{\text{fair}}$.
- However, $3s^e(p_{\text{fair}}) = 3\,p_{\text{fair}}/4 < p_{\text{fair}}$!
- Whatever Crusoe may regard as a fair price, Friday will (mis-)interpret his offer and reject it.

Does the result hinge on the parameters of the example?

One can set up a slightly more general model of the problem. For that, consider the following matrix of outcomes:

$$\text{Robinson's outcome} \quad \begin{cases} p & \textit{trade} \\ u_R(s) & \textit{no trade} \end{cases}$$

$$\text{Friday's outcome} \quad \begin{cases} u_F(s) - p & \textit{trade} \\ 0 & \textit{no trade} \end{cases}$$

- $u_R(s)$ and $u_F(s)$ are the respective values of milk for Robinson and Friday.
- Robinson knows the exact value of $u_R(s)$, but Friday has to use an estimate $u_F(s^e)$.
- If Robinson is willing to sell his goat at price p_{fair}, then $p_{\text{fair}} \geq u_R(s)$, and Friday concludes $s \leq u_R^{-1}(p_{\text{fair}})$.
- Friday's estimate will be $s^e = u_R^{-1}(p_{\text{fair}})/2$.
- On the other hand Friday accepts p_{fair} if $u_F(s^e) \geq p_{\text{fair}}$.
- So p_{fair} has to satisfy $u_F(u_R^{-1}(p_{\text{fair}})/2) \geq p_{\text{fair}}$.
- One further transformation yields $u_R^{-1}(p_{\text{fair}})/2 \geq u_F^{-1}(p_{\text{fair}})$.
- In other words, Friday will accept a offer if the gain of trade outweighs the risk of overestimation of s.

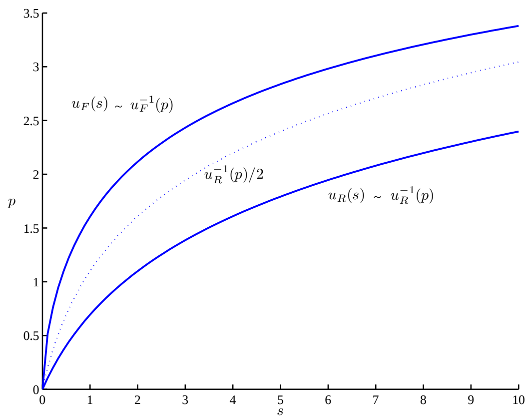

$u_R^{-1}(p_{\text{fair}})/2 \geq u_F^{-1}(p_{\text{fair}})$

- If u_F is not much larger than u_R, the risk of buying a low quality spoils the deal.

Is this a real problem?

You may be aware of the fact, that a while ago large German car producers ran into problems because for years their management had been cheating on the declaration of pollutant emission of their diesel vehicles. The public discussion of the problem caused a (slight) disarrangement of car markets world wide.

A car market obstacle

In a radio broadcast I heard the following statement:

Obviously demand will decrease as prices go down!

One may be puzzled because usually we expect demand to decrease with prices.

- The broadcast statement is suggesting the opposite
- However, here the question arises what the true value of a car may be?
- Even if the producer of the car may have to pay for part of the extra costs to solve the problem with the car, the resale value will probably decline.
- Current prices for new cars of this brand are likely to decrease because buyers are expecting inconveniences.

- The lower the current price of the car the more the potential buyers will expect the resale price to drop.
- Even worse: Potential buyers may be afraid that the current seller wants to get rid of the car because he is afraid the problem may be more severe than admitted by the seller.
- Hence, at least some buyers may refrain from buying this brand. And its is unlikely that you find new potential buyers.
- This phenomenon is called *adverse selection*.
- Can you build a little model to check whether the argument goes through?

Auctions
Now, Morteza wants to buy a very unique mug with the Eiffel Tower painted on it from a special souvenir shop. The shop owner is aware of his position as a monopolist, but has no idea about Morteza's willingness to pay and that of other potential clients. The number of interested customers is not very large, but too large to negotiate with each customer individually.

Therefore, he decides to auction off this special mug.

There are many forms of auctions. Here is the short list of the basic ones

- English Auction: Bidding starts at a low price and is raised incrementally as progressively higher bids are solicited, until no higher bids are received.

- Dutch Auction: Bidding starts with a high price and is decreased step by step by the auctioneer until the bid is accepted by a buyer.

What form of auction should the seller choose in order to maximize his revenue?

Assume that each bidder *privately* and *independently* forms an opinion of the value of the mug. Consider at first the case of an English Auction.

- Morteza and other bidders continue to participate until the price reaches their own private values.
- The auction stops when the bidder with the second-highest value drops out.
- Therefore, the seller's expected price is the expected value of the second-highest private value (plus the marginal increment of the last bidder).

Now consider the Dutch Auction.

- Morteza and his competitors plan to call out when the price has fallen slightly below their private valuations.
- The seller's expected price is the expected value of the highest private value minus the incremental amount by which this bidder allows the price to drop below his private value.
- Again, the seller's expected price turns out to be close to the expected the second-highest private value.

Which form of auctioning should the shop owner choose?

- In the case of an independent private value model, it doesn't matter!
- The result is known as the *revenue equivalence theorem*[1].

3 Basics of Game Theory

3.1 Normal Form Games

Definition of Games[2]

- What do we need to describe games?
    1. Actors (**players**)
    2. Set of possible **actions**

[1] A formal statement and proof needs a more elaborate specification of the framework.

[2] In this section we only recall fundamentals and terminology. You are supposed to know the basic concepts and you should be able to apply them.

3. **Rules** of the game
4. **Strategy** matters (at least in more sophisticated games)
5. Outcome with results for each single actor (**payoffs**)
6. Social interaction: the outcome of each player depends on the actions/strategy chosen by the others (**strategic interdependence**)
7. More or less **information** about all these aspects

- So a game is a rule-governed situation with a well-defined outcome, characterized by **strategic interdependence**.

Setting up a game takes several very different steps

- Choose a topic
- Name the players and define their goals
- List the actions they can take
- Describe the sequential structure
- Define the information sets and think about uncertainties
- Declare the payoffs

Last not least you arrange all this in the appropriate form. At the core of a game you now have

- Action sets, strategy sets, a dynamic structure, an information structure and probabilities or beliefs

The goal of the approach of game theory is to make reasonable statements about

- Strategy profiles or strategy-belief profiles
- The development of the game
- The outcome of the game and the payoffs

At any instant of a game, i.e. whenever a player is called up to decide on or to make a move, she picks an action out of the set of actions available.

Definition
A **strategy** s_i of player i is a particular choice of actions out of the set of available actions.[3]

The **strategy space** S_i of player i is the set of all strategies available to her.

[3]In sequential games actions have to be taken ät each instant of the game ∵ We will extend the definition later

A **strategy profile** $s = (s_1, \ldots, s_n)$ is a list of strategies for all players.

Heads up!
In simple games a player may be called up to make a move only once. In particular all players may be called up to decide on their strategy simultaneously at the beginning of the game. In this context, choosing a strategy is equivalent to choosing an action!

In more complex games, players may have to make their moves sequentially and when they exercise their move, they may have more or less information about their situation. We will have to update our definition for such games.

An equilibrium of a game is a strategy profile selected by some rule out of the set $\mathcal{E}$ of all strategy profiles .

In other words: An equilibrium picks a strtagey for each player. At the moment the question which pair is chosen is still open.

An **equilibrium concept** *or* **solution concept** *is a rule that defines a single equilibrium or a several equilibria based on the possible strategy profiles and the payoff functions.*[4]

$$F : \{S_1, \ldots, S_n, \pi_1, \ldots, \pi_n\} \longrightarrow \mathcal{E}$$

Player i's **best response** *or* **best reply** *to the strategies s_{-i} is the strategy s_i that yields him the greatest payoff given the strategy profile s_{-i} of the other players*

$$\pi_i(s_i, s_{-i}) \geq \pi_i(\tilde{s}_i, s_{-i}) \quad \forall \tilde{s}_i$$

This rather formal definition has a simple intuition: Given the choice of strategy of all other players, player i has no advantage from deviating from a *b*est response.

Detailed notation makes the dependency on the other players' strategies explicit

$$s_i(s_{-i})$$

The best response is strongly best if no other strategies are equally good, and weakly best otherwise.

Nash Equilibrium
The strategy profile s^* is a **Nash equilibrium** *if no player has an incentive to deviate from his strategy given that the other players do not deviate. Formally,*

$$\pi_i(s_i^*, s_{-i}^*) \geq \pi_i(\tilde{s}_i, s_{-i}^*), \;\; \forall i, \;\; \tilde{s}_i.$$

Heads up!
Make sure you are aware of a few details

[4] Notice that we use the word *rule* which is somehow informal. We avoid to go one step further and define a mapping from the set of games to the set of sets of profiles.

- A *Nash equilibrium* is a strategy profile, i.e. it is defined in terms of strategies, not in terms of actions nor in terms of payoffs! However, an equilibrium strategy profile generates *equilibrium payoffs.*
- The definition is a rule selecting one or perhaps more than one strategy profile. However, it may happen that no profile satisfies the equilibrium conditions.
- Nash equilibrium is not derived from maximizing behavior of players who only observe their environment! Nevertheless, it has a strong, deep rationale. A player maximizes given the other maximize as well.

Normal Form Games I

Player I is infectious carrying a dangerous virus. Player S is susceptible, she could become infected. They live at different places. For the weekend they had planned to go hiking together. Otherwise they could stay at home and study.

		$\mathcal{I}$	
		go hiking	*study*
$\mathcal{S}$	*go hiking*	(-300, 200)	(100, 50)
	study	(50, 100)	(50, 50)

- Analyze the game!

Solution:

- First assume $\mathcal{S}$ chooses *go hiking.* The best answer of $\mathcal{I}$ is *go hiking* as well: $s_{\mathcal{I}}(\textit{go hiking}) = \textit{go hiking}$.
- But $s_{\mathcal{S}}(\textit{go hiking}) = \textit{study}$. Hence, there is no Nash-equilibrium where $\mathcal{S}$ chooses *go hiking.*
- Now assume $\mathcal{S}$ choses *study.* The best answer of $\mathcal{I}$ again is *go hiking*: $s_{\mathcal{I}}(\textit{study}) = \textit{go hiking}$.
- And $s_{\mathcal{S}}(\textit{go hiking}) = \textit{study}$. Hence, the strategy profile (*study, go hiking*) forms the Nash-equilibrium of the game. The profile is a pair of mutually best responses.

There is an easier way to find this Nash-equilibrium:

- For player $\mathcal{I}$ the strategy *hike* always is the best answer! It is a dominant strategy.
- Seeing this $\mathcal{S}$ has to find his best answer to this dominant strategy only.
- Of course we get the same equilibrium (*study, go hiking*).

Normal Form Games II

Consider the game with a different pattern of outcomes:

		$\mathcal{I}$	
		go hiking	*study*
$\mathcal{S}$	*go hiking*	(-300, 200)	(30, 50)
	study	(50, 30)	(50, 50)

- Analyze the game!
- Compare to the former version!

Normal Form Games III
Consider the game with a third pattern of outcomes:

		$\mathcal{I}$	
		go hiking	*study*
$\mathcal{S}$	*go hiking*	(-300, -50)	(100, 40)
	study	(40, 100)	(50, 50)

- Analyze the game and compare to the former versions!

3.2 Extensive Form Games

Extensive form: Game tree
An elegant, common model to organize the different elements defining a sequential game is a game tree.[5]

- *A* **node** *is a point in the game at which some player or Nature takes an action, or the game ends.*
- *A* **successor** *to node X is a node that may occur later in the game if X has been reached.*
- *A* **predecessor** *to node X is a node that must be reached before X can be reached.*
- *A* **starting node** *is a node with no predecessors.*
- *An* **end node** *or* **end point** *is a node with no successors.*

[5] This section is closely related to Rasmusen, chapter 2.

- *A* **branch** *is associated with a single action in a player's action set at a particular node.*
- *A sequence of nodes and branches leading from the starting node to an end node is called a* **path**.[6]
- *A* **game tree** *is a configuration of nodes and branches running without any closed loops from the starting node to its end nodes.*

Subgames
Any single node apart from end notes defines a subgame starting from this node.

Notation

- *Nodes are labeled with the name of the player. If a player is supposed to move several times we may add and index to the label*
- *Branches are labeled with the name of the action.*
- *End nodes are labeled with a list of* **payoffs** *for all players.*

At any node of a game some player is called up to make a move. She picks an action out of the set of actions available at this node.

Definition
A **strategy** s_i of player i is a list of actions out of the set of the actions available at the respective node of the game where she is supposed to make a move.

The **strategy space** S_i of player i is the set of all strategies available to her.

A **strategy profile** $s = (s_1, \dots, s_n)$ is a list of strategies for all players.

Heads up!
Even in simple extensive form games a player may be called up to make a move only once. In such a case, choosing a strategy is equivalent to choosing an action!

However, in more complex games, players may have to make their moves sequentially, and may move several times. In such case a notation is required which refers to the nodes considered.

Notation
A node is identified by its label or through the path leading to it We propose several possible, equivalent notations.

Assume it is player i's turn to move at nodes $I_1, I_2, I_3, \dots$. The path from the starting node to any node I_j is $\mathcal{P}_j$

[6] A connected segment of a path from some initial node to some final node is called a path as well.

- A strategy of player i is a sequence of available actions at each node I_j: $s_i = (a_i(I_1), a_i(I_2), a_i(I_3), \dots)$.
- Equivalently we can associate each node I_j with the path $\mathcal{P}_j$ leading to the node and write: $s_i = (a_i(\mathcal{P}_1), a_i(\mathcal{P}_2), a_i(\mathcal{P}_3), \dots)$.
- A very short notation obviously is $s_i = (a_{i,1}, a_{i,2}, a_{i,3}, \dots)$ whereby $a_{i,j}$ is associated with I_j or equivalently with $\mathcal{P}_j$.
- We are free to use any convenient notation.

Extensive form of the entry deterrence game

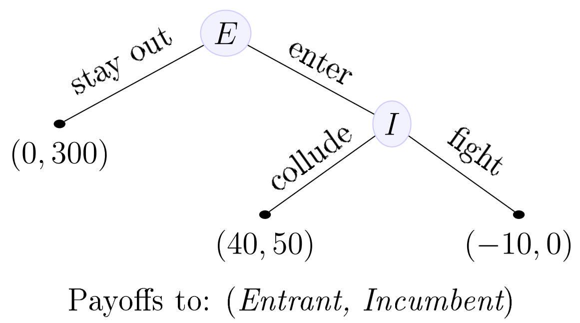

Figure 1: Entry deterrence game in extensive form

From game to solution

- There are two Nash equilibria
  (enter, collude) and (stay out, fight).
- The behavior of player I in the latter one is not plausible. It is not subgame perfect: Once player I is prompted to move the only rational action is *collude*.

4 Games with Imperfect Information

4.1 Introduction

Introduction

- An important extension of extensive form games is the inclusion of *imperfect information*.

- When it's their turn, players may not know, what the last move of an opponent or of Nature was. They did not directly observe it.
- Perhaps they can exclude some possible actions, but still cannot distinguish between others.
- Part of the history of the game may be covered by an opaque veil.
- This may be relevant for their potential outcomes.
- How can they decide?
- How can we include that in the description of the game?

From story to game

Consider a follow the leader game with two players. Let us start with the full information version of the sequential game.

- Smith and Jones are duopolists in a particular market.
- There is a possibility to change the design of the product from *small* to *large*.
- If both firms change the design, duopoly profits increase. If only one of the players changes the design, demand shrinks and both firms make losses.
- Coordinated action is required.

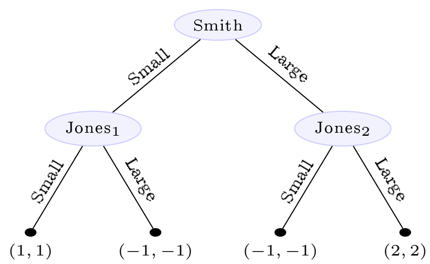

From game to solution

This game has two subgames originating at Jones$_1$ and Jones$_2$, respectively. In subgame one *Small* is the best answer of Jones, whereas in subgame two it is *Large*. In other words (*Small/Small, Large/Large*) is the dominant strategy of Jones. On the other hand, *Large* is Smith's best answer to Jones's dominant strategy. Hence, the game has a unique subgame perfect equilibrium { *Large*, (*Small/Small, Large/Large*) } with equilibrium payoffs (2,2).

4.2 Representing Imperfect Information

Information sets

Assume Jones does not know whether Smith has chosen *Small* or *Large*.

- In our (simple) example the lack of information on Jones' side makes the game equivalent to a (static) normal form game with simultaneous moves.
- In general we will have to include the representation of information in the game tree.
- Jones cannot distinguish between node $Jones_1$ and node $Jones_2$
- We can indicate that by a **dotted line** connecting the nodes.
- We say the connected nodes belong to the same *information set*.

Game tree with information sets

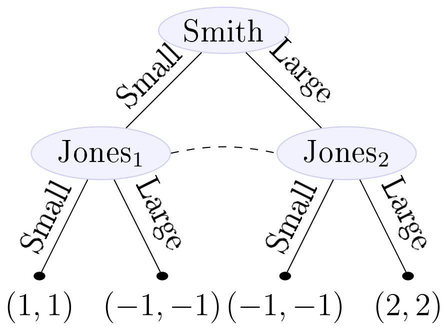

Figure 2: Follow-the-Leader in extensive form with asymmetric information

- This game has no proper subgame!
- The strategy space of Jones is smaller now, compared to the full information game! Like in the simultaneous move game it is { Small, Large }.

The extensive form of Follow-the-Leader turns into a game equivalent to the simple strategic form version of the game, when Jones cannot distinguish between the nodes $Jones_1$ and $Jones_2$.

	S	L
S	$(1,1)$	$(-1,-1)$
L	$(-1,-1)$	$(2,2)$

Payoffs to: (Smith, Jones)

This game has two Nash equilibria

$$\{ \ Large, \ Large \ \} \text{ and } \{ \ Small, \ Small \ \}$$

However, there some details we did not examine yet. They may be part of the dynamic structure of the game and have substantial implications.

Definition
Player i's **information set** at any particular point of the game is the set of different nodes in the game tree that he knows might be the actual node, but between which he cannot distinguish by direct observation.

Definition
Player i's **information partition** is a collection of his information sets such that

1. Each path is represented by one node in a single information set in the partition.
2. The predecessors of all nodes in a single information set are in one information set.

- In an obvious way we can characterize information partitions as being **finer** or **coarser**.
- A change in the information structure from coarser to fine is called an **information refinement**.
- The finest information structure is reached, when each information set consists of a **singleton**.

Obviously a particular question arises from this setting:

- How can player Jones find a best answer to a strategy of Smith?
- Or even more substantial: What is a strategy of a player suffering from incomplete information?

4.2.1 Information sets and strategies

Consider the following game as an example:

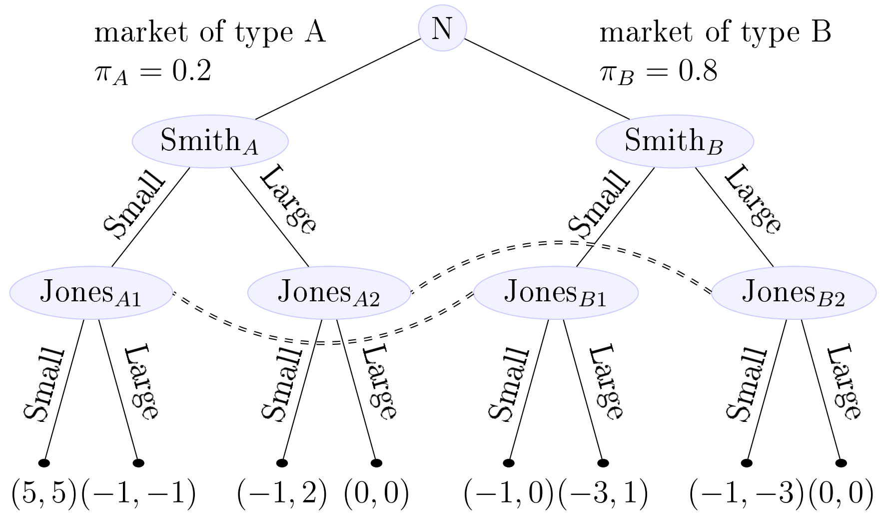

- We have to adjust the concept of strategies: Nodes can no longer be the anchor points of planned actions which make up strategies.
- We have to relate plans to information sets because a player cannot plan to behave different at different nodes within an information set. He simply would not know how to execute such a plan.
- This implies: At different nodes in an information set the player must have identical sets of actions he can choose from.

First we give a concise description the **information structure**:

$$\Omega_{Smith} = \{\omega_A, \omega_B\} = \{\{Smith_A\}, \{Smith_B\}\}$$

$$\Omega_{Jones} = \{\omega_1, \omega_2\} = \{\{Jones_{A1}, Jones_{B1}\}, \{Jones_{A2}, Jones_{B2}\}\}$$

Now we determine the **strategy set** of Smith

$$\begin{aligned} \sigma_{Smith} \quad = \quad & \{(Small|A, Small|B), (Small|A, Large|B), \\ & (Large|A, Small|B), (Large|A, Large|B)\} \end{aligned}$$

Obviously there is a shorter, equivalent notation

$$\sigma_{Smith} = \{\, (S|A, S|B)\,,\ (S|A, L|B)\,,\ (L|A, S|B)\,,\ (L|A, L|B)\,\}$$

and the **strategy set** of Jones

$$\begin{aligned}\sigma_{Jones} &= \{(Small|1, Small|2), (Small|1, Large|2),\\ &\quad (Large|1, Small|2), (Large|1, Large|2)\}\end{aligned}$$

Again we can use the shorter notation

$$\sigma_{Jones} = \{\, (S|1, S|2)\,,\ (S|1, L|2)\,,\ (L|1, S|2)\,,\ (L|1, L|2)\,\}$$

Heads up!
There are 1001 ways to label strategies. Consider the following ones:

- $$\begin{aligned}\sigma_{Jones,1} &= (Small|\omega_1, Small|\omega_2), \ldots,\\ \sigma_{Jones,3} &= (Large|\omega_1, Small|\omega_2), \ldots\end{aligned}$$
- $$\begin{aligned}\sigma_{Jones,1} &= (Small|Small, Small|Large), \ldots,\\ \sigma_{Jones,3} &= (Large|Small, Small|Large), \ldots\end{aligned}$$
- $$\begin{aligned}\sigma_{Jones,1} &= (Small\ after\ Small, Small\ after\ Large), \ldots,\\ \sigma_{Jones,3} &= (Large\ after\ Small, Small\ after\ Large), \ldots\end{aligned}$$
- and finally an extremely short notation
$$\begin{aligned}\sigma_{Jones,1} &= (S\,,\ S), \ldots,\\ \sigma_{Jones,3} &= (L\,,\ S), \ldots\end{aligned}$$
- The more stages the game has or the more complex the information structure is, the more useful is the explicit reference to information sets. Always try to avoid misunderstandings.

We realize:

- Irrespective of the lack of information of Jones, Smith can evaluate his pay-offs for both types of markets for any strategy profile.
- Jones faces the problem that his pay-offs depend on the random type of market. By assumption he does not observe the outcome of this random event.
- Jones has to consider expected pay-offs rather than certain pay-offs.

What are the probabilities Jones should use for his calculations?

- The simple answer is: He can use (π_A, π_B)!
- On the basis of these fundamental probabilities we can calculate mutual best responses to search for Nash-equilibria.
- The more sophisticated answer is: Maybe Jones can learn from Smith! Smith knows the type of market when he takes his move. The strategy choice of Smith may reveal what Smith knows.
- Jones may replace the probabilities (π_A, π_B) by more reliable probabilities called beliefs which are conditional on the strategy of Smith.
- It was Thomas Bayes who told our ancestors in the 18th century how to do this.

- We call the fundamental probabilities π specified in the game tree *priors*.
- Assume the goal of each player is to maximize expected pay-off based on the priors π.
- Notice that in the simple example player Smith knows the nature of the market when he has to take his move. For any strategy profile he can compute his pay-off without suffering from uncertainty.
- On the other hand, player Jones's pay-offs have to be expected pay-offs.

4.2.2 Expected Pay-Offs and Best Responses

- $$\begin{aligned} E\Pi_{Jones}(Small|Small) &= 0.2 \cdot 5 + 0.8 \cdot 0 = 1 \\ E\Pi_{Jones}(Large|Small) &= 0.2 \cdot (-1) + 0.8 \cdot 1 = 0.6 \end{aligned}$$
- The best response of Jones to *Small* is *Small*
- $$\begin{aligned} E\Pi_{Jones}(Small|Large) &= 0.2 \cdot 2 + 0.8 \cdot (-3) = -2 \\ E\Pi_{Jones}(Large|Large) &= 0.2 \cdot 0 + 0.8 \cdot 0 = 0 \end{aligned}$$
- The best response of Jones to *Large* is *Large*

Summing up:

- $\sigma^*_{Jones} = (Large \mid Large,\ \ Small \mid Small)$ is the best strategy of Jones.
- In a way it is a dominant strategy, i.e. a strategy independent of the strategy of Smith. This is because Jones is not interested in learning why Smith chooses *Large* or *Small*. There seems to be no advantage to him from knowing how Smith's move depends on the type of market.
- This is different for Smith. Obviously his best response to σ^*_{Jones} is $\sigma^*_{Smith} = (Small \mid A,\ \ Large \mid B)$,
- whereas Smith's best response to always *Large* is always *Large* as well.

Expected Pay-Offs Revisited
Let us recall the calculations of expected pay-offs:

- We assumed Jones to use the priors π to calculate his expected pay-off. In particular this makes him choose *Small* if he observes *Small*.
- Therefore his expected pay-off after observing *Small* is

$$E\Pi_{Jones}(Small|Small) = 0.2 \cdot 5 + 0.8 \cdot 0 = 1$$

- Imagine Jones knows the strategy of Smith! Observing *Small* he can now conclude that the market must be of type A! Therefore he should expect to get pay-off 5.
- It is rather unlikely that this happens because π_A is so small. But if he observes *Small* he should conclude that the unlikely event has happened and revise the expected pay-off.
- Remember: Within the information sets ω_1, ω_1 the probabilities π_A, π_B define one and the same probability distribution.
- The revision of probabilities may lead to different distributions in the two information sets.

4.2.3 Beliefs and Assessments

Beliefs
Beliefs are probability distributions for nodes within information sets. They are formed by players. A complete set of distributions for each non-trivial information set is called a belief profile. For convenience they should be listed player by player. However, in simple examples we usually have only one player suffering from a lack of information or other words only one player who has non-trivial information ssets.

Notation: We use μ to denote beliefs, opposed to π for priors.

Assessments
Beliefs typically change expected pay-offs. Hence, they interfere with equilibrium considerations. They should be derived in a rational way in relation to the opponents' strategies.

In our current example we specify rational beliefs for the two information sets (ω_1, ω_2) labeled

$$(\mu_1, \mu_2) = (\,(\mu_{1A,}, \mu_{1B}),\ (\mu_{2A,}, \mu_{2B})\,)$$

As discussed above a rational assessment derived from the dominance consideration can only be

$$(\,(\mu_{1A,}, \mu_{1B}),\ (\mu_{2A,}, \mu_{2B})\,) = (\,(1, 0),\ (0, 1)\,)$$

Notice that the assessment is a system of conditional probabilities, conditional on the strategy of Smith.

- In our example only Jones forms beliefs. The beliefs above make a full belief profile.
- The observation itself - either *Small* or *Large* - does not reveal anything to Jones!
- The interpretation of the observation in context with a strategy may reveal something:
  The moves are assumed to be the outcome of a strategy *(Small / A, Large / B)* of Smith!
- The combination of a strategy and beliefs constitutes an assessment.

4.3 Bayesian Inference

How do we derive beliefs from strategies?

- Bayes rule is the statistical tool for the rational method to derive beliefs from a combination of priors and additional information.
- In our example the moves of Smith come with probabilities zero or one according to his strategy.

$$prob(Small|A) = 1\,,\ prob(Large|A) = 0$$
$$prob(Small|B) = 0\,,\ prob(Large|B) = 1$$

- And the market is of type A or B according to the priors

$$prob(A) = \pi_A,\ prob(B) = \pi_B$$

Recall Bayes' rule for a particular type of market $t_0 \in \{A, B\}$, and an action of Smith $a_0 \in \{Small, Large\}$

$$prob(t_0|a_0) = \frac{prob(a_0|t_0)\, prob(t_0)}{prob(a_0|A)\, prob(A) + prob(a_0|B)\, prob(B)}$$

Evaluation with respect to the strategy considered

$$prob(A|Small) = 1,\ prob(B|Small) = 0$$
$$prob(A|Large) = 0,\ prob(B|Large) = 1$$

Bayes rule inverts the direction of reactions on information: Smith decides how to move when he observes the type. Jones infers the conditional probabilities of types when he observes a move.

Refinements of Subgame Perfect Equilibrium

Decades after early contributions by Harsanyi (1967) and Selten (1975) the most often used concept today is that by Kreps and Wilson (1982). Whereas the early discussion was more about robustness the latter shifted the focus towards asymmetry of information together with rational use of information. Beyond that one may look for robustness in that context as well.

In a perfect Bayesian equilibrium

- some players observe Nature's move, and that may completely close their information gap. At least it will lead to a natural update of priors.
- Others may not be informed about Nature's move, but can monitor the informed players and try to make use of their observations. They may update priors on the basis of observed moves or strategies of the others.

4.4 Perfect Bayesian Equilibrium

From story to game
Recall the Follow the Leader example with imperfect, asymmetric information

- Smith is the first to decide whether to adhere to large or change to small design.
- Jones does not learn about Smith's choice before he has to decide.

Now replace the incomplete information concerning the move of Smith by imperfect information about Smith's type.

- Nature assigns type A, B, or C to Smith with probabilities π_A, π_B, π_C.
- In case of type A both are better off with *Large*; in case of type B the players have conflicting preferences; in case of C both are better off with *Small*.
- Smith knows his type and decides for *Small* or *Large*.
- Jones observes Smith's move and tries to maximize his expected payoff.
- He takes into account the probabilities of Nature's assignment and the observed move of Smith.

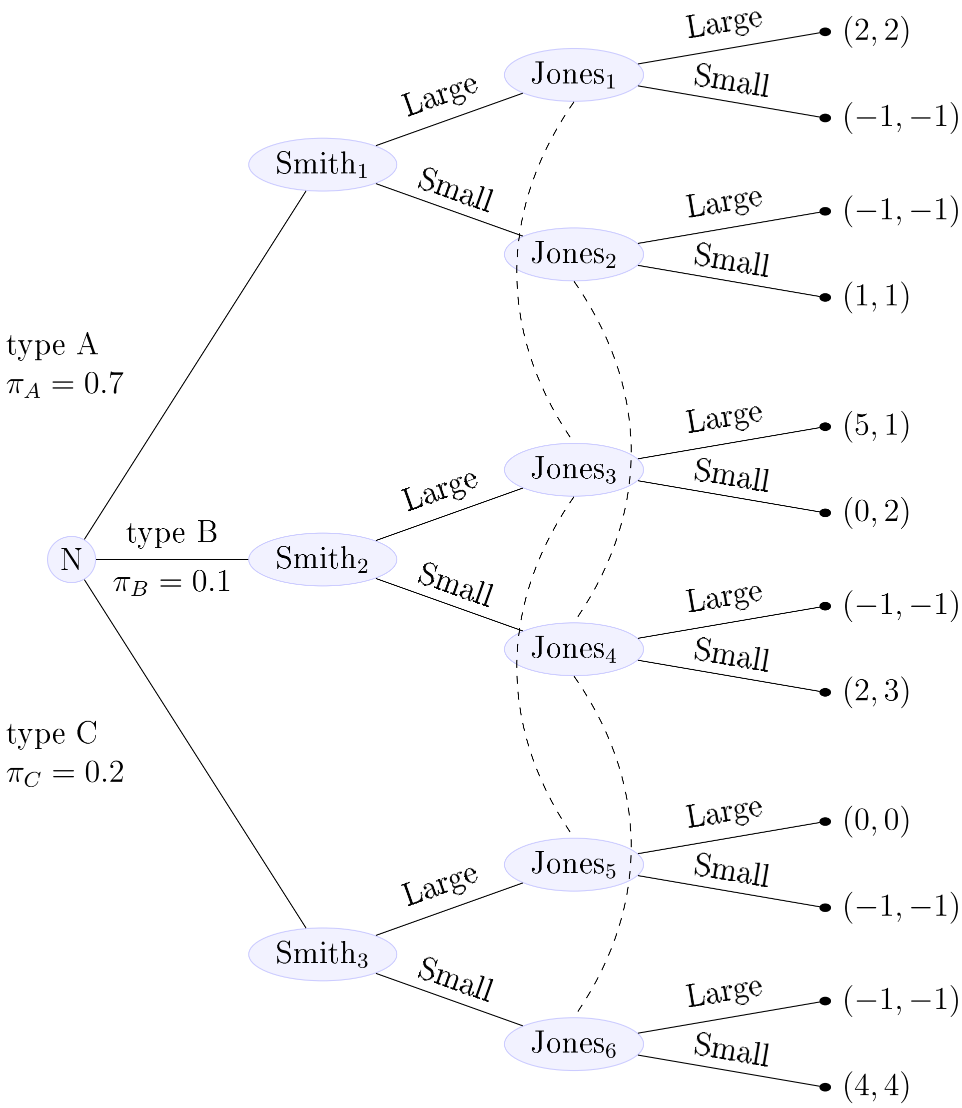

Payoffs to: (Smith, Jones)

Figure 3: Follow-the-Leader III

4.4.1 Strategy-Belief Profile

Player Jones has two non-trivial information sets, i.e. information sets which are not a singleton. They are reached by Smith's move L or S, respectively. We may denote them by $\omega_J(L) = \{$ Jones$_1$, Jones$_3$, Jones$_5$ $\}$ and $\omega_J(S) = \{$ Jones$_2$, Jones$_4$, Jones$_6$ $\}$.

The concept of Bayesian equilibrium requires to attach beliefs to the nodes in information sets. In the game considered here they may be denoted by $\mu = (\mu_J(L), \mu_J(S))$ with $\mu_J(L) = (\mu_{A|L},\, \mu_{B|L},\, \mu_{C|L})$ and $\mu_J(S) = (\mu_{A|S},\, \mu_{B|S},\, \mu_{C|S})$.

$\mu_J(L)$ and $\mu_J(S)$ are conditional probability distributions, i.e. the probabilities per information set add up to one.

We call μ a belief profile.

A plausible candidate for a strategy profile in equilibrium is

$$\sigma = \{\sigma_S, \sigma_J\} = \{\ (\mathrm{L|A,\ L|B,\ S|C}),\ (\mathrm{L|L,\ S|S})\ \}$$

Yet, what are the expected payoffs?

The beliefs have to be part of the equilibrium specification. We call the combination of strategies and beliefs a *strategy-belief profile.*

With the help of beliefs expected payoffs are well defined and we can apply the concept of subgame perfect Nash equilibrium.

4.4.2 Bayesian updating of priors

Beliefs should not be arbitrary. We consider them to be rational in equilibrium if they are formed according to Bayes rule, whenever this is possible. This links the beliefs profile to the priors and a particular strategy profile.

When Jones assumes Smith plays the equilibrium strategy σ_S, and when he observes S, he should conclude that Smith is of type C with probability $\mu_C = 1$

When Jones assumes Smith plays the equilibrium strategy σ_S, and when he observes L, he should conclude that Smith is of type A or B with probabilities: $(\mu_A, \mu_B) = (7/8, 1/8)$

Recall the strategy profile of our game

$$\sigma = \{(L|A, L|B, S|C)\,,\ (L|L, S|S)\}$$

We apply Bayes' rule and use labels $\mu_{A|L}$ to denote $prob(A|L)$, ...

$$\begin{aligned}
\mu_{A|L} &= \frac{prob(L|A)\,\pi_A}{prob(L|A)\,\pi_A + prob(L|B)\,\pi_B + prob(L|C)\,\pi_C} \\
&= \frac{1 \cdot 0.7}{1 \cdot 0.7 + 1 \cdot 0.1 + 0 \cdot 0.2} = \frac{0.7}{0.7 + 0.1} = \frac{7}{8}
\end{aligned}$$

$$\mu_{B|L} = \frac{1 \cdot 0.1}{1 \cdot 0.7 + 1 \cdot 0.1 + 0 \cdot 0.2} = \frac{0.1}{0.7 + 0.1} = \frac{1}{8}$$

$$
\begin{aligned}
\mu_{C|L} &= 0 & \mu_{C|S} &= 1 \\
\mu_{A|L} &= 7/8 & \mu_{A|S} &= 0 \\
\mu_{B|L} &= 1/8 & \mu_{B|S} &= 0
\end{aligned}
$$

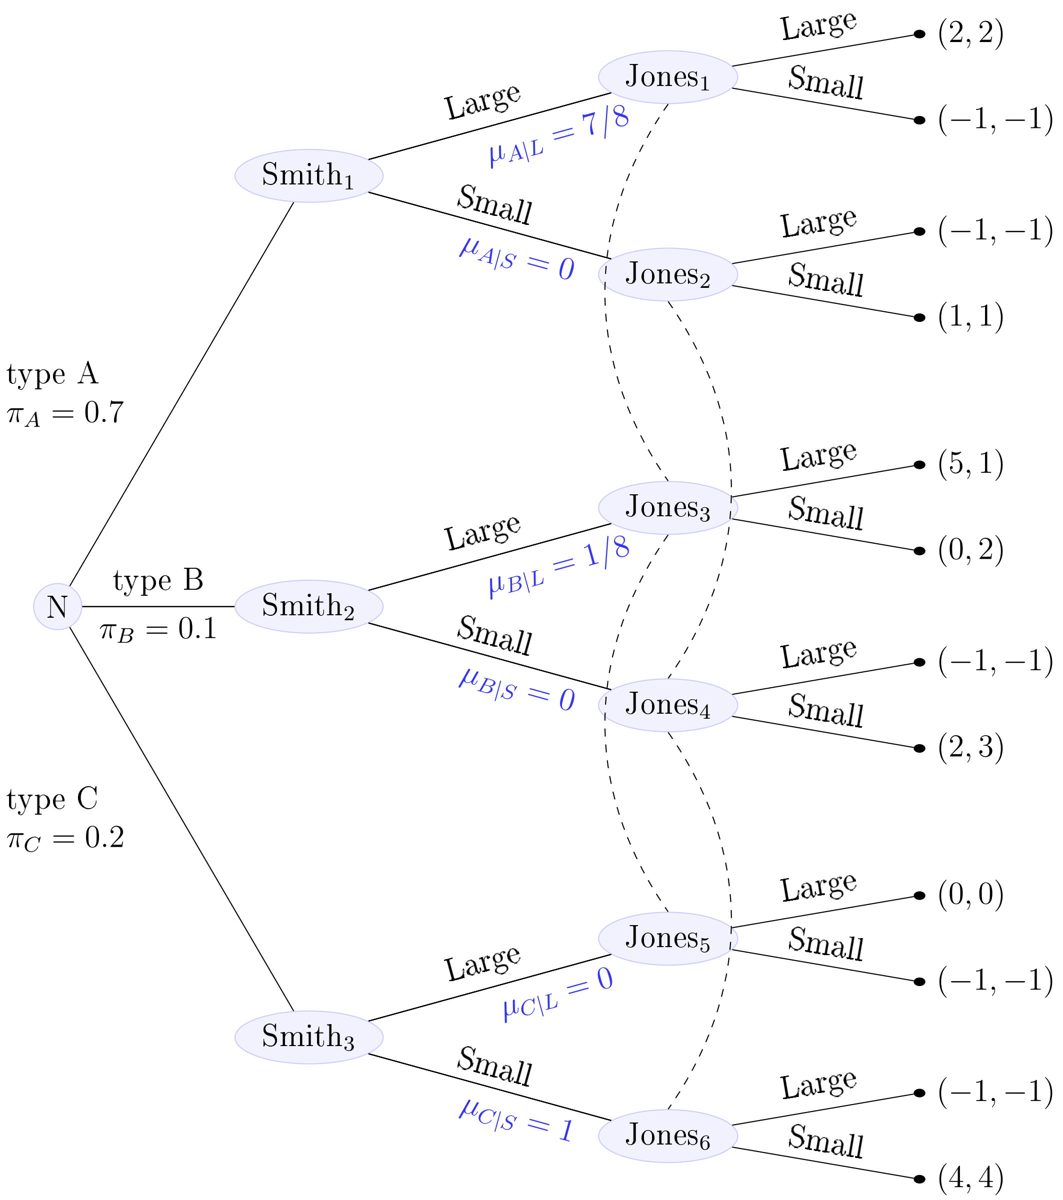

Payoffs to: (Smith, Jones)

Figure 4: Follow-the-Leader III

To verify the equilibrium, we compute the expected payoffs for Jones based on Bayesian beliefs. These expected payoffs may either reject or confirm the strategy profile combined with the respective Bayesian beliefs as a strategy-belief profile of a subgame perfect Bayesian Nash equilibrium.

1. σ_J is best answer to σ_S:
   - $(L|L)$ yields $\Pi_J(L|L) = \frac{7}{8}\,2 + \frac{1}{8}\,1 = \frac{15}{8}$ and
   - $(S|L)$ yields $\Pi_J(S|L) = \frac{7}{8}\,(-1) + \frac{1}{8}\,2 = -\frac{5}{8}$
   - Hence, L is best answer to L.

   - $(S|S)$ yields $\Pi_J(S|S) = 1 \cdot 4 = 4$ and
   - $(L|S)$ yields $\Pi_J(L|S) = 1 \cdot (-1) = -1$.
   - Hence, S is best answer to S.

   Summing up, we confirm, that to follow the leader is Jones' best response.

2. σ_S is best answer to σ_J:
   - In case Smith is of type A or B, L followed by L yields a higher payoff to Smith than S followed by S because $2 > 1$ and $5 > 2$.
   - In case Smith is of type C, S followed by S yields a higher payoff to Smith than L followed by L. because $4 > 0$

Uniqueness of equilibrium

Is this the only Bayesian equilibrium of this game? There are eight strategies for Smith, four strategies for Jones, and therefore 32 strategy profiles (verify!).

Only the expected payoffs of Jones depend on beliefs. To Jones' four strategies $\sigma_{J_1} = (L|L, S|S)$, $\sigma_{J_2} = (L|L, L|S)$, $\sigma_{J_3} = (S|L, S|S)$ and $\sigma_{J_4} = (S|L, L|S)$ we find Smith's best answers without going through the process of Bayesian updating. This will identify four candidates for Smith's equilibrium strategies.

The equilibrium above already includes the first one, σ_{J_1}, and its best answer.

Let us continue with the remaining ones.

Can σ_{J_2} be an equilibrium strategy of Jones?

- $\sigma_{J_2} = (L|L, L|S)$ yields a best answer $\sigma_{S_2} = (L|A, L|B, L|C)$.

Can σ_{J_2} be an equilibrium strategy of Jones?
The uniform strategy of Smith gives no rise to effective updating. Indeed, all conditional probabilities $prob(L|.)$ are equal to one and therefore $\mu_{A|L} = \frac{\pi_A}{\pi_A+\pi_B+\pi_C} = \pi_A$ and so forth.

- In this case Bayesian updating is possible but does not change anything. The denominator of the Bayesian formula is equal to one. The Bayesian beliefs are identical to the corresponding priors.

The opposite can happen as well. According to a strategy of a player, an information set may never be reached. The denominator then is equal to zero, and we cannot apply Bayes' formula. We will deal with that problem later.

Can σ_{J_2} be an equilibrium strategy of Jones?
It remains to check, whether σ_{J_2} is a best answer to σ_{S_2} using priors as beliefs.

- $(L|L)$ yields an expected payoff $\Pi_J(L|L) = 0.7 \cdot 2 + 0.1 \cdot 1 + 0.2 \cdot 0 = 1.5$, whereas $(S|L)$ yields $\Pi_J(S|L) = 0.7 \cdot (-1) + 0.1 \cdot 2 + 0.2 \cdot (-1) = -0.7$
- Hence, L is the best answer to L.
- $(L|S)$ yields $\Pi_J(L|S) = 0.7 \cdot (-1) + 0.1 \cdot (-1) + 0.2 \cdot (-1) = -1$,
- whereas $(S|S)$ yields $\Pi_J(S|S) = 0.7 \cdot 1 + 0.1 \cdot 3 + 0.2 \cdot 4 = 1.8$
- Hence, S is best answer to S. $\hookrightarrow$ *contradiction!*

$\sigma = \{(L|A, L|B, L|C), (L|L, L|S)\}$ is no equilibrium profile

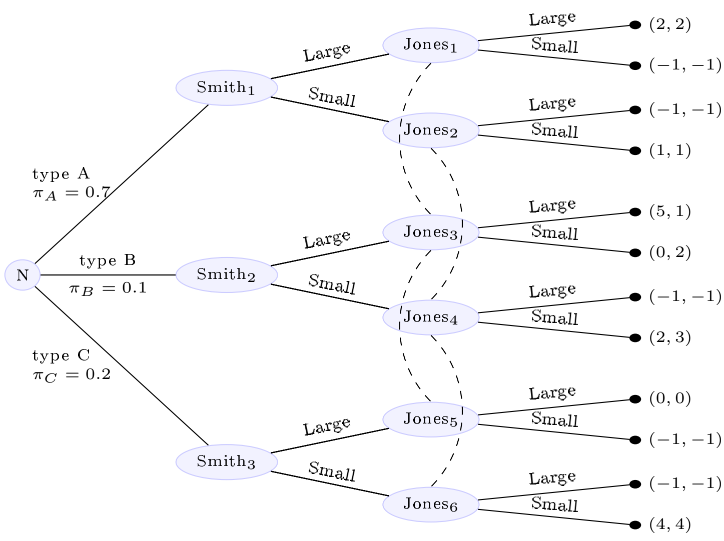

Summing up, $\sigma_{J_2} = (L|L, L|S)$ cannot be part of an equilibrium strategy-belief profile.

Can σ_{J_3} be an equilibrium strategy of Jones?

- $\sigma_{J_3} = (S|L, S|S)$ yields a best answer

$$\sigma_{S_3} = (S|A, S|B, S|C)$$

  Again uniformity gives no rise to updating and we use priors to compute expected payoffs. By the same reasoning as above we conclude, that $\sigma_{J_3} = (S|L, S|S)$ cannot be part of an equilibrium strategy-belief profile either.

$\sigma = \{(S|A, S|B, S|C), (S|L, S|S)\}$ is no equilibrium profile

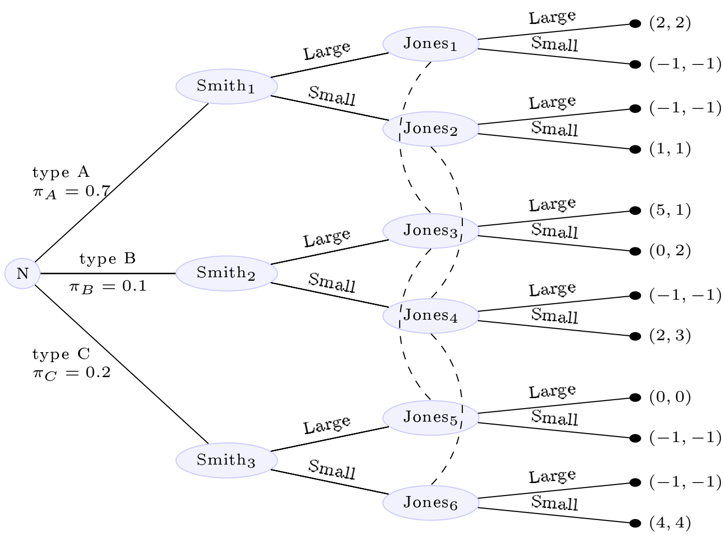

Can σ_{J_4} be an equilibrium strategy of Jones?

- $\sigma_{J_4} = (S|L, L|S)$ yields a best answer $\sigma_{S_4} = (\ \{S \text{ or } L\}|A, L|B,\ \{S \text{ or } L\}|C)$. In case he observes A or C Smith is indifferent between L and S. Neither $(L|A, L|B, S|C)$ nor $(L|A, L|B, L|C)$ can be part of an equilibrium together with $\sigma_{J_4} = (S|L, L|S)$, as seen before. Hence we are left with $(S|A, L|B, S|C)$ and $(S|A, L|B, L|C)$.

- Let us examine them one by one
  - Consider $(S|A, L|B, S|C)$. Bayesian updating yields three non-zero beliefs $\mu_{B|L} = 1$, $\mu_{A|S} = 7/9$ and $\mu_{C|S} = 2/9$. We compute: $\Pi_J(S|L) = 2$ is larger than $\Pi_J(L|L) = 1$, which supports the equilibrium. However, $\Pi_J(L|S) = \frac{7}{9}\cdot(-1) + \frac{2}{9}\cdot(-1) = -1$ and $\Pi_J(S|S) = \frac{7}{9}\cdot 1 + \frac{2}{9}\cdot 4 = \frac{15}{9}$ reject the equilibrium.
  - Consider $(S|A, L|B, L|C)$. Now Bayesian updating yields three non-zero beliefs $\mu_{A|S} = 1$, $\mu_{B|L} = 1/3$ and $\mu_{C|L} = 2/3$. We compute: $\Pi_J(S|S) = 1$ is larger than $\Pi_J(L|S) = -1$, which rejects the equilibrium. Even more, $\Pi_J(L|L) = \frac{1}{3}\cdot 1 + \frac{2}{3}\cdot 0 = \frac{1}{3}$ and $\Pi_J(S|L) = \frac{1}{3}\cdot 2 + \frac{2}{3}\cdot(-1) = 0$ reject the equilibrium.

We finally conclude that the game has a unique perfect Bayesian equilibrium.

4.5 Sequential Rationality and Consistency of Beliefs

Another Example
In order to work out what the single requirements for a perfect Bayesian equilibrium contribute to the concept we consider the following further example of a game in extensive form with complete but imperfect information

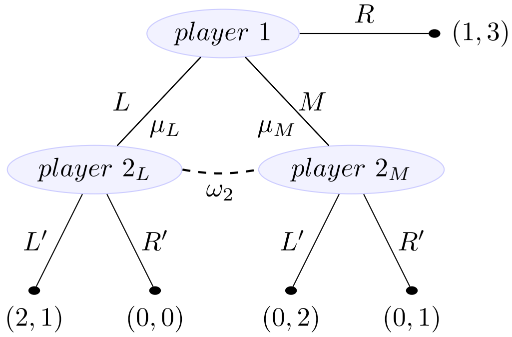

The strategy sets of the players are

$$\sigma_1 = \{L, M, R\}, \quad \sigma_2 = \{L', R'\}$$

Normal Form Representation
We may represent the game in normal form

		player 2 L'	R'
player 1	L	2, 1	0, 0
	M	0, 2	0, 1
	R	1, 3	1, 3

The Nash equilibria in pure strategies are: (L, L') and (R, R')

Notice that the reduction to the normal game wipes out potential sequential action. In general this raises the question whether the equilibria are subgame perfect in the extensive form game?

However, in our example there are no subgames due to the imperfect information of player two!

Beliefs
At the information set ω_2 (which is not a singleton) player two must have beliefs to be able to compute expected pay-offs.

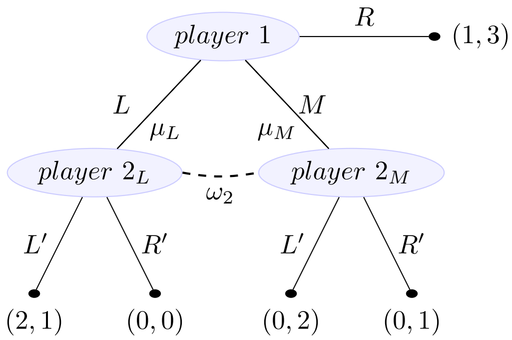

- $E\Pi_2(L') = \mu_L \cdot 1 + (1 - \mu_L) \cdot 2 = 2 - \mu_L$
- $E\Pi_2(R') = \mu_L \cdot 0 + (1 - \mu_L) \cdot 1 = 1 - \mu_L$
- Notive that in our example $E\Pi_2(L') > E\Pi_2(R')$ no matter what the beliefs are!

Definition 1 (Sequential Rationality)**.** A player is said to be *sequentially rational* if and only if, at each information set where he is to move, he tries to maximize his expected pay off given his beliefs and the other players ' subsequent strategies.

In our example

- this eliminates the equilibrium (R, R') for any belief $(\mu_L,\ \mu_M)$;
- the beliefs do not affect the selection of equilibrium, but without them player 2 cannot act sequentially rational;
- obviously, the Bayesian beliefs for (L, L') are $(\mu_L,\ \mu_M) = (1,\ 0)$.

Indeed, the evaluation of Bayes' rule at the information set ω_2 of player 2 with probabilities $p(.)$ gives

$$\begin{aligned} p(L\,|\,\omega_2) &= \frac{p(\omega_2|L) \cdot p(L)}{p(\omega_2|L)p(L) + p(\omega_2|M)p(M) + p(\omega_2|R)p(R)} \\ &= \frac{1 \cdot 1}{1 \cdot 1 + 1 \cdot 0 + 0 \cdot 0} = 1 \end{aligned}$$

whereas

$$p(M\,|\,\omega_2) = \frac{1 \cdot 0}{1 \cdot 1 + 1 \cdot 0 + 0 \cdot 0} = 0$$

Remark: We included the term $p(\omega_2|R)p(R)$ only to demonstrate the implication of a move which does not lead to the information set considered. As the conditional probability is zero, we could have dropped it from the beginning.

Mixed Strategies

The formula can be applied in case of mixed strategies as well. Consider probabilities $p(L) = \alpha_L$, $p(M) = \alpha_M$, and $p(R) = \alpha_R$ for the moves $L,\ M,\ R$ to get

$$\begin{aligned} p(L\,|\,\omega_2) &= \frac{p(\omega_2|L) \cdot \alpha_L}{p(\omega_2|L)\alpha_L + p(\omega_2|M)\alpha_M + p(\omega_2|R)\alpha_R} \\ &= \frac{1 \cdot \alpha_L}{1 \cdot \alpha_L + 1 \cdot \alpha_M + 0 \cdot \alpha_R} = \frac{\alpha_L}{\alpha_L + \alpha_M} \end{aligned}$$

whereas

$$p(M \mid \omega_2) = \frac{1 \cdot \alpha_M}{1 \cdot \alpha_L + 1 \cdot \alpha_M + 0 \cdot \alpha_R} = \frac{\alpha_M}{\alpha_L + \alpha_M}$$

Notation
Given a (possibly mixed) strategy profile σ, an information set is said to be on the path of play if and only if this information set is reached with positive probability according to σ.

- In case of pure strategies it happens regularly that some information sets are on the path of play others are not.
- In case of mixed strategies it still may happen.
- In case of trembling hands it will not happen anymore.

Definition 2 (Consistency on a path)**.** Given any strategy profile σ and an information set ω on the path of play of σ a player's beliefs at ω are said to be consistent on the path of σ if and only if the beliefs are derived using Bayes' rule and σ.

Now consider a variation of the previous game. The exit decision is given to a new player and we have a proper subgame beginning at node *player 2*.

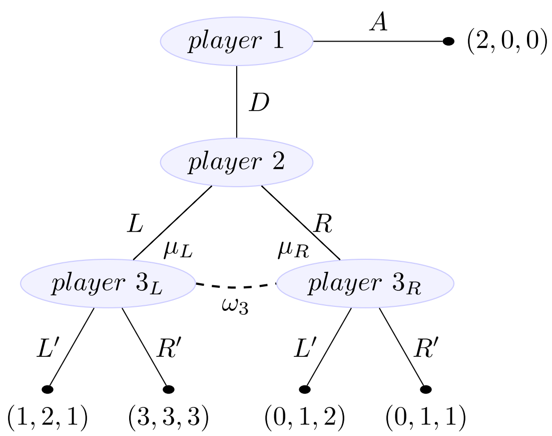

The search for a perfect Bayesian equilibrium of this game is straight forward.

- The equilibrium has to be subgame perfect.
- The subgame between players 2 and 3 has a unique Nash-equilibrium (L, R').
- Player 1 prefers to enter the subgame of the other players because $3 > 2$. Hence $\sigma^* = (D,\ L,\ R')$ is the unique subgame perfect Nash equilibrium of the (entire) game.
- The Bayesian beliefs of player 3 given the strategy profile σ are $\mu = (\mu_L,\ \mu_R) = (1, 0)$.
- σ and μ form a perfect Bayesian equilibrium.

Is there an equilibrium with action A?

- All strategy pairs of player 2 and player 3 are best responses to A.
- A is best response to all strategy pairs except to (L, R').
- Consider $\sigma = (A,\ L,\ L')$ which is a profile of mutually best responses, a Nash equilibrium.
- Notice that the information set ω_3 is off the equilibrium path.
- Assume player 3's belief is $\mu = (0, 1)$. Given this belief, L' is optimal.
- But, σ is not subgame perfect as we found out earlier.

Can we extend the requirement to form rational or at least reasonable beliefs off the equilibrium path?

Definition 3 (Consistency off a path)**.** At information sets off the equilibrium path of a strategy profile σ, beliefs must be determined by Bayes' rule and σ where possible.

In the example above $\sigma = (A,\ L,\ L')$ with belief is $\mu = (0, 1)$ is not consistent off the path of σ.

- L chosen by player 2 together with consistency off the path enforces belief $\mu = (1, 0)$.
- And, L' is no longer the best response of player 3 because only R' is sequentially rational with μ.

Now consider a combination of the previous games. The exit decision is given to *player 2* asides from *player 1*.

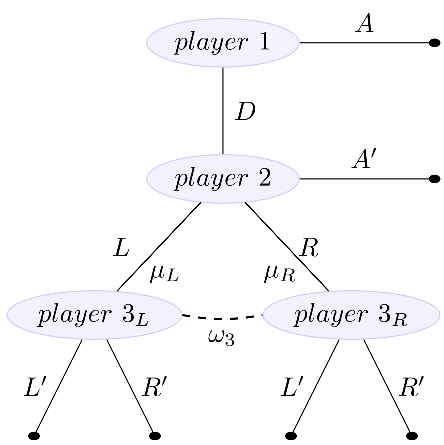

- Assume *player 1*'s equilibrium strategy is A. Then ω_3 is off the equilibrium path.
- If *player 2*'s strategy is A' in this version of the game, *player 3* cannot apply Bayes rule to form his belief.

- On the other hand, if *player 2* uses a mixed strategy $\alpha = (\alpha_L, \alpha_R, \alpha_{A'})$ with $\alpha_L + \alpha_R > 0$, then the Bayesian belief of *player 3* is

$$\mu = \left(\frac{\alpha_L}{\alpha_L + \alpha_R}, \ \frac{\alpha_R}{\alpha_L + \alpha_R} \right)$$

- In a Nash equilibrium no player chooses a strictly dominated strategy.
- In a perfect Bayesian equilibrium, the requirements of having beliefs and sequential rationality are equivalent to insisting that no player's strategy be strictly dominated beginning at any information set.
- Nash and Bayesian Nash equilibrium do not share this feature at information sets off the equilibrium path.
- Even a subgame perfect Nash equilibrium may not share this feature at some information sets off the equilibrium path, such as information sets that are not contained in any subgame.
- In a perfect Bayesian equilibrium, players cannot threaten to play strategies that are strictly dominated beginning at any information set off the equilibrium path.