Part III: Conditional beliefs in dynamic games

Lecture 5:

Common initial belief in rationality

7 Common initial belief in rationality

Until now, we have focused on **static games:**

- every player chooses only **once,**
- if a player makes a choice, he has **no information** about the choices made (or to be made) by others.

So, the belief that you have about the opponents' choices can **never be contradicted** during the game.

We shall now turn to **dynamic games:**

- players choose **sequentially,**
- a player may choose **more than once** during the game,
- if a player makes a choice, he may have **full, or partial, information** about the choices made by others so far.

So, it may happen that you **initially** believe that your opponent chooses a, but **later** you observe that your opponent has **not** chosen a.

In this case, you must **revise** your belief about the opponent.

The way in which you **revise your belief** will be important for the **eventual choice(s)** you make.

Questions:

How can we formalize **belief revision** in an epistemic model?

What are **reasonable** ways to revise your belief?

What **behavioral consequences** do they have?

7.1 Example: Father and son

Story: Suppose, you have a twelve year old son.

On the table there are two enormous piles of papers that have to be put in alphabetical order, and this seems a perfect job for him.

You first ask him whether he wants to do the job.

If he accepts, then he will receive a reward if he does the job properly.

However, you will not be at home when your son is doing the job, and checking both piles when you arrive at home requires quite some work.

This may give your son an incentive to cheat, by doing only one pile instead of two.

Utilities for the son:

If you find out that your son has only done one pile, he must do the other pile as well, but will receive no reward.

If you do not catch him cheating, he will get a reward of 5 euros.

Doing a pile reduces his utility by 3.

Utilities for you:

Checking a pile reduces your utility by 1.

If you reward your son unjustly, this would reduce your utility by 5.

If you check no pile, this would appear immoral to you, and would reduce your utility by 5.

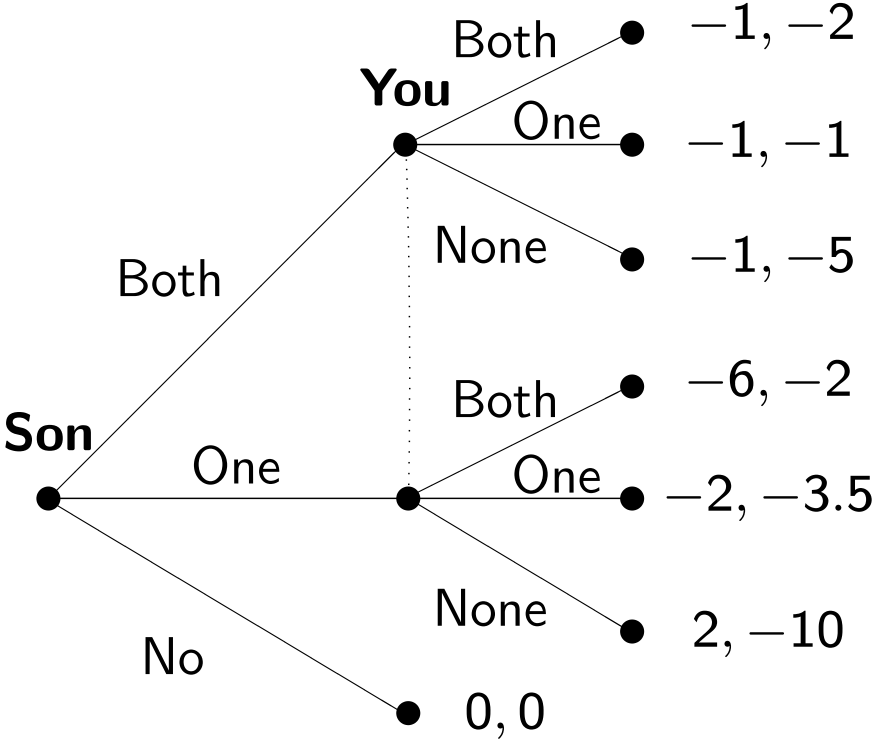

Initially, you believe that your son believes that you will check at least one pile.

Therefore, you initially believe that your son will reject the job.

So, at your information set you must **revise** your belief about your son.

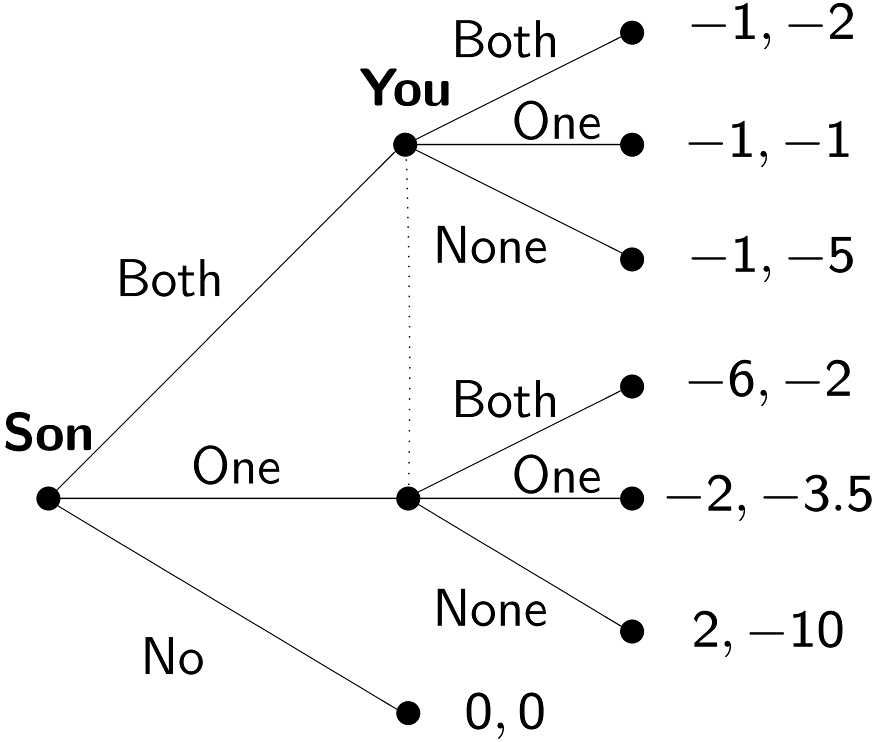

Belief revision scenario 1:

You conclude that your son must have done **only one pile,** since doing two piles can never be optimal (given the depicted utilities).

Therefore, you will **check both piles.**

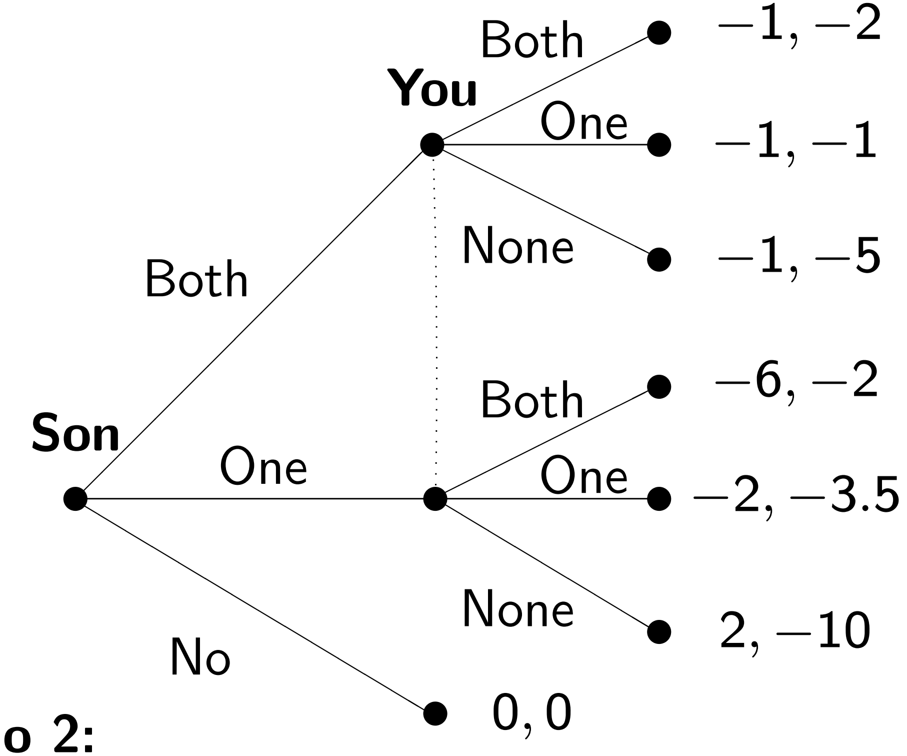

Belief revision scenario 2:

Initially, you believe that your son prefers doing two piles over doing one pile.

You **still believe so** at your information set.

Hence you will **check only one pile.**

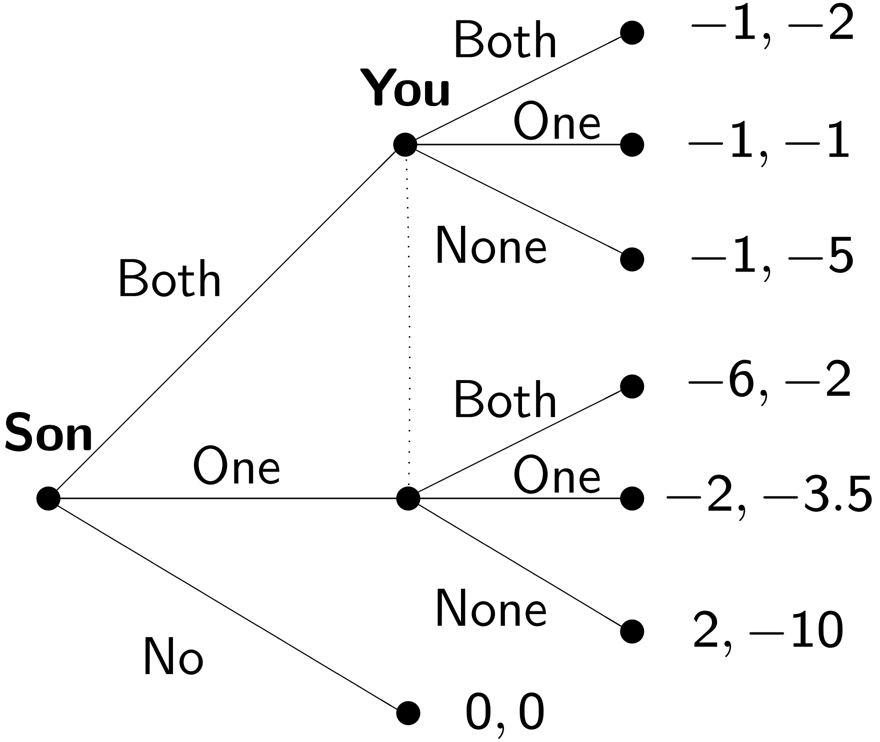

So, the **way in which you revise your belief** crucially matters for the choice you eventually make!

7.2 Dynamic games

The rules of the game are specified by a **finite game tree** $Tr = (X, Z, E)$**,** consisting of:

- finite set X of **non-terminal nodes,** specifying the **decision nodes,**
- finite set Z of **terminal nodes,** specifying the different situations in which the game may end,
- finite set E of **edges** $(x, y) \in X \times (X \cup Z)$, specifying the **choices** for the players. Choice (x, y) moves the game from x to y.

A **finite dynamic game** Γ consists of:

- a finite game tree $Tr = (X, Z, E)$,
- a finite set of players I,
- for every player i, a set $X_i \subseteq X$ of **decision nodes,** such that $\cup_{i \in I} X_i = X$, and $X_i \cap X_j = \emptyset$ for all $i \neq j$. The set X_i represents the set of nodes at which i must choose.
- for every player i, an **information partition** $H_i = (h_i^1, ..., h_i^K)$ of X_i. Every $h_i^k \subseteq X_i$ is an **information set.** Player i cannot distinguish between any two nodes in h_i^k. Every two nodes $x, y \in h_i^k$ must have the same number of outgoing edges.

- for every player i, and every information set $h_i \in H_i$, a set of **available choices** $C_i(h_i)$ at h_i. Every choice $c_i \in C_i(h_i)$ corresponds to a set of outging edges, one for every node $x \in h_i$.

- for every player i, a utility function $u_i : Z \rightarrow \mathbf{R}$, assigning to every terminal node $z \in Z$ a utility $u_i(z)$.

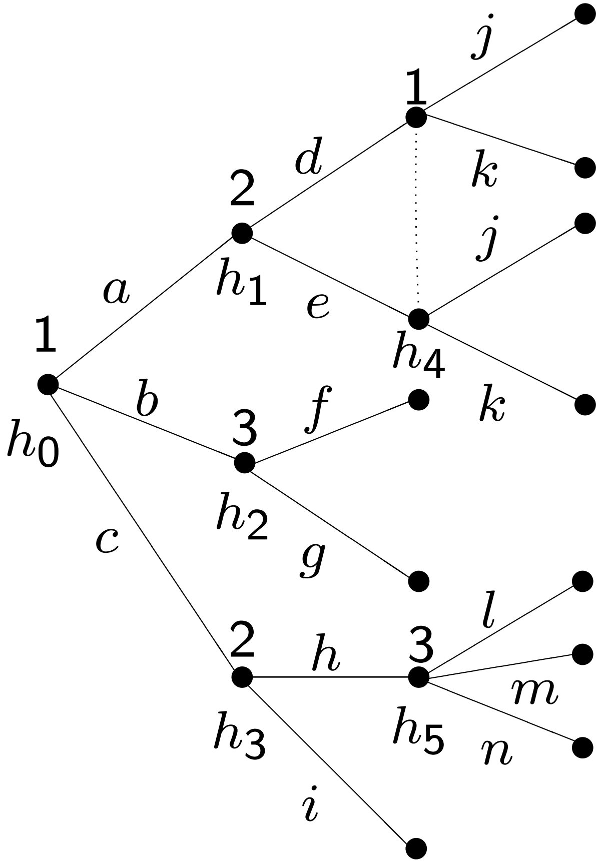

Information partitions: $H_1 = \{h_0, h_4\}$, $H_2 = \{h_1, h_3\}$, $H_3 = \{h_2, h_5\}$.

Choices: $C_1(h_0) = \{a, b, c\}$, $C_1(h_4) = \{j, k\}$, $C_2(h_1) = \{d, e\}$, $C_2(h_3) = \{h, i\}, C_3(h_2) = \{f, g\}$, $C_3(h_5) = \{l, m, n\}$.

A **strategy** for player i is a function s_i that assigns to every information set $h_i \in H_i$ **that is reachable by** s_i an available choice $s_i(h_i) \in C_i(h_i)$.

We denote by S_i the set of strategies for player i.

Be careful:

Traditionally, a strategy s_i is defined as a function that assigns to **every** information set $h_i \in H_i,$ whether reachable by s_i or not, an available choice $c_i \in C_i(h_i)$.

However, I find choices prescribed at information sets h_i that are **not reachable** by s_i problematic.

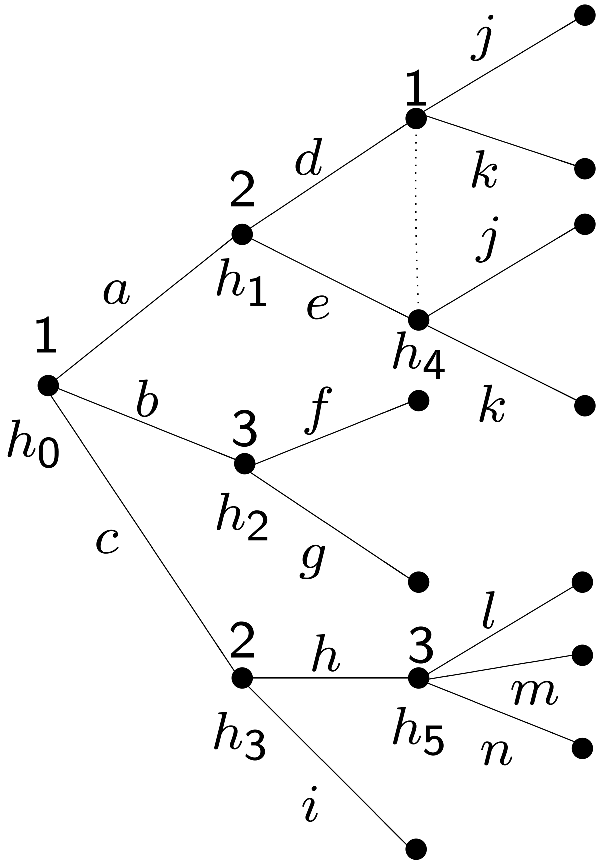

Strategies: $S_1 = \{(a, j), (a, k), b, c\}$

$S_2 = \{(d, h), (d, i), (e, h), (e, i)\}$

$S_3 = \{(f, l), (f, m), (f, n), (g, l), (g, m), (g, n)\}$

7.3 Conditional beliefs

Informally, a **conditional belief** about the opponents' strategies specifies at the beginning of the game, and at each of your information sets h_i, a belief about the opponents' strategies that are **possible** given that h_i has been reached.

Let h_0 denote the **beginning** of the game. Let $H_i^* := H_i \cup \{h_0\}$.

For every player i, and every information set $h_i \in H_i^*$, let $S_{-i}(h_i) \subseteq S_{-i}$ be the set of opponents' strategy profiles that are **possible,** given that h_i has been reached.

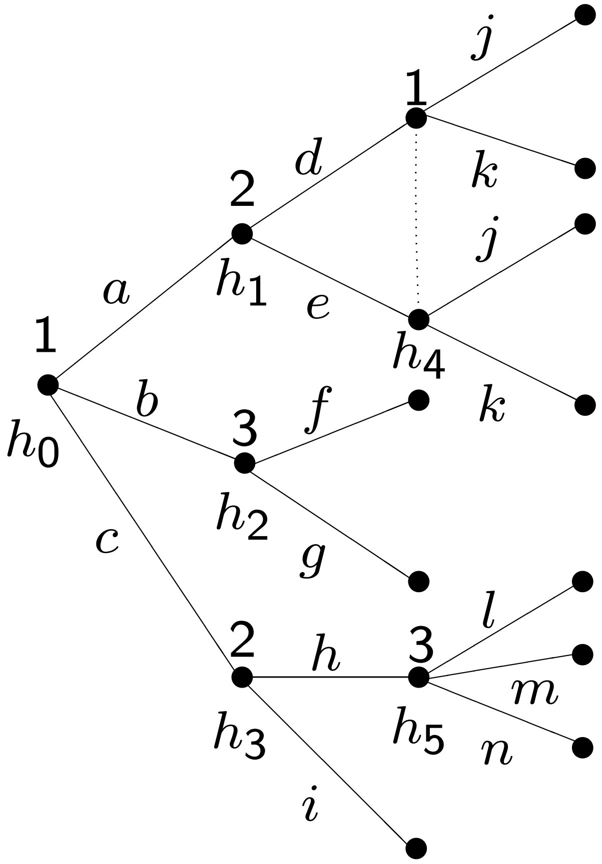

If h_1 is reached, player 1 must have chosen a strategy from $\{(a, j), (a, k)\}$.

So, $S_{-2}(h_1) = \{(a, j), (a, k)\} \times S_3$.

Similarly, $S_{-2}(h_3) = \{c\} \times S_3$.

Clearly, $S_{-i}(h_0) = S_{-i}$ for every player i.

Formally, a **conditional belief** for player i about the opponents' strategies is a vector $b_i = (b_i(h_i))_{h_i \in H_i^*}$ where

$$b_i(h_i) \in \Delta(S_{-i}(h_i))$$

for all information sets $h_i \in H_i^*$.

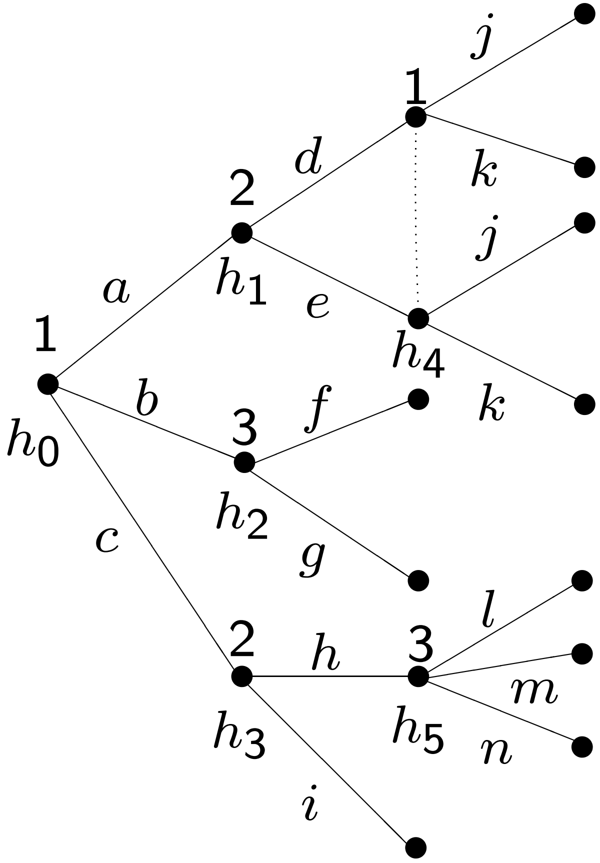

Example: $b_2(h_0) = \frac{1}{2}((a, j), (f, l)) + \frac{1}{2}(b, (g, m))$,

$b_2(h_1) = ((a, j), (f, l))$, $b_2(h_3) = (c, (f, m))$.

7.4 Optimal strategies

For player i, consider a conditional belief b_i about the opponents' strategies.

For a given strategy $s_i \in S_i$, let $H_i(s_i)$ be the set of information sets for player i that are reachable by s_i (including h_0).

For an information set $h_i \in H_i(s_i)$, let $u_i(s_i, b_i|h_i)$ be the **expected utility at** h_i from choosing s_i, given the belief $b_i(h_i)$ at h_i about the opponents' strategy choices.

The strategy s_i is **optimal at** h_i under the conditional belief b_i if

$$u_i(s_i, b_i|h_i) \geq u_i(\hat{s}_i, b_i|h_i)$$

for all $\hat{s}_i \in S_i(h_i)$.

Here, $S_i(h_i) \subseteq S_i$ are the player i strategies for which h_i is reachable.

The strategy s_i is **rational** under the conditional belief b_i if s_i is **optimal at every information set** $h_i \in H_i(s_i)$ under b_i.

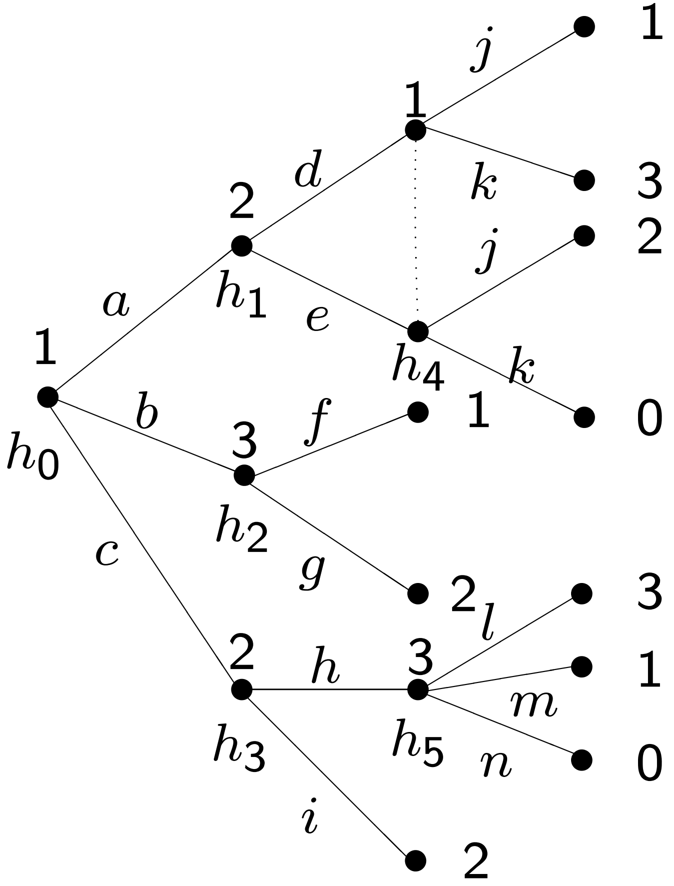

Conditional belief b_2:

$b_2(h_0) = \frac{1}{2}((a,j),(f,l)) + \frac{1}{2}(b,(g,m))$, $b_2(h_1) = ((a,j),(f,l))$, $b_2(h_3) = (c,(f,m))$.

Only strategy (e,i) is rational for player 2 under the conditional belief b_2.

7.5 Epistemic model

At every information set h_i, player i holds a conditional belief about:

- the opponents' strategy choices,
- the conditional belief that j holds, at every h_j, about his opponents' strategies,
- the conditional belief that j holds, at every h_j, about the conditional belief that k holds, at every h_k, about his opponents' strategy choices, and so on.

How can we model this in a compact way?

Let Γ be a finite dynamic game.

A **finite epistemic model** for Γ is a tuple $\mathbf{M} = (T_i, b_i)_{i \in I}$ where

- T_i is the finite set of types for player i, and
- b_i is a function that assigns to every type $t_i \in T_i$ and **to every information set** $h_i \in H_i^*$ a **conditional belief**

$$b_i(t_i, h_i) \in \Delta(S_{-i}(h_i) \times T_{-i}).$$

So, as the game proceeds, player i can **revise his belief** about the opponents' strategies, but also **about the opponents' conditional beliefs.**

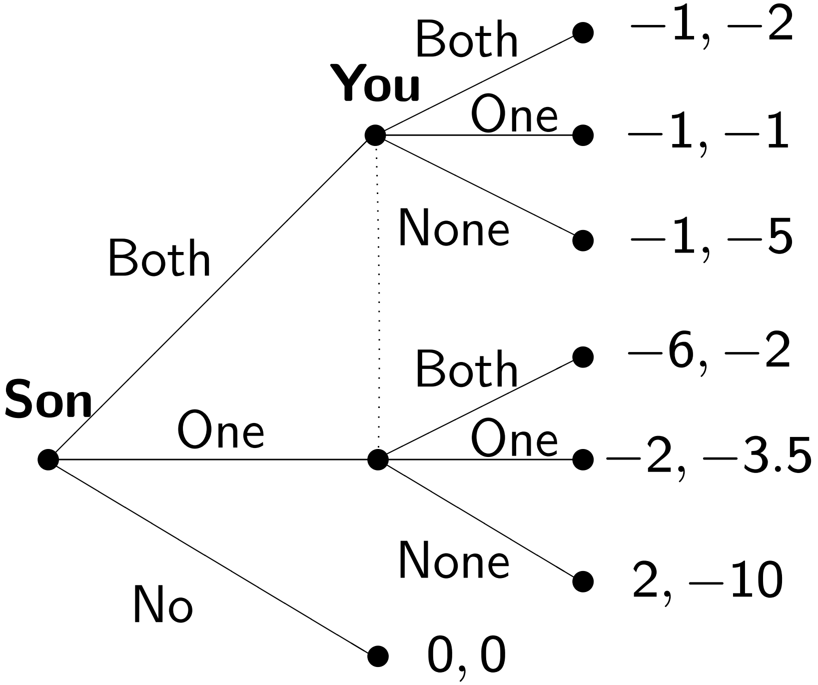

Epistemic model: $T_1 = \{t_1^{both}, t_1^{none}\}$ (you), $T_2 = \{t_2^{one}, t_2^{no}\}$ (son).

$b_1(t_1^{both}, h_0) = (no, t_2^{no})$, $b_1(t_1^{both}, h_1) = (one, t_2^{one})$,

$b_1(t_1^{none}, h_0) = b_1(t_1^{none}, h_1) = (one, t_2^{one})$.

$b_2(t_2^{one}, h_0) = (none, t_2^{none})$, $b_2(t_2^{no}, h_0) = (both, t_1^{both})$.

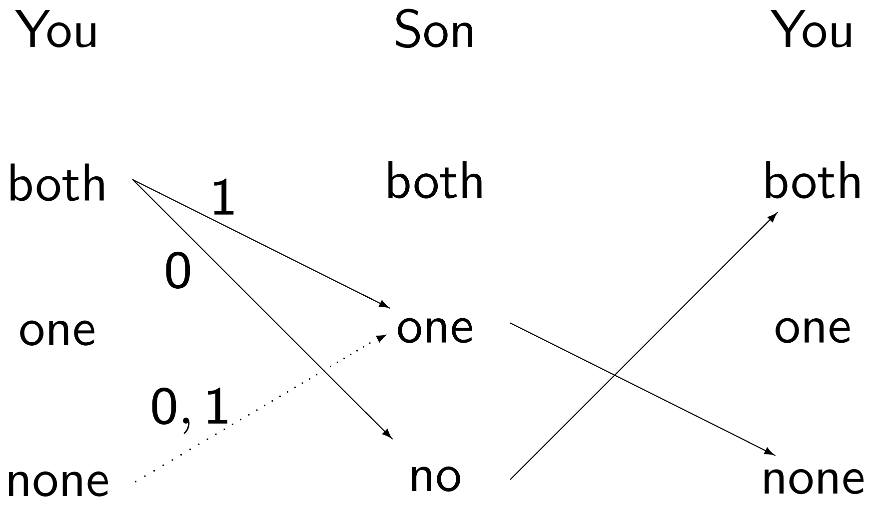

You Son You

both both both

1

0

one one one

0, 1

none no none

Epistemic model: $T_1 = \{t_1^{both}, t_1^{none}\}$ (you), $T_2 = \{t_2^{one}, t_2^{no}\}$ (son).

$b_1(t_1^{both}, h_0) = (no, t_2^{no})$, $b_1(t_1^{both}, h_1) = (one, t_2^{one})$,

$b_1(t_1^{none}, h_0) = b_1(t_1^{none}, h_1) = (one, t_2^{one})$.

$b_2(t_2^{one}, h_0) = (none, t_2^{none})$, $b_2(t_2^{no}, h_0) = (both, t_1^{both})$.

7.6 Common initial belief in rationality

Strategy s_i is **rational** for type t_i if s_i is rational for t_i's conditional belief about the opponents' strategies.

Type t_i **initially believes in the opponents' rationality** if the initial belief $b_i(t_i, h_0)$ only assigns positive probability to strategy-type pairs (s_j, t_j) where s_j is rational for type t_j.

Note: A type t_i that **initially** believes in the opponents' rationality may believe **later** that his opponents have chosen irrational strategies.

Type t_i **structurally believes in the opponents' rationality** if **at every information set** $h_i \in H_i^*$, the conditional belief $b_i(t_i, h_i)$ only assigns positive probability to strategy-type pairs (s_j, t_j) where s_j is rational for type t_j.

So, you **always** believe that your opponents have chosen rationally, and will choose rationally in the future.

Problem:

Common **structural** belief in the opponents' rationality is in general not possible !

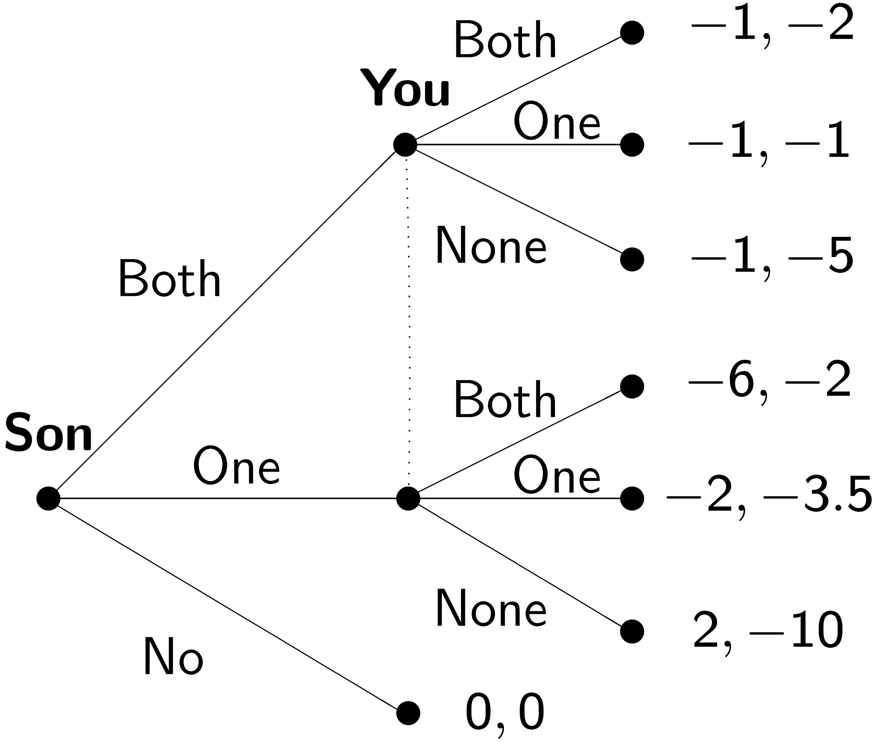

At your information set h_1 (after observing that your son has accepted the job) you **cannot** simultaneously believe that your **son chooses rationally,** and that your **son believes that you choose rationally.**

Let $E_j \subseteq T_j$ be a subset of opponent j's types.

For instance, $E_j :=$ set of j's types that initially believe in the opponents' rationality.

Type t_i **initially believes in** E_j if the initial belief $b_i(t_i, h_0)$ only assigns positive probability to player j types in E_j.

Type t_i expresses **common initial belief in rationality** (Ben-Porath, 1997) if:

- t_i initially believes in the opponents' rationality,
- t_i initially believes that every opponent initially believes in his opponents' rationality,
- t_i initially believes that every opponent initially believes that every opponent initially believes in his opponents' rationality,

and so on.

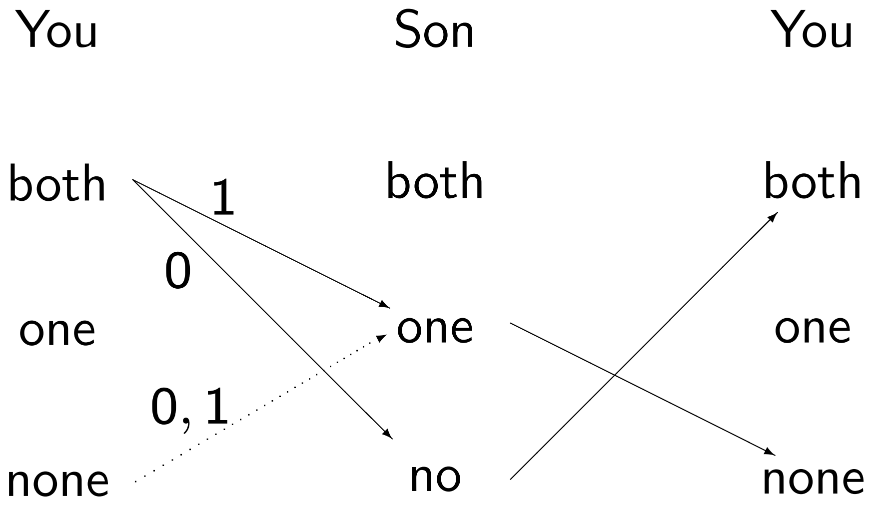

You Son You

both 1 both both

0

one one one

0, 1

none no none

Your type t_1^{both} expresses **common initial belief in rationality.**

So, you can rationally choose to check both piles under common initial belief in rationality.

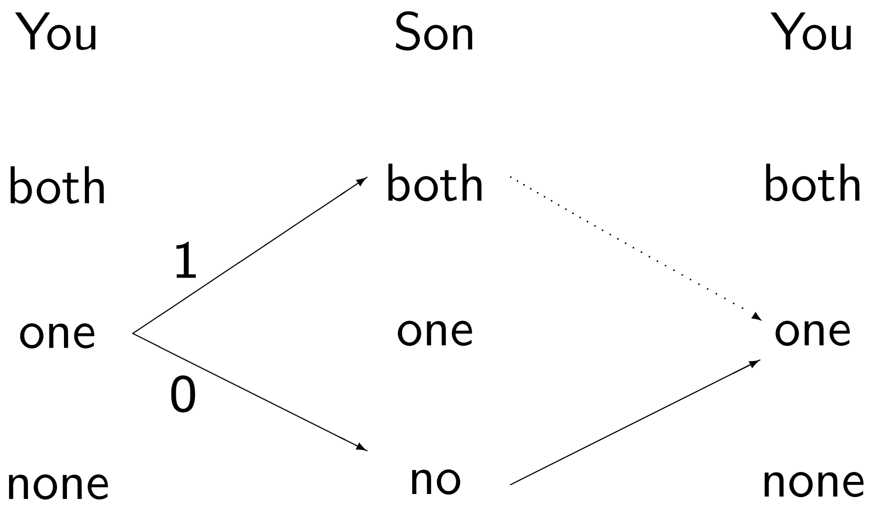

You

both

one

none

1

0

Son

both

one

no

You

both

one

none

Your type t_1^{one} expresses **common initial belief in rationality.**

So, you can rationally choose to check one pile under common initial belief in rationality.

Hence, under **common initial belief in rationality** you can rationally choose to **check both piles,** or to **check only one pile.**

7.7 Algorithm: Ben-Porath procedure

Can we find an algorithm that, for every finite dynamic game, computes all the strategies that can rationally be chosen under **common initial belief in rationality?**

Say that a strategy s_i is **rational** in the dynamic game Γ if s_i is rational for some conditional belief b_i about the opponents' strategies.

For every opponent j, let $D_j \subseteq S_j$ be a subset of j's strategies, and let $D_{-i} := \times_{j \neq i} D_j$.

Say that strategy s_i is **strictly dominated on** D_{-i} if there is some randomized strategy $\mu_i \in \Delta(S_i)$ such that

$$u_i(s_i, s_{-i}) < u_i(\mu_i, s_{-i})$$

for all $s_{-i} \in D_{-i}$.

Algorithm: Ben-Porath procedure (Ben-Porath (1997))

$$
\begin{aligned}
S_i^1 &:= \{s_i \in S_i \mid s_i \textbf{ rational} \text{ in the dynamic game}\} \\
S_i^2 &:= \{s_i \in S_i^1 \mid s_i \text{ not } \textbf{strictly dominated} \text{ on } S_{-i}^1\} \\
S_i^3 &:= \{s_i \in S_i^2 \mid s_i \text{ not } \textbf{strictly dominated} \text{ on } S_{-i}^2\} \\
&\vdots \\
S_i^k &:= \{s_i \in S_i^{k-1} \mid s_i \text{ not } \textbf{strictly dominated} \text{ on } S_{-i}^{k-1}\} \\
&\vdots
\end{aligned}
$$

Theorem 7.1: (Based on Ben-Porath (1997))

Let Γ be a finite dynamic game.

Then strategy s_i can rationally be chosen under **common initial belief in rationality**

if and only if

strategy s_i survives the **Ben-Porath procedure.**

Proof:

(a) **Assume:** s_i is rational for some type t_i that expresses common initial belief in rationality.

To show: s_i survives the Ben-Porath procedure.

$$S_i^1 := \{s_i \in S_i \mid s_i \textbf{ rational} \text{ in the dynamic game}\}.$$

Since s_i is rational for t_i, we must have that $s_i \in S_i^1$.

$S_i^2 := \{s_i \in S_i^1 \mid s_i$ not **strictly dominated** on $S_{-i}^1\}$.

Since t_i **initially believes in the opponents' rationality,** $b_i(t_i, h_0)$ should only assign positive probability to strategies in S_{-i}^1.

Since s_i is rational for t_i, strategy s_i must in particular be optimal for $b_i(t_i, h_0)$.

So, s_i is optimal for some $b_i(h_0) \in \Delta(S_{-i}^1)$.

Hence, s_i is not strictly dominated on S_{-i}^1, which means that $s_i \in S_i^2$.

$$S_i^3 := \{s_i \in S_i^2 \mid s_i \text{ not } \textbf{strictly dominated} \text{ on } S_{-i}^2\}.$$

Since t_i **initially believes in the opponents' rationality,** and **initially believes that his opponents initially believe in their opponents' rationality,** $b_i(t_i, h_0)$ should only assign positive probability to strategies in S_{-i}^2.

Since s_i is rational for t_i, strategy s_i must in particular be optimal for $b_i(t_i, h_0)$.

So, s_i is optimal for some $b_i(h_0) \in \Delta(S_{-i}^2)$.

Hence, s_i is not strictly dominated on S_{-i}^2, which means that $s_i \in S_i^3$.

An so on. So, $s_i \in S_i^k$ for all k, and hence survives the Ben-Porath procedure.

(b) **Assume:** $\hat{s}_i$ survives the Ben-Porath procedure.

To show: $\hat{s}_i$ is rational for some type $\hat{t}_i$ that expresses **common initial belief in rationality.**

Let $(D_i)_{i \in I}$ be the sets of strategies that survive the Ben-Porath procedure. By construction:

- every $s_i \in D_i$ is rational for some conditional belief about the opponents' strategies,

- every $s_i \in D_i$ is not strictly dominated on D_{-i}.

Hence:

- for every $s_i \in D_i$ there is some conditional belief $\hat{b}_i^{s_i}$ about the opponents' strategies such that at every $h_i \in H_i$ strategy s_i is optimal under $\hat{b}_i^{s_i}(h_i) \in \Delta(S_{-i}(h_i))$,

- every $s_i \in D_i$ is optimal under some belief $\tilde{b}_i^{s_i}(h_0) \in \Delta(D_{-i})$.

For every $s_i \in D_i$, construct the conditional belief $b_i^{s_i}$ about the opponents' strategies by

$$b_i^{s_i}(h_i) = \begin{cases} \tilde{b}_i^{s_i}(h_i), & \text{if } h_i = h_0 \\ \hat{b}_i^{s_i}(h_i), & \text{if } h_i \neq h_0 \end{cases}.$$

Hence, every $s_i \in D_i$ is rational under $b_i^{s_i}$, and $b_i^{s_i}$ initially believes that opponents choose from D_{-i}.

For every $s_i \notin D_i$, choose an arbitrary conditional belief $b_i^{s_i}$ about the opponents' strategies.

Now, construct an epistemic model $(T_i, b_i)_{i \in I}$, where $T_i = \{t_i^{s_i} \mid s_i \in S_i\}$ for every player i, and the conditional belief $b_i(t_i^{s_i})$ is given

$$b_i(t_i^{s_i}, h_i)((s_j, t_j)_{j \neq i}) = \begin{cases} b_i^{s_i}(h_i)((s_j)_{j \neq i}), \text{ if } t_j = t_j^{s_j} \text{ for all } j \\ 0, \text{ otherwise} \end{cases} .$$

For every $s_i \in D_i$, we have:

- s_i is rational for $t_i^{s_i}$,

- the initial belief $b_i(t_i^{s_i}, h_0)$ only assigns positive probability to strategy-type pairs $(s_j, t_j^{s_j})$, where $s_j \in D_j$.

Hence, for every $s_i \in D_i$, the type $t_i^{s_i}$ expresses **common initial belief in rationality.**

So, every $s_i \in D_i$ can rationally be chosen under common initial belief in rationality.

This completes the proof.

7.8 Connection with lexicographic beliefs

Recall: A **lexicographic belief** for player i about the opponents' strategies is a list $\lambda_i = (\lambda_i^1, ..., \lambda_i^K)$ where $\lambda_i^k \in \Delta(S_{-i})$ for all $k \in \{1, ..., K\}$.

The lexicographic belief λ_i is **cautious** if for every $s_{-i} \in S_{-i}$ there is some $k \in \{1, ..., K\}$ with $\lambda_i^k(s_{-i}) > 0$.

Every **cautious lexicographic** belief induces, in a natural way, a **conditional** belief !

Take a cautious lexicographic belief $\lambda_i = (\lambda_i^1, ..., \lambda_i^K)$ about the opponents' strategies.

For every $h_i \in H_i^*$, let $k(h_i)$ be the first k for which $\lambda_i^k(S_{-i}(h_i)) > 0$.

Here, $\lambda_i^k(S_{-i}(h_i)) := \sum_{s_{-i} \in S_{-i}(h_i)} \lambda_i^k(s_{-i})$.

Define $b_i^{\lambda_i}(h_i) \in \Delta(S_{-i}(h_i))$ by

$$b_i^{\lambda_i}(h_i)(s_{-i}) := \frac{\lambda_i^{k(h_i)}(s_{-i})}{\lambda_i^{k(h_i)}(S_{-i}(h_i))} \text{ for all } s_{-i} \in S_{-i}(h_i).$$

Then, $b_i^{\lambda_i} = (b_i^{\lambda_i}(h_i))_{h_i \in H_i^*}$ is the **conditional belief induced by** λ_i.

Define $b_i^{\lambda_i}(h_i) \in \Delta(S_{-i}(h_i))$ by

$$b_i^{\lambda_i}(h_i)(s_{-i}) := \frac{\lambda_i^{k(h_i)}(s_{-i})}{\lambda_i^{k(h_i)}(S_{-i}(h_i))} \text{ for all } s_{-i} \in S_{-i}(h_i).$$

In particular, $b_i^{\lambda_i}(h_0) = \lambda_i^1$.

So, **initial belief** in $b_i^{\lambda_i}$ corresponds to **weak belief** in λ_i.

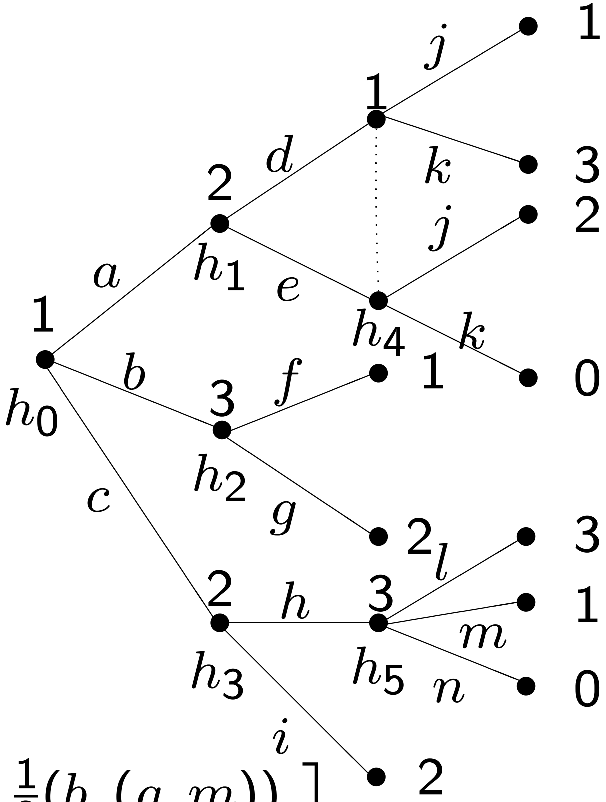

$$\lambda_2 = \left[\begin{array}{c} \frac{1}{2}((a,j),(f,l)) + \frac{1}{2}(b,(g,m)) \\ (c,(f,m)) \\ \cdots \end{array}\right] \text{induces}$$

$b_2(h_0) = \frac{1}{2}((a,j),(f,l)) + \frac{1}{2}(b,(g,m))$, $b_2(h_1) = ((a,j),(f,l))$, $b_2(h_3) = (c,(f,m))$.

Lemma 7.2:

Let Γ be a finite dynamic game, and $\lambda_i = (\lambda_i^1, ..., \lambda_i^K)$ a cautious lexicographic belief about the opponents' strategies. Take a strategy $s_i \in S_i$.

If s_i is **rational** for the **lexicographic** belief λ_i, then s_i is also **rational** for the induced **conditional** belief $b_i^{\lambda_i}$.

With the above lemma in mind, the following result should not come as a surprise:

Theorem 7.3:

Let Γ be a finite dynamic game, and $s_i \in S_i$ a strategy for player i.

If s_i can rationally be chosen by a cautious "lexicographic" type t_i that expresses **common weak belief in rationality** (cf. Lecture 3),

then s_i can rationally be chosen by a "conditional type" that expresses **common initial belief in rationality**.

7.9 Dynamic games with perfect information

Informally, a dynamic game is said to be with **perfect information** if all players, whenever they have to choose, know exactly which choices have been made by their opponents until then.

Formally, a finite dynamic game Γ is of **perfect information** if every information set $h_i \in H_i$ consists of only one node.

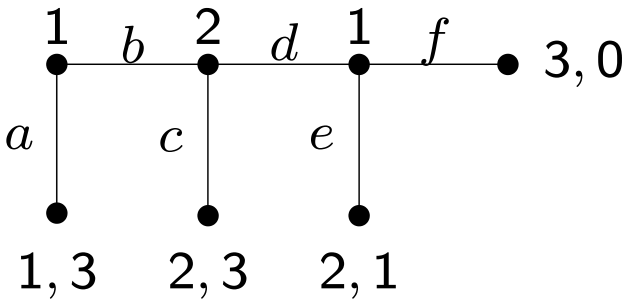

A dynamic game Γ with perfect information is **free of ties** if for every player i, every information set $h_i \in H_i$, every two choices $c_i, c_i' \in C_i(h_i)$, every terminal node z following c_i, and every terminal node z' following c_i', it holds that $u_i(z) \neq u_i(z')$.

Lemma 7.4: (Ben-Porath (1997))

Consider a finite dynamic game Γ with **perfect information** that is **free of ties.**

Then, a strategy s_i is **rational** (for some conditional belief)

if and only if

s_i is **not weakly dominated.**

We know: s_i is not weakly dominated if and only if s_i is rational for some cautious lexicographic belief.

Theorem 7.5:

Let Γ be a finite dynamic game with **perfect information** that is **free of ties.**

Then, a strategy s_i can rationally be chosen by a cautious “lexicographic” type t_i that expresses **common weak belief in rationality** (cf. Lecture 3),

if and only if

s_i can rationally be chosen by a “conditional type” that expresses **common initial belief in rationality**.

Theorem 7.6: (Ben-Porath (1997))

Let Γ be a finite dynamic game with **perfect information** that is **free of ties.**

Then, a strategy s_i can rationally be chosen under **common initial belief in rationality**

if and only if

s_i survives the **Dekel-Fudenberg procedure.**

7.10 Related Models

Reny (1992a) presents a **weakening** of sequential equilibrium **(Kreps and Wilson (1982)).**

The main idea in Reny's concept if **weak sequential rationality.**

Stated in terms of conditional beliefs, weak sequential rationality requires that a player **initially believes in the opponents' rationality**, but may believe later on that his opponents have made, and will make, irrational choices.

Reny (1992b, 1993) illustrates why common structural belief in rationality is in general impossible in dynamic games.

Hammond (1994) and **Halpern (2003)** formally investigate the relationship between **lexicographic beliefs** and **conditional beliefs.**

Hammond (1994) shows that the space of conditional beliefs is “equivalent” to the space of lexicographic beliefs with non-overlapping supports.

7.11 References

Elchanan Ben-Porath (1997): “Rationality, Nash equilibrium and backwards induction in perfect-information games”, Review of Economic Studies 64, 23-46.

Joseph Y. Halpern (2003): “Lexicographic probability, conditional probability, and nonstandard probability”, Available on his webpage: http://www.cs.cornell.edu/home/halpern/papers/lex.pdf

Peter J. Hammond (1994): “Elementary non-Archimedean representations of probability for decision theory and games”. In P. Humphreys (Ed.), Patrick Suppes: Scientific Philosopher; Volume 1. Dordrecht, Netherlands: Kluwer.

David Kreps and Robert Wilson (1982): “Sequential equilibria”, Econometrica 50, 863-894.

Philip J. Reny (1992a): “Backward induction, normal form perfection and explicable equilibria”, Econometrica 60, 627-649.

Philip J. Reny (1992b): "Rationality in extensive-form games", Journal of Economic Perspectives 6, 103-118.

Philip J. Reny (1993): "Common belief and the theory of games with perfect information", Journal of Economic Theory 59, 257-274.