Operations Research B:

Dynamic Optimization

and

Extensive Games

Joachim Rosenmüller

****

Lecture Notes
Sommer 2004

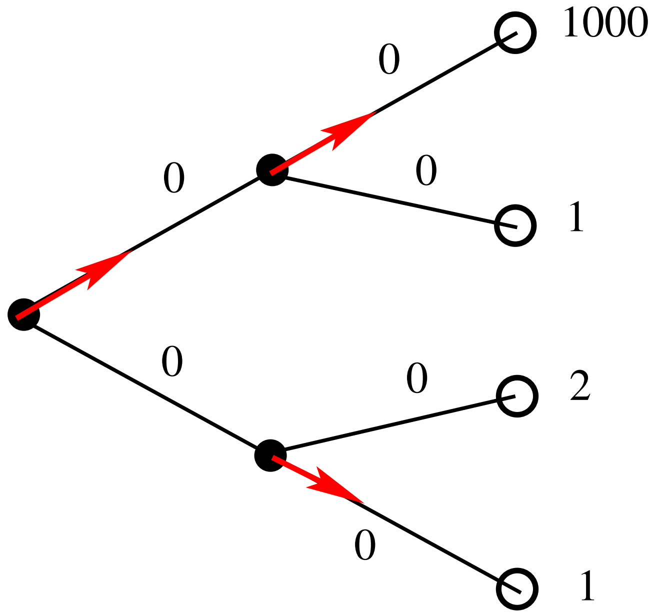

Contents

Chapter 1

The Principle of Dynamic Optimization

With this chapter we present the basic ingredients of the Optimality Priciple; the procedure based on this principle is sometimes called “backwards induction”. The model is the most simple one such that a dynamic process appears which admits of control by a decision maker and yields some reward depending on decisions made during the propagation of the process.

To be more precise, we wish to introduce three ingredients that more or less constitute a “controllable process”; The dynamic system, the incentive, and the concept of strategic behavior.

The dynamic system specifies the law of motion of some process of economic (or physical) nature. To some economically active individual (the decision maker, the controller, the player) there is a possiblity of influencing the motion of the process. Actions to controll the process are repeatedly feasible.

The incentive defines some reward or payoff available to the controller depending on the actions taken during time.

Finally, one has to develop a notion of “stategy”. A strategy is meant to be a well defined plan for long range control of the process, taking into account the possibly developments of motion, information on the past motion at some instant when intermediate action is possible and other relevant data. Action is instantaneous, strategy is long range and well planned.

1 Backwards Induction on a Graph

Within this section the dynamic process or controllable system is of a most simple nature. We start out with a ***state space*** $\mathcal{X}$ which is just a nonempty finite set, say $\mathcal{X} = \{\xi, \eta, \ldots, \kappa\}$.

A ***binary relation*** on $\mathcal{X}$ is a subset

$$\prec \ \subseteq \mathcal{X} \times \mathcal{X}. \tag{1}$$

We write $\xi \prec \eta$ instead of $(\xi, \eta) \in \ \prec$ and interprete this as "ξ is a predecesseor of η" or "η is a successor of ξ". The pair $(\mathcal{X}, \prec)$ is called a ***graph***. The elements of $\mathcal{X}$ are sometimes referred as ***nodes***.

Immediately we want the binary relation to satisfy certain requirements. E.g. we call $\prec$ ***asymmetric*** if, for some pair $\xi, \eta \in \mathcal{X}$ the relation $\xi \prec \eta$ implies that $\eta \prec \xi$ is *not* the case. This implies obviously that $\xi \prec \xi$ is *never* satisfied for some $\xi \in \mathcal{X}$.

A sequence

$$\boldsymbol{x} \ = \ (x_0, x_1, \ldots, x_T)$$

with $x_t \in \mathcal{X}$ $(t = 0, 1, \ldots, T)$ is a ***path*** in $\mathcal{X}$ (connecting x_0 and x_T) if each node (except the first one) is a successor of the previous one in the sense of the binary relation, i.e., if

$$x_{t-1} \prec x_t \ (t = 1, \ldots, T)$$

hols true. A path is a ***loop*** if the first and the last node coincide ($x_0 = x_T$).

Given some graph $(\mathcal{X}, \prec)$, we create a further binary relation written $\overset{\mathsf{T}}{\prec}$ and called the ***transitive hull*** of $\prec$ by the definition

$$\begin{array}{c} \xi \overset{\mathsf{T}}{\prec} \eta \\ \text{if and only if there exists a path} \\ \boldsymbol{x} \ = \ (x_0, x_1, \ldots, x_T) \\ \text{with } x_0 = \xi, \ \xi_T = \eta \ . \end{array} \tag{2}$$

We note that the transitive hull of an asymmetric binary relation is not necessarily asymmetric: if $\prec$ has loops, then the transitive hull is not asymmetric.

Now we come to our first

Definition 1.1. *A **tree** is a graph $(\mathcal{X}, \prec)$ with the following properties:*

1. *$\prec$ and $\overset{\mathsf{T}}{\prec}$ are asymmetric, thus, there are no loops..*
2. *There is a specific node, $\xi_0 \in \mathcal{X}$ called the **root** such that*

   (3) $$\xi_0 \overset{\mathsf{T}}{\prec} \xi \ (\xi \in \mathcal{X})$$

   holds true. That is, ξ_0 transitively precedes every other element. Hence, ξ_0 has no predecessor in view of the first requirement.
3. *For every $\eta \in \mathcal{X}, \ \eta \neq \xi_0$ we have*

   (4) $$|\{\xi | \xi \prec \eta\} = 1,$$

   that is, each node apart from the root has exactly one predecessor.

There is quite some structure defined for a graph to be a tree. We can typically imagine the shape of a tree to be represented by the following sketch (Figure 3.1). Here we tacitly assume that points further to the right and connected to some point immediately are successors. There are arcs of the graph or branches of the tree indicated, these objects do not occur in the formal set–up, yet they can at once be identified with pairs of nodes that are elements of the binary relation defining successors/predecessors.

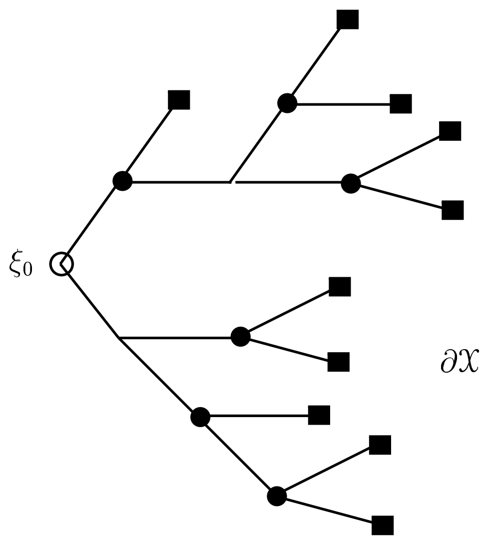

Figure 1.1: Graphical Representation of a Tree

Some useful notation is introduced as follows. We write

(5) $$\mathbf{V}(\eta) \ := \ \{\xi | \xi \prec \eta\}$$

for the set of predecessors of some node $\xi \in \mathcal{X}$,

$$\mathbf{N}(\eta) \ := \ \{\xi | \eta \prec \xi\} \tag{6}$$

and

$$\partial\mathcal{X} \ := \ \{\xi | \mathbf{N}(\xi) = \emptyset\} \tag{7}$$

for the set of points without successors, i.e., the ***boundary*** of $\mathcal{X}$.

Theorem 1.2. *For every $\xi \in \mathcal{X}$ there exists a unique path connecting the root ξ_0 and ξ.*

The simple proof will be omitted.

Theorem 1.3. *Let $\xi_0^\star \in \mathcal{X}$, and let*

$$\begin{aligned} \mathcal{X}^\star \ &:= \ \{\xi_0^\star\} \cup \{\xi \in \mathcal{X} | \xi_0^\star \overset{\mathsf{T}}{\prec} \xi\} \\ \prec^\star \ &:= \ \prec \cap (\mathcal{X}^\star \times \mathcal{X}^\star), \end{aligned} \tag{8}$$

then $(\mathcal{X}^\star, \prec^\star)$ is a tree.

Again the proof is omitted.

Definition 1.4. *The tree $(\mathcal{X}^{\xi_0^\star}, \prec^{\xi_0^\star}) \ := \ (\mathcal{X}^\star, \prec^\star)$ given in the context of Theorem 1.3 via formula (8) is called the **subtree** of $(\mathcal{X}, \prec)$ generated by $\xi_0^\star$.*

At this stage we have finished the definition of our dynamic system for the present purpose. We imagine that a “process” is equivalent to a path connecting the root with a boundary node. Such a path we may sometimes call a ***play***. The decision maker or player may, at each node, choose the successor, this way generating plays or “playing”. This is the complete description of a controllable process.

Why should someone take the trouble of controlling the process or generating a play? There should be an incentive provided by some “reward” or “utility”.

Definition 1.5. *Let $(\mathcal{X}, \prec)$ be a tree. A function*

$$f \ : \ \prec \to \mathbb{R} \tag{9}$$

*is called an **intermediate reward** function. A function*

$$u \ : \ \partial\mathcal{X} \to \mathbb{R} \tag{10}$$

*a **terminal reward** function.*

Thus, whenever the decision maker or player at some node ξ chooses to proceed with node $\eta \in \mathbf{N}(\xi)$, then he receives a (monetary) payment, utility, or reward $f(\xi, \eta)$. And finally, if the play ends at $\kappa \in \partial\mathfrak{X}$, then, in addition, he receives a reward $u(\kappa)$.

Of course we do not require the functions f and u to be necessarily nonnegative. A negative reward can be called a cost, fee, or disutility according to interpretation.

We close this part with

Definition 1.6. *Given a tree* $(\mathfrak{X}, \prec)$ *as well as reward functions* f *and* u*, the quadruplet*

$$\Sigma = (\mathfrak{X}, \prec, f, u) \tag{11}$$

is called a ***decision tree****.*

The decision tree suggests a time dependent interpretation of a process: this is just a path connecting the root and the boundary, or, as we have said, a play. The idea is that the decision maker, at each stage of the process, makes a decision by choosing a successor thus generating a play. Given such a play, the player can *evaluate* the total payoffs resulting. this idea is captured by the following

Definition 1.7. *Let* Σ *be a decision tree and let*

$$\underline{\overline{\mathbf{X}}} := \{\boldsymbol{x} = (x_0, x_1, \ldots, x_T) \mid \boldsymbol{x} \textit{ is a play in } \mathfrak{X}\}\,. \tag{12}$$

The ***evaluation*** *is the mapping*

$$\begin{aligned} \boldsymbol{C}^0 &= \boldsymbol{C}^{0x_0} : \underline{\overline{\mathbf{X}}} \to \mathbb{R} \\ \boldsymbol{C}^0(\boldsymbol{x}) &:= f(x_0, x_1) + f(x_1, x_2) + \ldots + f(x_{T-1}, x_T) + u_(x_T) \end{aligned} \tag{13}$$

This way we extend the reward function to a function on the plays by just summing up the total reward along a particular play.

Now, the next task is to present a precise meaning of the decision maker's/player's strategic behavior. This in effect should result in the definition of strategies and payoffs. In order to provide a simple approach first, we will tentatively assume that the player just chooses a path. This may not be too far off in the present context as the player controls the process at each node. Yet, eventually this idea of strategy cannot be maintained, but it suffices in order to describe the principle of dynamic optimization – which is our main purpose presently.

If, for the moment, we identify plays and strategies then his decision problem is clearly specified.

Definition 1.8. *1. Let Σ by a decision tree. The corresponding* ***normal form*** *(the decision problem in normal form) is the pair*

$$\Gamma^0 \;=\; \Gamma^0_\Sigma \;:=\; (\underline{\overline{\mathbf{X}}}, \boldsymbol{C}^0). \tag{14}$$

We refer to $\underline{\overline{\mathbf{X}}}$ as the ***alternatives*** *(available to the decision maker or player) and to $\boldsymbol{C}^0$ as the evaluation.*

2. The ***value*** *of Γ^0_Σ is given by*

$$v_0(\xi_0) \;:=\; max\left\{\boldsymbol{C}^0(\boldsymbol{x}) \,\middle|\, \boldsymbol{x} \in \boldsymbol{X}\right\} \tag{15}$$

3. A play (or strategy) $\bar{\boldsymbol{x}} \in \underline{\overline{\mathbf{X}}}$ is called ***optimal*** *if*

$$\boldsymbol{C}^0(\bar{\boldsymbol{x}}) \;=\; v_0(\xi_0) \tag{16}$$

holds true.

Thus, the player now faces an maximizatin problem: find the maximum as well as maximizers of a well defined function $\boldsymbol{C}^0$ on a (finite) set $\underline{\overline{\mathbf{X}}}$. The next problem is to find a procedure for the computation of a maximizer (an optimal path) and the maximum (the value). In course of this development it will immediately become clear why we connect the value of the normal form with the root of the tree within the notational convention.

At this stage the conceptual discussion is completed. The presentation of the solution procedure is now based on the simple

Principle of Optimality

which we first state verbally:

A path is optimal if and only if every subpath is optimal.

As there is presently no exact definition of a subpath available, we supply this verbally: a subpath is a remainder of a path connecting some node on the path to the boundary. Now we can we extend the value of the normal form in order to obtain a function on the nodes as follows.

Let $\xi \in \mathcal{X}$. For every path connecting ξ and the boundary

$$\boldsymbol{x} \;=\; (\xi, x_1, \ldots, x_S) \quad \text{with} \quad x_S \in \partial\mathcal{X}. \tag{17}$$

we define the evaluation quite analogously to Definition 1.7 via

$$\boldsymbol{C}^{0\xi}(\boldsymbol{x}) \;:=\; f(\xi, x_1) + f(x_1, x_2) + \ldots + f(x_{S-1}, x_S) + u(x_S) \tag{18}$$

The value "at ξ" is then obviously given by

$$v_0(\xi) \;:=\; \max \left\{ \boldsymbol{C}^{0\xi}(\boldsymbol{x}) \,\middle|\, \boldsymbol{x} \;\; \begin{array}{c} \text{is a path connecting } \xi \\ \text{and the boundary} \end{array} \right\}. \tag{19}$$

and an optimal path $\boldsymbol{x}$ "at ξ" is a maximizer in (19).

Now we have

Theorem 1.9. *Let Σ be a decision tree and let $\xi \in \mathcal{X}$ be a node. Then $\bar{\boldsymbol{x}} = (\xi, \overline{x}_1, \ldots, \overline{x}_t, \ldots, \overline{x}_T)$ is an optimal path at ξ if and only if, for every node $\overline{x}_t$ of $\bar{\boldsymbol{x}}$, the path $(\overline{x}_t, \ldots, \overline{x}_T)$ is optimal at x_t.*

The proof is quite obvious: if one could improve upon the payoff at some $\overline{x}_t$ by deviating from the remainder of the path $\bar{\boldsymbol{x}}$, then on could obviously improve the payoff for the path $\bar{\boldsymbol{x}}$ as well and vice versa. Yet, for the purpose of illustrating later developments, we wish to provide the formal proof at least partially.

Proof: Pick t and a node $\overline{x}_t$ of $\bar{\boldsymbol{x}}$. Assume that some subpath $\widetilde{\boldsymbol{x}} \;=\; (\overline{x}_t, \ldots, \overline{x}_T)$ is not optimal at $\overline{x}_t$. Then there is a path

$$\widehat{\boldsymbol{x}} \;=\; (\overline{x}_t, \widehat{x}_{t+1,}, \ldots, \widehat{x}_S)$$

satisfying

$$\boldsymbol{C}^{0\overline{x}_t}(\widehat{\boldsymbol{x}}) > \boldsymbol{C}^{0\widehat{x}_t}(\widetilde{\boldsymbol{x}}).$$

It follows that

$$\begin{aligned}
\boldsymbol{C}^{0\xi}(\bar{\boldsymbol{x}}) \;&=\; f(\xi, \overline{x}_1) + \ldots + f(\overline{x}_t, \overline{x}_{t+1}) + \ldots + f(\overline{x}_{T-1}, \overline{x}_T) + u(\overline{x}_T) \\
&=\; f(\xi, \overline{x}_1) + \ldots + C^{0\overline{x}_t}(\widetilde{\boldsymbol{x}}) \\
&<\; f(\xi, \overline{x}_1) + \ldots + C^{0\overline{x}_t}(\widehat{\boldsymbol{x}}) \\
&=\; f(\xi, \overline{x}_1) + \ldots + f(\overline{x}_t, \widehat{x}_{t+1}) + \ldots + f(\widehat{x}_{S-1}, \widehat{x}_S) \\
&=:\; \boldsymbol{C}^{0\xi}(\boldsymbol{x}^\star)
\end{aligned} .$$

Clearly the last expression is the payoff at some path $\boldsymbol{x}^\star$ connecting ξ and the boundary which contradicts the optimality of $\bar{\boldsymbol{x}}$. This proves that every subpath of an optimal path is optimal. The reverse direction of the proof is omitted.

q.e.d.

Remark 1.10. As we have seen in the context of Definition 1.4, any node $\xi^\star$ generates a subtree $(\mathfrak{X}^{\xi^\star}, \prec^{\xi^\star}) = (\mathfrak{X}^\star, \prec^\star)$ It is rather obvious that the restrictions of the reward functions, i.e.,

$$f^\star := f_{|\prec^\star} \qquad u^\star := u_{|\partial\mathfrak{X}} \tag{20}$$

together with the subtree define a decision tree

$$\Sigma^\star := \Sigma^{\xi^\star} := (\underline{\overline{\mathbf{X}}}^\star, \prec^\star, f^\star, u^\star) \tag{21}$$

with some obvious adjustment of notation for the set of plays connecting the root and the boundary in $\Sigma^\star$. There are now actually three possible interpretations of the function

$$\begin{aligned} v_0 \ &: \quad \underline{\overline{\mathbf{X}}} \to \mathbb{R} \\ v_0(\xi) \ &= \quad \max\left\{\boldsymbol{C}^{0,\xi}(\boldsymbol{x}) \,\middle|\, \boldsymbol{x} \text{ connects } \xi \text{ and the boundary } \partial\underline{\overline{\mathbf{X}}}.\right\} \end{aligned} \tag{22}$$

These are as follows:

- $v_0(\xi)$ reflects the maximal payoff of a path in Σ starting in ξ. This is the tentative approach we have favoured so far.
- $v_0(\xi)$ reflects the maximal total *remaining* payoff of a play in Σ that starts in ξ_0 and passes ξ.
- $v_0(\xi)$ is the value of Σ^ξ in the sense of definition 1.8.

The difference may seem to be subtle and not too relevant. Yet, one should consider the fact that formula (22) actually can be seen to define a function v_Σ and analogously can be used to define a version $v^\star = v_{\Sigma^\star}$ for every node $\xi^\star$ of the original tree. We manifest this idea by the following

Definition 1.11. *The function* $v_0 = v_{0\Sigma}$ *defined by* (22) *is the* ***value function*** *of* Σ.

Clearly, the above discussion leads to a whole family of value function, each one defined on a subtree generated by some node. Obviously the following is satisfied.

Corollary 1.12. *If Σ is a decision tree and, for some $\xi, \eta \in \mathcal{X}$, the relations*

$$\xi_0 \prec^{\mathsf{T}} \xi \prec^{\mathsf{T}} \eta$$

are satisfied, then

$$v_{0\Sigma\xi}(\eta) = v_{0\Sigma\xi_0}(\eta) = v_0(\eta) \tag{23}$$

is satisfied.

Now we come to the basic equation of "backwards induction."

Theorem 1.13 (The Optimality Equation). *Let Σ be a decision tree. Then the value function $v_0 : \overline{\underline{\mathbf{X}}} \to \mathbb{R}$ satisfies*

$$v_0(\xi) = \max_{\eta \in \mathbf{N}(\xi)} \{f(\xi, \eta) + v_0(\eta)\} \quad (\xi \in \mathcal{X} - \partial\mathcal{X}) \tag{24}$$

$$v_0(\xi) = u(\xi) \quad (\xi \in \partial\mathcal{X}) \ . \tag{25}$$

Moreover, v_0 is uniquely defined by (24) and (25).

We call (24) the ***Optimality Equation*** and (25) the ***boundary condition***.

There is a rather obvious interpretation attached to this Theorem.

The meaning of the boundary condition is rather obvious: the value of the decision tree that starts at a boundary node is the terminal payoff at this point as the process terminates immediately and no decision is asked for.

As for the optimality equation, suppose we know the value of the problem for all successors of some node ξ which is *not* located at the boundary. Then the value at ξ can be obtained by maximizing the one step payoff obtained by moving from ξ to a successor plus the value obtainable at this successor.

The argument shows that we can compute v_0 *inductively*, starting with the boundary nodes and moving backwards, that is, by ***backwards induction***.

Proof of Theorem 1.13 The last statement of the Theorem (uniquenes) is quite easy to check by induction: actually, the above interpretation provides the clue: v_0 is uniquely defined on $\partial\mathcal{X}$ as it equals u. Once v_0 is defined for all successors of some node, it is defined for this node itself by the optimality equation.

We prove that the value function as defined by (19) satisfies (24) and (25), again the last property is obvious. So we pick some $\xi \notin \partial\mathcal{X}$ and verify (24).

Let $\bar{\boldsymbol{x}} = (\xi, \overline{x}_1, \ldots, \overline{x}_T)$ be an optimal path for the start in ξ. In view of Theorem 1.9 (the optimality principle made precise), we know that $(\overline{x}_2, \ldots, \overline{x}_T)$ is optimal for start in $\overline{x}_1$, hence we obtain

$$
\begin{aligned}
v_0(\xi) &= f(\xi, \overline{x}_1) + f(\overline{x}_1, \overline{x}_2) + \ldots + f(\overline{x}_{T-1}, \overline{x}_T) + u(\overline{x}_T) \\
&= f(\xi, \overline{x}_1) + v_0(\overline{x}_1) \ .
\end{aligned} \tag{26}
$$

On the other hand we may consider an *arbitrary* path $(\xi, x_1, \ldots, x_S)$ connecting ξ and the boundary. Then we obtain

$$
v_0(\xi) \ \geq \ f(\xi, x_1) + f(x_1, x_2) + \ldots + f(x_{S-1}, x_S) + u(x_S). \tag{27}
$$

Consequently, for any $\eta \in \mathbf{N}(\xi)$, we have

$$
\begin{aligned}
v_0(\xi) \ &\geq \ f(\xi, \eta) + \max \{ f(\eta, x_2) + \ldots + f(x_{S-1}, x_S) + u(x_S) \mid \\
&\qquad\qquad \mid (\eta, x_2, \ldots, x_S) \text{ connects } \eta \text{ and } \partial\mathfrak{X} \} \\
&= \ f(\xi, \eta) + v_0(\eta) \ .
\end{aligned} \tag{28}
$$

In view of (26) we have actually an equation in (28), which proves indeed the optimality equation.

q.e.d.

Remark 1.14. Using the notation of the last proof, let again be $\bar{\boldsymbol{x}} = (\xi, \overline{x}_1, \ldots, \overline{x}_T)$ an optimal path for the start at ξ (which is not necessarily unique). As we have now an equation in (28) we conclude using (26) that

$$
\begin{aligned}
v_0(\xi) &= \max_{\eta \in \mathbf{N}(\xi)} \left(f(\xi, \eta) + v_0(\eta) \right) \\
&= f(\xi, \overline{x}_1) + v_0(\overline{x}_1)
\end{aligned} \tag{29}
$$

holds true, that is, $\overline{x}_1$ is a maximizer of (24). This we write

$$
\overline{x}_1 \in \operatorname*{argmax}_{\mathbf{N}(\xi)} \{ f(\xi, \bullet) + v_0(\bullet) \} \ =: \ \mathbf{M}_{\mathbf{N}(\xi)} \{ f(\xi, \bullet) + v_0(\bullet) \} \ . \tag{30}
$$

Hence, a successor of ξ at some optimal path starting at ξ is a maximizer of the optimality equation.

On the other hand, suppose that $\widehat{\eta} \in \mathbf{M}_{\mathbf{N}(\xi)} \{ f(\xi, \bullet) + v_0(\bullet) \}$ is a maximizer (29). Then there exists an optimal path $\widehat{\boldsymbol{x}} = (\xi, \widehat{\eta}, \widehat{x}_2, \ldots, \widehat{x}_S)$ such that $\widehat{\eta}$ is the first node following ξ. Indeed, if $(\xi, \eta, x_2, \ldots, x_R)$ is an arbitrary path starting at ξ, then we come up with

(31)

$$
\begin{aligned}
f(\xi,\eta) + f(\eta, x_2) + \ldots + f(x_{R-1}, x_R) + u(x_T) \leq & \; f(\xi,\eta) + v_0(\eta) \\
\leq & \; \max_{\eta'} f(\xi,\eta') + v_0(\eta') \\
= & \; f(\xi,\widehat{\eta}) + v_0(\widehat{\eta}) \\
= \; f(\xi,\widehat{\eta}) + f(\widehat{\eta},\widehat{x}_2) + \ldots + f(\widehat{x}_{S-1},\widehat{x}_S), &
\end{aligned}
$$

with a suitable optimal path $(\widehat{\eta}, \widehat{x}_2, \ldots, \widehat{x}_S)$ which starts at $\widehat{\eta}$.

∘ ∼∼∼∼∼∼∼ ∘

Combining our results we may state

Corollary 1.15. *The maximizers of the optimality equation are exactly the first nodes of optimal paths. More precisely, if, for some* $\xi \in \mathcal{X}$

$$\widehat{\eta} \in \;\; \boldsymbol{M}_{\mathbf{N}(\xi)} \left\{ f(\xi, \bullet) + v_0(\bullet) \right\}$$

holds true, then $\widehat{\eta}$ *together with an optimal path* $(\widehat{\eta}, \widehat{x}_2, \ldots\ldots, \widehat{x}_S)$ *yields an optimal path* $(\xi, \widehat{\eta}, \widehat{x}_2, \ldots\ldots, \widehat{x}_S)$

Remark 1.16. Combining these results, we can describe a procedure which yields the value function and an optimal path starting at the boundary and procceeding recursively “backwards”; this procedure is sometimes referred to as “backwards induction”.

1. By equation (25) we have $v_0(\xi) = u(\xi)$ for all ξ located in the boundary.

2. If v_0 has been determined on all nodes within $\mathbf{N}(\xi)$ for some ξ *not* on the boundary, then compute $v_0(x)$ by means of formula (24). In doing so, determine in particular a maximizer $\overline{\eta}$ of (24).

3. Given an optimal path that has been determined for the maximizer $\overline{\eta} \in \mathbf{N}(\xi)$, an optimal path for the start at ξ is obtained by augmenting this path with ξ.

There is an obvious graphical representation of this procedure as follows:

1. Represent the tree by a two–dimensional graph and the value function by a set of real numbers affixed to the boundary nodes.

2. Proceed to a predecessor of a boundary point. Choose a direction which yields a maximum of the values plus the intermediate reward and apply a mark (an arrow) in this direction. Fix the new value at this node and proceed likewise with all other nodes.

3. Once the root is reached, an optimal play is obtained by just following successively the arrows attached at each node.

◦ ∼∼∼∼∼∼∼ ◦

Example 1.17. Figure 1.2 indicates a tree and the reward functions. Final rewards are attached to boundary nodes and intermediate rewards, as they occur at transition from on node to a successor, are listed at the arcs of the graph.

Figure 1.3 describes the inductive computation of the value function as well as the choice of optimal successors. The value function at boundary points equals the terminal reward. When moving one step backwards, we consider we consider the sum of the final and the intermediate reward in each direction.We mark a maximizing direction with an arrow and attach the maximizing value to th node under consideration.

After finitely many steps we reach the root. The value of the decision problem is 31 and it is obtain by following the arrows from the root to a terminal node. That is, the "lowermost" path in the sketch is optimal.

◦ ∼∼∼∼∼∼∼ ◦

We consider the last example to provide a motivation to change our viewpoint concerning the structure of a global plan of decisions. So far we have identified plays and global plans. The example (and the previous discussion) however, suggest that the decision maker or player actually constructs global plans of another nature. The set of arrows in Figure 1.3 points to the following: a global plan, or as we shall say a ***strategy*** is complet set of local decisions deciding at each node how to proceed. Formally:

Definition 1.18. *Let Σ be a decision tree. A* ***strategy*** *or* ***control*** *is a mapping*

$$\begin{aligned} \boldsymbol{\alpha} \; : & \quad \mathfrak{X} - \partial\mathfrak{X} \to \mathfrak{X} \\ & \textit{satisfying} \\ \boldsymbol{\alpha}(\xi) \in & \quad \mathbf{N}(\xi) \quad (\xi \in \mathfrak{X} - \partial\mathfrak{X}) \; . \end{aligned} \tag{32}$$

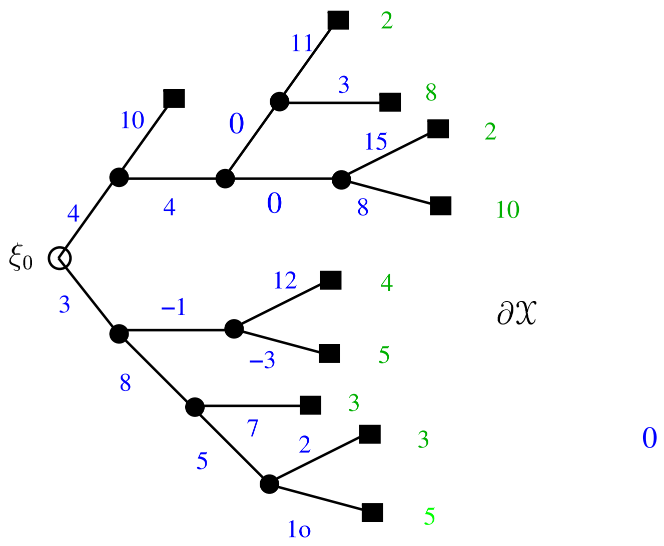

Figure 1.2: A decision tree

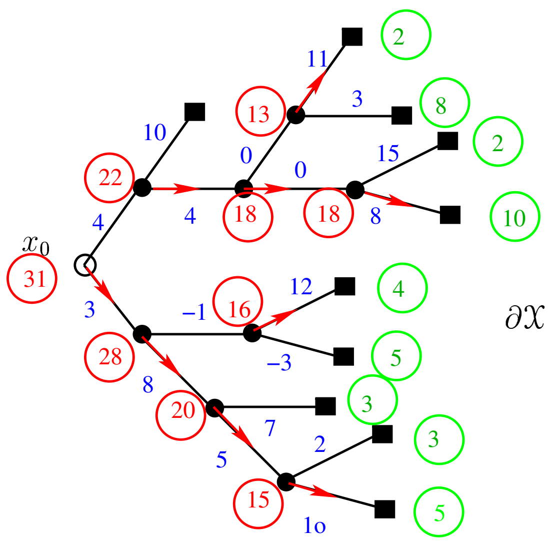

Figure 1.3: The value function and an optimal path

That is, a strategy is a set of arrows, one located at each node that requires an action.

Again we have to specify the payoff that results from a strategy. Clearly, any strategy $\boldsymbol{\alpha}$ induces a path (play)

$$X(\boldsymbol{\alpha}) \;:=\; (x_0, x_1, x_2, \ldots, x_T) \tag{33}$$

such that

$$x_t = \boldsymbol{\alpha}(x_{t-1}) \;\;, \; x_{t-1} \prec x_t \;\; (t = 1, \ldots, T), \quad x_T \in \partial\mathfrak{X} \tag{34}$$

holds true. Composing this with the evaluation (see Definition 1.7), we obtain

Definition 1.19. *Let Σ be a decision tree and denote by*

$$\mathfrak{S} \;:=\; \{\boldsymbol{\alpha} \mid \boldsymbol{\alpha} \textit{ is a strategy } \} \tag{35}$$

the set of strategies in the sense of Definition 1.18.

1. *The* ***payoff function*** *is the mapping*

   $$\boldsymbol{C}^{\xi_0}_{\bullet} \;:\; \mathfrak{S} \to \mathbb{R} \tag{36}$$

   given by

   $$\boldsymbol{C}^{\xi_0}_{\boldsymbol{\alpha}} \;:=\; \boldsymbol{C}^{0\xi_0}(X(\boldsymbol{\alpha})) \;\; (\boldsymbol{\alpha} \in \mathfrak{S}) \;. \tag{37}$$

2. *The* ***value*** *of the decision problem (relativ to $\mathfrak{S}$) is given by*

   $$v(x_0) \;:=\; max\left\{\boldsymbol{C}^{\xi_0}(\boldsymbol{\alpha}) \,\middle|\, \boldsymbol{\alpha} \in \mathfrak{S}\right\} \quad . \tag{38}$$

3. *A strategy $\bar{\boldsymbol{\alpha}}$ is* ***optimal*** *whenever it satisfies*

   $$\boldsymbol{C}^{0}(\bar{\boldsymbol{\alpha}}) \;=\; v(x_0) \quad . \tag{39}$$

Needless to say that the value actually yields the same result though it is defined via a different function. Also, we can repeat our discussion concerning the dependence on the coordinate x_0: there is actualy a value $v(\xi)$ defined for every $\xi \in \mathfrak{X}$ and it may be seen as a result obtained by starting from ξ as well as attached to a restricted decision problem Σ^{ξ}.

The ***The Priciple of Optimality***

is now rephrased as follows:

A strategy $\boldsymbol{\alpha}$ is optimal if and only if it is optimal for every subtree that is reached by the path $X(\boldsymbol{\alpha})$ generated.

There is a slight change in this version when compared to the naive one that concerned paths only. Note that i *cannot* be phrased to mean " ... in every subtree".

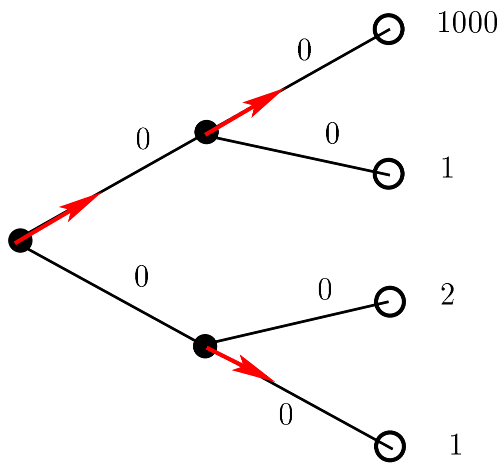

Figure 1.4: An Optimal Strategy

Therefore, albeit this means even more redundancy, we are going to provide a precise version of the Optimality Equation in the present context. To this end, observe that, for every $\xi^\star \in \mathfrak{X} - \partial\mathfrak{X}$ and the decision problem $\Sigma^{\xi^\star}$ generated thereby, we can also speak about the ***restriction of a strategy*** $\boldsymbol{\alpha}$ which is given by

$$\boldsymbol{\alpha}^\star \ := \ \boldsymbol{\alpha}\,\big|_{\mathfrak{X}^\star - \partial\mathfrak{X}^\star} \tag{40}$$

which is of course a strategy *beginning at* $\xi^\star$ or *in* $\Sigma^{\xi^\star}$. It is actually quite irrelevant whether, when dealing in the restricted graph, one considers $\boldsymbol{\alpha}^\star$ or $\boldsymbol{\alpha}$ as long as the impact on nodes in $\mathfrak{X}^\star$ is studied. Hence phrases like "$\boldsymbol{\alpha}$ is optimal for a start in ξ " or "$\boldsymbol{\alpha}$ is optimal in $\Sigma^\star$ " can be used equivalently. This permits us to state the following Version of Theorem 1.13

Theorem 1.20. *1. Let $\bar{\boldsymbol{\alpha}}$ be an optimal strategy for the start at $\xi \in \mathfrak{X}$ and let $\bar{\boldsymbol{x}} \ := \ X(\bar{\boldsymbol{\alpha}}) \ = \ (\xi, \overline{x}_1, \ldots, \overline{x}_T)$ be the path determined by $\bar{\boldsymbol{\alpha}}$ (as in* (33),(34)*). Then $\bar{\boldsymbol{\alpha}}$ is optimal for the start at every $\overline{x}_T$ $(t = 1, \ldots, T-1)$.*

2. *Let* $\boldsymbol{\alpha}$ *be a strategy and let* $x := X(\boldsymbol{\alpha}) = (x_0, x_1, \ldots, x_T)$ *be the path determined by* $\boldsymbol{\alpha}$*. If* $\boldsymbol{\alpha}$ *is optimal for the start at any node* x_t*, then* $\boldsymbol{\alpha}$ *is optimal*

3. *The function*

$$
\begin{aligned}
v &: \quad \mathfrak{X} \to \mathbb{R} \\
v(\xi) &:= \quad \max\left\{ \boldsymbol{C}^{\xi}_{\boldsymbol{\alpha}} \,\middle|\, \boldsymbol{\alpha} \in \mathfrak{S} \right\} \quad (\xi \in \mathfrak{X})
\end{aligned}
\tag{41}
$$

satisfies

$$
\begin{aligned}
v(\xi) &= \max_{\eta \in \mathbf{N}(\xi)} \{ f(\xi, \eta) + v(\eta) \} \quad (\xi \in \mathfrak{X} - \partial\mathfrak{X}) \\
v(\xi) &= u(\xi) \quad (\xi \in \partial\mathfrak{X})
\end{aligned}
\tag{42}
$$

4. *If, for every* $\xi \in \mathfrak{X} - \partial\mathfrak{X}$*, a maximizer of* (41) $\bar{\boldsymbol{\alpha}}(\xi)$ *is given,i.e., if*

$$
v(\xi) = f(\xi, \bar{\boldsymbol{\alpha}}(\xi)) + v(\bar{\boldsymbol{\alpha}}(\xi))
\tag{43}
$$

is satisfied, then the strategy $\bar{\boldsymbol{\alpha}}$ *defined this way is optimal* for the start at every $\xi \in \mathfrak{X} - \partial\mathfrak{X}$.

5. *Finally, if a strategy* $\bar{\boldsymbol{\alpha}}$ *is optimal for the start at every* $\xi \in \mathfrak{X} - \partial\mathfrak{X}$*, then, for every such* ξ *the successor* $\bar{\boldsymbol{\alpha}}(\xi)$ *is a maximizer in* (42)*, i.e., satisfies* (43)*.*

We do not offer a proof, this theorem reflects more or less a reformulation of the previous results in terms of strategies. Of course it shows that the value function v and v_0 are equal. One might argue that the machinery is much more comprehensive. Clearly, so is the reward: one gets a simultaneous solution and values for the whole family of decision problems. And the development shows, that there is a drift towards the definition of strategies; to deal with just the path concept is inapropriate.

We would like to repeat that there are essentially 3 steps to be taken in order to deal with a dynamic decision problem. Essentially, we follow the scheme:

$1^{st}STEP$: Set up the extensive form or dynamic model.

$2^{nd}STEP$: Define the concept of a strategy.

$3^{rd}STEP$: Set up the normal form.

2 Games on a Tree

Within this section the dynamic process is to be controlled by at least two players: we discuss dynamic games. Again we start out with a finite ***state space*** $\mathcal{X}$ and a ***binary relation*** $\prec$ on $\underline{\overline{\mathbf{X}}}$ which reflects the possible succession of states. The pair $(\mathcal{X}, \prec)$ is assumed to constitute a a ***tree***.

Now, as there are more that one individuals involved in governing the process, we introduce the ***set of players*** which we denote by $\boldsymbol{I} = \{1, \ldots, n\}$. This set reflects just an enumeration of independent decision entities which are supposed to have an incentive running with the process – by no means do we want to infer that "persons" are supposed to enter the model.

As a consequence, we assume thatat each node decisions for a successor should is assigned to one player. Moreover, the incentives to participate should vary with players, that is, we have reward functions defined for *each* player. Thus we end up with the following

Definition 2.1. *A* ***game tree*** *or game on a tree in* ***extensive form*** *for n persons is a five–tuple*

$$\Sigma \;=\; (\mathcal{X}, \prec; \boldsymbol{\iota}; f, u) \tag{1}$$

with the following data:

1. $(\mathcal{X}, \prec)$ *is a tree,*
2. $\boldsymbol{\iota} \;:\; \mathcal{X} - \partial\mathcal{X} \to \boldsymbol{I}$ *is mapping that assigns each node to a player, this mapping is called the* ***player assignment****. We assume that there is at least one node assigned to each player.*
3. $f \;=\; (f^i)_{i\in\boldsymbol{I}}$ *is a family of mappings*

   $$f^i \;:\; \prec \to \mathbb{R} \quad (i \in \boldsymbol{I}), \tag{2}$$

   the (family of) (intermediate) ***reward functions*** *for the players.*
4. *Finally,* $u \;=\; (u^i)_{i\in\boldsymbol{I}}$ *is a family of mappings*

   $$u^i \;:\; \partial\mathcal{X} \to \mathbb{R} \quad (i \in \boldsymbol{I}), \tag{3}$$

   the (family of) final or ***terminal reward functions*** *for the players.*

Remark 2.2. Recall that a binary relation $\prec$ is a formally a subset of $\mathcal{X} \times \mathcal{X}$, thus intermediate rewards are defined for each player $i \in \boldsymbol{I}$ and all pairs (ξ, η) with η succeding ξ to be $f^i(\xi, \eta)$.

Similarly, terminal payoffs are defined for each player i and every terminal node $\xi \in \partial\mathcal{X}$ to be $u^i(\xi)$. This is obviously a description of the situation quite analogous to the one in SECTION 1.

We can, therefore, attempt to sketch a tree game as previously by representing nodes as points in the plane and relations by arcs connecting nodes.

The rewards are reflected by simaltaneously noting all rewards for the players at each arc or terminal point. The following is sketch representing a tree game.

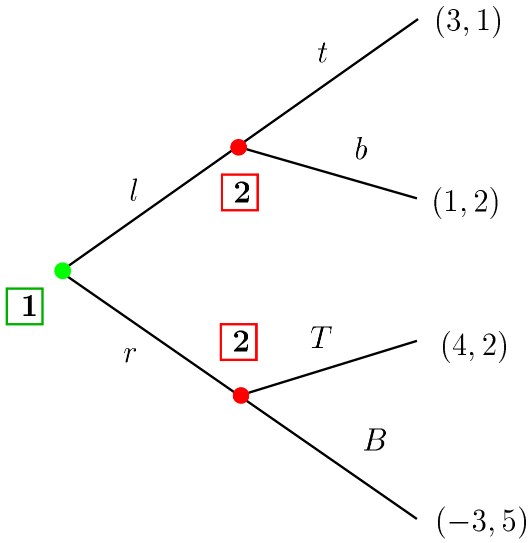

Figure 2.1: A Tree Game

In Figure 2.1 the successors of a node are identified by assigning labels to the corresponding arcs (like l for left, t for top or the like). These labels also refer to "decisions" or "actions" of the player "in charge".

◦ ∼∼∼∼∼∼∼ ◦

The assignment function obviously induces a partition of the nonterminal nodes according to assignment as follows:

Definition 2.3. *Let Σ be a game tree. The set of* ***player i's nodes*** *is given* *by*

$$\text{(4)} \qquad \mathcal{X}^i := \boldsymbol{\iota}^{-1}(\{i\}) = \{\xi \in \mathcal{X} - \partial\mathcal{X} \,|\, \boldsymbol{\iota}(\xi) = i\} \quad (i \in \boldsymbol{I}).$$

*The **player partion** is the partition of $\mathcal{X} - \partial\mathcal{X}$ given by*

$$\text{(5)} \qquad \underline{\underline{\mathcal{X}}} := \left(\mathcal{X}^i\right)_{i\in\boldsymbol{I}}.$$

We can speak of a partition proper as each player is assigned at least one node. This assumption we have indeed provided in Definition 2.1.

At this stage we have finished the first step: the dynamic decision problem, in this case the game tree, has been specified. Of course there is an interpretation that goes with the formal definition:

At each stage of the process, that is, at each node reached during the development, there is a player to whom the correponding node is assigned. This player has the right to choose a successor. In doing so, the player will keep in mind that he and all the opponents obtain an intermediate reward, but also that the process is moved one step ahead, thus certain alternatives are now outside of reach for anyone.

The succesive choice of nodes generates a path connecting the root and the boundary, or, as we have said, a play. Therefore, we have good reason to specify evaluations for the single players.

Definition 2.4. *Let Σ be a game tree or extensive form. Also, let*

$$\text{(6)} \qquad \overline{\underline{\mathbf{X}}} := \{\boldsymbol{x} = (x_0, x_1, \ldots, x_T) \,|\, \boldsymbol{x} \text{ is a play in } \mathcal{X}\}.$$

*For each $i \in \boldsymbol{I}$, the **evaluation** for player i is the mapping*

$$\text{(7)} \qquad \begin{aligned} \boldsymbol{C}^i &= \boldsymbol{C}^{i\xi_0} : \overline{\underline{\mathbf{X}}} \to \mathbb{R} \\ \boldsymbol{C}^i(\boldsymbol{x}) &:= f^i(x_0, x_1) + f^i(x_1, x_2) + \ldots + f^i(x_{T-1}, x_T) + u^i(x_T) \\ &=: F^i(\boldsymbol{x}) + u^i(x_T), \end{aligned}$$

the last row suggesting a shorthand notation for the total of all intermediate rewards and the final reward along a path. The index ξ_0 will be emphasized later on when we introduce subtrees as previously.

Now, the second step is to present a precise meaning of the player's strategic behavior. In view of our previous discussion in SECTION 1, we attempt the definition of a "strategy" to describe a global plan of decisions or actions, deciding at each node which is assigned to the player how to proceed.

Formally:

Definition 2.5. *Let Σ be a game tree. A (pure)* ***strategy*** *for player i is a mapping*

$$\begin{aligned} \boldsymbol{\alpha}^i \; : \quad & \mathfrak{X}^i \to \mathfrak{X} \\ \textit{satisfying} & \\ \boldsymbol{\alpha}^i(\xi) \in \quad & \mathbf{N}(\xi) \quad \xi \in \mathfrak{X}^i \;. \end{aligned} \tag{8}$$

Intuitively, a strategy for player i is a set of arrows, one located at each node that is assigned to player i.

Example 2.6. Figure 2.2 represents the simple tree that was introduced in Figure 2.1. Player 1 is assigned the root, and player 2 the other two non–terminal nodes.

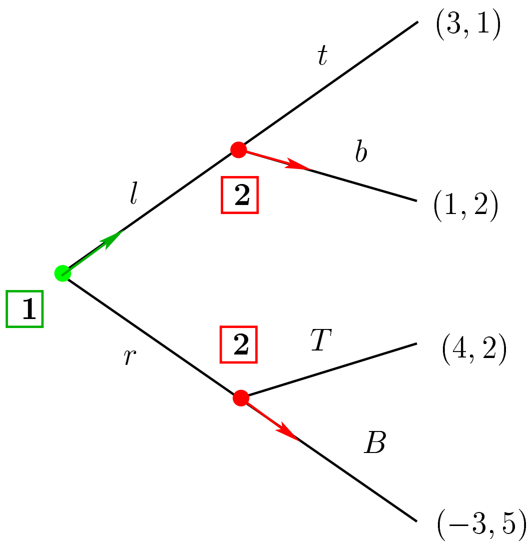

Figure 2.2: Strategies of 2 Players

Hence, a strategy of player 1 consists of just one decision or "action" at the root. As for player 2, a strategy is a pair of decisions or "actions" each one specified at one of the nodes that are assigned to him. Clearly, if player 1 decides to move l, player 2's decision to decide for B at the "lower" node is irrelevant. Yet, strategic behavior is modelled in way such that actions are assigned to all possible moves.

◦ ∼∼∼∼∼∼ ◦

Henceforth we will clearly distiguish between strategies as defined above and ***actions***. The latter ones are the local decisions of a player at some specified state, i.e., in the present context the choice of a particular successor at some particular node. Thus, a strategy is a set of actions, each one for one of the players nodes. In Figure 2.2 it so happens, the a strategy for player 1 coincides with an action while a strategy for player 2 is a set of two actions.

Now let us clarify the payoff that results from a set of strategies each one chosen from a player. Clearly, any strategy n–tuple $\boldsymbol{\alpha} = (\boldsymbol{\alpha}^1, \ldots, \boldsymbol{\alpha}^n)$ induces a path (play)

$$X(\boldsymbol{\alpha}) = X^{\xi_0}(\boldsymbol{\alpha}) := (x_0, x_1, x_2, \ldots, x_T) \tag{9}$$

such that

$$x_0 = \xi_0 \tag{10}$$

and

$$x_t = \boldsymbol{\alpha}^i(x_{t-1}) \;\;,\; x_{t-1} \prec x_t,\; x_{t-1} \in \mathfrak{X}^i, \;\; (t = 1, \ldots, T), \;\; x_T \in \partial\mathfrak{X} \tag{11}$$

holds true. Composing this with the evaluation (see Definition 2.7), we obtain

Definition 2.7. *Let Σ be a game tree and let*

$$\mathfrak{S}^i := \{\boldsymbol{\alpha}^i \mid \boldsymbol{\alpha}^i \textit{ is a strategy for player } i\} \quad (i \in \boldsymbol{I}) \tag{12}$$

denote the sets of strategies in the sense of Definition 2.5. Also, let

$$\mathfrak{S} = \mathfrak{S}^1 \times \ldots \times \mathfrak{S}^n$$

denote the set of strategy n–tuples.

*Then player i-s **payoff function** is the mapping*

$$\boldsymbol{C}^{i\xi_0}_{\bullet} : \mathfrak{S} \to \mathbb{R} \tag{13}$$

given by

$$\begin{aligned} \boldsymbol{C}^{i\xi_0}_{\boldsymbol{\alpha}} &:= \boldsymbol{C}^{i\xi_0}(X(\boldsymbol{\alpha})) \\ &= F^{i\xi_0}(X(\boldsymbol{\alpha})) + u^i(X_T(\boldsymbol{\alpha}) \\ &=: \boldsymbol{F}^{i\xi_0}(\boldsymbol{\alpha}) + \boldsymbol{u}^i_T(\boldsymbol{\alpha}) \;\; (\boldsymbol{\alpha} \in \mathfrak{S}) \;. \end{aligned} \tag{14}$$

The last notation is meant to distiguish between the intermediate and the final payoffs whenever we choose to employ $\boldsymbol{\alpha}$ as the argument, cf. Definition 2.7.

Now the principle difference between one person decision problems and multipersonal decision theory (Game Theory, that is) concerns the solution concept. In Optimization the techniques may be quite involved, the ideology is in general quite simple based on the idea of finding a maximum.

In Game Theory, the noncooperative solution concept is foremost the Nash Equilibrium which, however, can only be formulated in context with a game in normal form. At this stage of the discussion there emerges indeed a normal form: after we having specified strategies and the resulting payoffs we are clearly in the position to write down the "matrix game" resulting.

Definition 2.8. *Let Σ be game tree (or a tree game in extensive form). Let $\mathfrak{S}^i$ denote the set of (pure) strategies of player $i \in \boldsymbol{I}$ and let $\boldsymbol{C}^i_\bullet = \boldsymbol{C}^{i\xi_0}_\bullet$ be the payoff function as defined by* (13) *and* (14)*. Then*

$$\Gamma = \Gamma_\Sigma = \Gamma^{\xi_0}_\Sigma \;=\; (\mathfrak{S}^1, \ldots, \mathfrak{S}^n; \boldsymbol{C}^1_\bullet, \ldots, \boldsymbol{C}^n_\bullet) \tag{15}$$

the ***tree game in normal form*** *or the* ***normal form game*** *resulting from Σ.*

Note that there only finitely many (pure) strategies. Hence the normal form game is a matrix game.

Example 2.9. Let us return to Example 2.6.

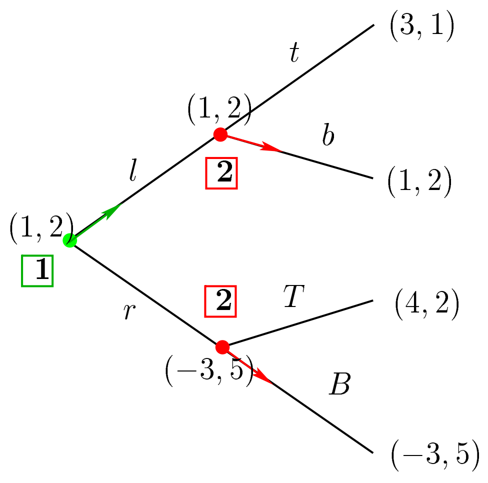

Figure 2.3: The Game Tree of Example 2.6

The strategies of player 1 are just his actions; those for player 2 are pairs of actions. Hence, the latter ones can be identified with pairs tT, tB, bT, bB.

(16)
$$\boldsymbol{C}^{1\xi_0}_{\bullet} = \begin{array}{c} \\ l \\ r \end{array} \begin{array}{c} \begin{array}{cccc} tT & tB & bT & bB \end{array} \\ \begin{pmatrix} 3 & 3 & 1 & \mathbf{1} \\ 4 & -3 & 4 & -3 \end{pmatrix} \end{array}$$

$$\boldsymbol{C}^{2\xi_0}_{\bullet} = \begin{array}{c} \\ l \\ r \end{array} \begin{array}{c} \begin{array}{cccc} tT & tB & bT & bB \end{array} \\ \begin{pmatrix} 1 & 1 & 2 & \mathbf{2} \\ 2 & 5 & 5 & 5 \end{pmatrix} \end{array}.$$

Computing the payoffs resulting we obtain the normal form corresponding to the game tree of Figure 2.5 represented by a matrix as in (16). Observe that the pair of (pure) strategies (l, bB) constitutes a Nash equilibrium. This is a well defined concept once we look at a noncooperative game "in normal form".

∘ ~~~~~~~ ∘

Remark 2.10. 1. Let Σ be a tree game and let $\xi \in \mathfrak{X}$. Then there is a well defined 'subgame' or 'remainder game' starting at ξ. The formal definition (similar to the one in SECTION 1) is rather obvious:

$\mathfrak{X}^{\xi}$ is the set of successors of ξ w.r.t. the transitive hull of $\prec$.

$\overset{\xi}{\prec}$ is the restriction of $\prec$ on $\mathfrak{X}^{\xi} \times \mathfrak{X}^{\xi}$.

$f^{i\xi}$ and $u^{i\xi}$ $(i \in \boldsymbol{I})$ are obtained as the restrictions of f^i on $\overset{i}{\prec}$ and of u^i on $\partial\mathfrak{X}^{\xi}$.

2. From this we infer that we obtain indeed a whole family of tree games (normal forms) Σ^{ξ} and a corresponding family of normal forms $\Gamma_{\Sigma^{\xi}}$.

3. Any strategy n-tuple $\boldsymbol{\alpha}$ may be restricted to Σ^{ξ} (so one should write $\boldsymbol{\alpha}^{\xi}$) and the payoffs received are the to be denoted by

$$\boldsymbol{C}^{i,\xi}_{\boldsymbol{\alpha}^{\xi}}$$

However, there are 'canonical' isomorphisms (as indicated in SECTION 1) and we shall simplify our arguments (and notations) by ignoring this notational obligation. Thus,

$$\boldsymbol{C}^{i,\xi}_{\boldsymbol{\alpha}}$$

will be used although $\boldsymbol{\alpha}$ is 'employed' only on vertices following ξ in the sense of $\overset{\mathsf{T}}{\prec}$.

4. If $\boldsymbol{\alpha}$ is a strategy n-tuple, then the payoffs of the present and the immediate future are connected via

$$C_{\boldsymbol{\alpha}}^{l,\xi} = C_{\boldsymbol{\alpha}}^{l,\boldsymbol{\alpha}^i(\xi)} + f^l(\xi, \boldsymbol{\alpha}^i(\xi)) \tag{17}$$

   for $i, l \in \mathcal{X}$ and $\xi \in \mathcal{X}^i$.

5. There is a family of normal form games $\Gamma_\Sigma^\xi = \Gamma_{\Sigma^\xi}$, each attached to some $\xi \in \mathcal{X}$.

∘ ∼∼∼∼∼∼∼ ∘

Definition 2.11 (Isaacs). *Let Σ be a game tree. Let*

$$\bar{\boldsymbol{\alpha}} = (\bar{\boldsymbol{\alpha}}^1, \ldots, \bar{\boldsymbol{\alpha}}^n)$$

*be an n-tuple of strategies. We shall say that $\bar{\boldsymbol{\alpha}}$ admits of **backwards induction** if there exists a family of functions.*

$$v^i : \mathcal{X} \to \mathbb{R} \quad (i \in \boldsymbol{I})$$

with the following properties

1. *For $i, l \in \boldsymbol{I}$ and $\xi \in \mathcal{X}^i$:*

$$\begin{aligned} v^i(\xi) &= \max_{\eta \in \mathbf{N}(\xi)} \{v^i(\eta) + f^i(\xi, \eta)\} \\ &= v^i(\bar{\boldsymbol{\alpha}}^i(\xi)) + f^i(\xi, \bar{\boldsymbol{\alpha}}^i(\xi)) \end{aligned} \tag{18}$$

$$v^l(\xi) = v^l(\bar{\boldsymbol{\alpha}}^i(\xi)) + f^l(\xi, \bar{\boldsymbol{\alpha}}^i(\xi)) \tag{19}$$

2.

$$v^l(\xi) = u^l(\xi) \quad (l \in I, \xi \in \partial\mathcal{X}) \tag{20}$$

The term "backwards induction" is rather a concept of dynamic programming or optimization. It seems that ISAACS was the first one to introduce this concept in the frame work of Game Theory if not at all.

We shall also use the term ***subgame–perfect*** when referRing to an equilibrium obtained by backwards induction. The latter term goes back to SELTEN and indicates indeed that the restriction to a subgame does not disturbe the equilibrium property. It is easy to see that there may occur non subgame–perfect equilibria. Example 2.12 shows such a situation.

As previously, we indicate a strategy by a set of arrows – however, a strategy for player i yields arrows only at the nodes in $\mathfrak{X}^i$. The family of functions $(v^i)_{i\in \boldsymbol{I}}$ as computed recursively, is indicated by a vector that appears at every node.

Example 2.12.

Figure 2.4 shows a simple game tree. The normal form is given by the

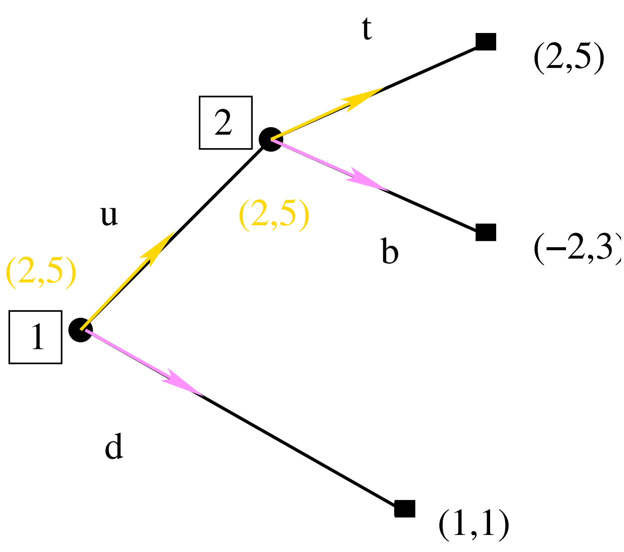

Figure 2.4: A Subgame–perfect Equilibrium and Another One

two matrices

$$
\begin{array}{rcl}
\boldsymbol{C}_{\bullet}^{1x_0} & = & \begin{array}{c} \\ u \\ d \end{array} \begin{array}{c} \begin{array}{cc} t & b \end{array} \\ \begin{pmatrix} 2 & -2 \\ 1 & 1 \end{pmatrix} \end{array} \\
 & & \\
\boldsymbol{C}_{\bullet}^{2x_0} & = & \begin{array}{c} \\ u \\ d \end{array} \begin{array}{c} \begin{array}{cc} t & b \end{array} \\ \begin{pmatrix} 5 & 3 \\ 1 & 1 \end{pmatrix} \end{array}.
\end{array}
\tag{21}
$$

The subgame perfect equilibrium in the upper left corners yields the payoffs $(2,5)$. There is a non subgame perfect equilibrium in the lower right corner which yields the payoffs $(1,1)$. This coreponds to player 1 moving "down" at once. Player 2 stabilizes this decision by threatening that he will choose "b" should he be called upon at all. Note that the threat is "not credible " as player 2, when called upon for a decision at his node, should opt for "t" when he intends to maximize his payoff. Player 1, anticipating this, may decide to move "u" even if the non subgame perfect equilibrium had been agreed upon ...

This is a frequent interpretation found for non subgame perfect equilibria: they are stabilized by threats, but the the threats are not credible when the plyer involved acts rationally, i.e., maximizes his payoff.

∘ ∼∼∼∼∼∼∼ ∘

We do have to provide a formal proof for the fact that backwards induction yields a Nash equilibrium. Of course this runs along the lines indicated in Section 1, yet one has to respect the multi player situation.

Theorem 2.13 (Zermelo–von Neumann–Kuhn). *Let Σ be a game tree and let $(\Gamma_{\Sigma^\xi})_{\xi\in\mathcal{X}}$ be the family of derived normal forms. Assume that a strategy n-tuple $\bar{\boldsymbol{\alpha}}$ admits of backwards induction. Then, for any $\xi \in \mathcal{X}$, it follows that $\bar{\boldsymbol{\alpha}}$ (or rather the restriction to Σ^ξ) is a Nash equilibrium in Γ_Σ^ξ. Moreover, the family v^i attached to $\bar{\alpha}$ via Definition 2.11 yields the equilibrium payoff, i.e.,*

$$
C_{\boldsymbol{\alpha}}^{i,\xi} = v^i(\xi) \ (i \in I, \xi \in \mathcal{X}). \tag{22}
$$

Proof:

By induction according to the maximal length of a paths in $\mathcal{X}$. Obviously nothing has to be proven if it so happens that

$$\mathfrak{X} = \{x_0\}$$

is true. We proceed by three steps as follows.

1$^{\mathbf{st}}$STEP :

Let $i \in I$ and consider $\xi \in \mathfrak{X}^i$. Then, for all $l \in I$ (including i!) we use (18) and (19) of Definition 2.11 in order to obtain

$$\begin{aligned} v^l(\xi) = \;& v^l(\bar{\boldsymbol{\alpha}}^i(\xi)) + f^l(\xi, \bar{\boldsymbol{\alpha}}^i(\xi)) \\ = \;& \boldsymbol{C}_{\bar{\boldsymbol{\alpha}}}^{l\bar{\boldsymbol{\alpha}}^i(\xi)} + f^l(\xi, \bar{\boldsymbol{\alpha}}^i(\xi)) \quad \text{(by induction hypothesis)} \\ = \;& \boldsymbol{C}_{\bar{\boldsymbol{\alpha}}}^{l\xi} \quad \text{(by Remark 2.10)} \end{aligned} \tag{23}$$

2$^{\mathbf{nd}}$STEP :

Now, in the same situation (i.e., for $i \in I, \xi \in \mathfrak{X}^i$), let player i deviate from $\bar{\boldsymbol{\alpha}}^i$ to $\widehat{\boldsymbol{\alpha}}^i \in \mathfrak{S}^i$.

Then we conclude that the following equations and inqualities hold true.

$$\begin{aligned} C_{\bar{\boldsymbol{\alpha}}}^{i,\xi} = \;& v^i(\xi) && \text{by the } 1^{st} \text{ step} \\ = \;& v^i(\bar{\boldsymbol{\alpha}}^i(\xi)) + f^i(\xi, \bar{\boldsymbol{\alpha}}^i(\xi)) && \text{by (23)} \\ \geq \;& v^i(\widehat{\boldsymbol{\alpha}}^i(\xi)) + f^i(\xi, \widehat{\boldsymbol{\alpha}}^i(\xi)) && \text{by (18) and as } \xi \in \mathfrak{X}^i, \\ = \;& C_{\bar{\boldsymbol{\alpha}}}^{i,\widehat{\boldsymbol{\alpha}}^i(\xi)} + f^i(\xi, \widehat{\boldsymbol{\alpha}}^i(\xi)) && \text{by the } 1^{st} STEP \\ \geq \;& C_{(\widehat{\boldsymbol{\alpha}}^i, \bar{\boldsymbol{\alpha}}^{-i})}^{i,\widehat{\boldsymbol{\alpha}}^i(\xi)} + f^i(\xi, \widehat{\boldsymbol{\alpha}}^i(\xi) && \text{by induction,} \\ & && \text{as } \bar{\boldsymbol{\alpha}} \text{ is an equilibrium in } \Gamma_{\Sigma\widehat{\boldsymbol{\alpha}}^i(\xi)} \\ = \;& C_{\widehat{\boldsymbol{\alpha}}^i, \bar{\boldsymbol{\alpha}}^{-i}}^{i,\xi} && \text{by Remark 2.10.} \end{aligned} \tag{24}$$

We have used the notation $(\widehat{\boldsymbol{\alpha}}^i, \bar{\boldsymbol{\alpha}}^{-i})$ for the strategy n–tuple resulting from $\bar{\boldsymbol{\alpha}}$ when player i switches to $\widehat{\boldsymbol{\alpha}}^i$. This proves indeed that player i does not improve his payoff.

3$^{\mathbf{rd}}$STEP :

Now the same has to be checked within the context that a player deviates who is not in charge of the next move at ξ, i. e., that we have $i \neq l \in I, \xi \in \mathfrak{X}^i$. Again we want to show that the inequality corresponding to (24) holds true, i.e.

$$C_{\bar{\boldsymbol{\alpha}}}^{l,\xi} \geq C_{\widehat{\boldsymbol{\alpha}}^l, \bar{\boldsymbol{\alpha}}-l}^{l\xi}.$$

However, if we compose (24) with the same line of equations and inequalities refering to l instead to i, then nothing changes but the first inequality (second line of (24)) which is an equation in view of (19),

q.e.d.

Corollary 2.14 (Zermelo–von Neumann–Kuhn). *Any tree game has Nash equilibria. Any zero–sum tree game has value and optimal strategies.*

Example 2.15. Again we return to Examples 2.6. and 2.9.

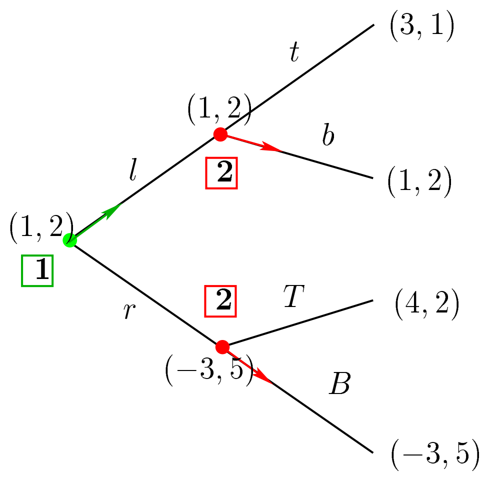

Figure 2.5: The Game Tree of Example 2.6

Recall that strategies are described by t, r for player 1 and pairs tT, tB, bT, bB for player 2.

$$\begin{aligned} \boldsymbol{C}^{1\xi_0}_{\bullet} &= \begin{array}{c} \\ l \\ r \end{array} \begin{array}{c} \begin{array}{cccc} tT & tB & bT & bB \end{array} \\ \begin{pmatrix} 3 & 3 & 1 & \mathbf{1} \\ 4 & -3 & 4 & -3 \end{pmatrix} \end{array} \\ \boldsymbol{C}^{2\xi_0}_{\bullet} &= \begin{array}{c} \\ l \\ r \end{array} \begin{array}{c} \begin{array}{cccc} tT & tB & bT & bB \end{array} \\ \begin{pmatrix} 1 & 1 & 2 & \mathbf{2} \\ 2 & 5 & 2 & 5 \end{pmatrix} \end{array}. \end{aligned} \tag{25}$$

Computing the payoffs resulting we obtain the normal form corresponding to the game tree of Figure 2.5 represented by a matrix as in (16). Observe that the pair of (pure) strategies (l, bB) constitutes a Nash equilibrium. This is a well defined concept once we look at a noncooperative game "in normal form".

∘ ∼∼∼∼∼∼∼ ∘

Remark 2.16. Let $n = 2$ and consider the case that

$$f^2 = -f^1 \,, \quad u^1 = -u^2$$

holds true. In this case we can say that the game tree has the ***zero sum*** property. It is then seen at once that $\Gamma_\Sigma = \Gamma^\xi_\Sigma$ is a ***zero sum game*** in pure strategies (or ***matrix game***).

Now it is sensible to omit the index $i = 1, 2$ whenever it simplifies the notation and no ambiguity is to be expected. Thus the terms f and u do not refer to a family of functions, instead we use the notation

$$f := f^1 = -f^2 \,, \quad u := u^1 = -u^2 \,.$$

Similarly, the payoff to player 1 is the only one that matters, we write

$$\boldsymbol{C}^\xi_\bullet := \boldsymbol{C}^{1\xi}_\bullet = -\boldsymbol{C}^{2\xi}_\bullet,$$

thus the normal form game is now

$$\Gamma = \Gamma_\Sigma = \Gamma^\xi_\Sigma = (\mathfrak{S}^1, \mathfrak{S}^2; \boldsymbol{C}^\xi_\bullet) \,. \tag{26}$$

A Nash equilibrium $\bar{\boldsymbol{\alpha}}$ is now tantamount to a ***saddlepoint***, i.e., it is characterized by a set of inequailities

$$\boldsymbol{C}_{\bar{\boldsymbol{\alpha}}^1\boldsymbol{\alpha}^2} \geq \boldsymbol{C}_{\bar{\boldsymbol{\alpha}}^1\bar{\boldsymbol{\alpha}}^2} \geq \boldsymbol{C}_{\boldsymbol{\alpha}^1\bar{\boldsymbol{\alpha}}^2} \quad ((\boldsymbol{\alpha}^1, \boldsymbol{\alpha}^2) \in \mathfrak{S}^1 \times \mathfrak{S}^2). \tag{27}$$

As is well known from the theory of matrix games, the existence of a Nash equilibrium or saddlepoint is equivalent with the existence of value and optimal strategies. In particular, we have

$$\begin{aligned} v_{\Gamma^\xi_\Sigma} &:= \operatorname*{val}_{\mathfrak{S}^1\times\mathfrak{S}^2} \boldsymbol{C}^\xi_\bullet \\ &= \max_{\mathfrak{S}^1} \min_{\mathfrak{S}^2} \boldsymbol{C}^\xi_\bullet = \min_{\mathfrak{S}^2} \max_{\mathfrak{S}^1} \boldsymbol{C}^\xi_\bullet \\ &= \boldsymbol{C}^\xi_{\bar{\boldsymbol{\alpha}}} = v(\xi) \,. \end{aligned} \tag{28}$$

Here, the first term is the value of the game Γ_Σ^x which is defined by the terms following subsequently. The last term refers to the value function which refers to player 1, the upper index 1 being omitted.

The equations in (28) are a consequence of Theorem 2.13, which establishes the existence of Nash equilibria as well as the equality of the equilibrium payoff (i.e. the value of the game) and the value function recursively defined.

∘ ∼∼∼∼∼∼∼ ∘

Corollary 2.17. *Let Σ be a game tree with zero sum property. Then, for every $\xi \in \mathcal{X}$, the normal form zero sum game Γ_Σ^ξ has value and optimal strategies. The value function (which is as well the value of $\Gamma_\Sigma^\bullet$ as a function in ξ, cf.* (28)*) satisfies the recursive equations*

1.

$$
\begin{aligned}
v(\xi) &= \max_{\eta \in \mathbf{N}(\xi)} \{v(\eta) + f(\xi, \eta)\} \\
&= v(\bar{\boldsymbol{\alpha}}^1(\xi)) + f(\xi, \bar{\boldsymbol{\alpha}}^1(\xi)) \quad (\xi \in \mathcal{X}^1),
\end{aligned} \tag{29}
$$

$$
\begin{aligned}
v(\xi) &= \min_{\eta \in \mathbf{N}(\xi)} \{v(\eta) + f(\xi, \eta)\} \\
&= v(\bar{\boldsymbol{\alpha}}^2(\xi)) + f(\xi, \bar{\boldsymbol{\alpha}}^2(\xi)) \quad (\xi \in \mathcal{X}^2),
\end{aligned} \tag{30}
$$

2.

$$
v(\xi) = u(\xi) \quad (\xi \in \partial\mathcal{X}), \tag{31}
$$

which in turn define optimal strategies $\bar{\boldsymbol{\alpha}}^1, \bar{\boldsymbol{\alpha}}^2$ for both players.

3 Dynamic Games: The Operator of Backwards Induction

Now we discuss the extensive form that it will appear in the particular shape of a dynamic game. We specify the model as previously: first of all, there is a model of the dynamic motion of a process. In the present version the time parameter is introduced explicitely.

We will define a state space as well as a set of possible actions for the players involved. Then there is the law transition, which, depending on the actions chosen by the players, determines the next state given the previous state.

As previously, $\boldsymbol{I} = \{1, \ldots, n\}$ represents the set of "players" or "individuals".

We start out with the

1st STEP : Definition of the extensive form.

1. $T > 0$ is an integer. T denotes the ***duration*** or ***horizon*** of the process. We introduce the (discrete) time intervals

$$\mathbf{T}_0 = \{0, \ldots, T\} \quad , \quad \mathbf{T} = \{1, \ldots, T\}$$

2. Next, $(\underline{\overline{\mathbf{X}}}_t)_{t \in \mathbf{T}}$ is a collection of finite sets. $\underline{\overline{\mathbf{X}}}_t$ is the state space for the process at time t. The Cartesian product

$$\underline{\overline{\mathbf{X}}} := \underline{\overline{\mathbf{X}}}_0 \times \cdots \times \underline{\overline{\mathbf{X}}}_T.$$

is the (global) **state space** containing all possible paths of a process.

3. For every $i \in \boldsymbol{I}$ and $t \in \mathbf{T}_1$, the finite set $\underline{\overline{\mathbf{Y}}}^i_t$ represents the possible ***actions*** of player $i \in \boldsymbol{I}$ at time t. We use abbreviations

$$\underline{\overline{\mathbf{Y}}}^i := \underline{\overline{\mathbf{Y}}}^i_0 \times \cdots \times \underline{\overline{\mathbf{Y}}}^i_T \quad (i \in \boldsymbol{I})$$
$$\underline{\overline{\mathbf{Y}}}_t := \underline{\overline{\mathbf{Y}}}^1_t \times \cdots \times \underline{\overline{\mathbf{Y}}}^n_t \quad (t \in \mathbf{T}),$$

and

$$\underline{\overline{\mathbf{Y}}} := \underline{\overline{\mathbf{Y}}}_0 \times \cdots \times \underline{\overline{\mathbf{Y}}}_T.$$

4. Next $\boldsymbol{a} = (\boldsymbol{a}_t)_{t\in\mathbf{T}}$ is a family of functions

$$\boldsymbol{a}_t \ : \ \underline{\overline{\mathbf{X}}}_{t-1} \times \underline{\overline{\mathbf{Y}}}_t \to \underline{\overline{\mathbf{X}}}_t$$

repesenting the ***law of motion***. We attach the following interpretation: Given a state of the process $\xi \in \underline{\overline{\mathbf{X}}}_{t-1}$ and an n–tuple $\eta = (\eta^1, \ldots, \eta^n) \in \underline{\overline{\mathbf{Y}}}_t$, of actions by all n players, the next state is

$$\boldsymbol{a}_t(\xi, \eta) \in \underline{\overline{\mathbf{X}}}_t.$$

Thus, $\boldsymbol{a}$ is supposed to describe the law of motion governing the process. Note that actions are thought of to take place "intermediately" between instant $t-1$ and instant t, yet the argument is associated with t – we write $\underline{\overline{\mathbf{Y}}}_t, y_t$ etc. These definitions so far describe the state space and the transition laws. Preferably one should combine the data into

$$\mathcal{X} \ := \ \left(\underline{\overline{\mathbf{X}}}\right)^{i\in\boldsymbol{I}}_{t\in\mathbf{T}_0}, \quad \mathcal{Y} \ := \ \left(\underline{\overline{\mathbf{Y}}}\right)^{i\in\boldsymbol{I}}_{t\in\mathbf{T}} \tag{1}$$

Now we turn to the rewards, incentives, or utilities provided for the players when they are jointly controlling the process.

1. For $i \in \boldsymbol{I}$ the function

$$f^i_t : \underline{\overline{\mathbf{X}}}_{t-1} \times \underline{\overline{\mathbf{Y}}}_t \longrightarrow \mathbb{R}.$$

denotes the ***intermediate reward function*** for player i at time t. When all players jointly choose the action n–tuple η and the process is in state $\xi \in \underline{\overline{\mathbf{X}}}_{t-1}$ at time $t-1$, then player i receives the utility $f^i_t(\xi, \eta)$. The complete family is denoted by

$$f = (f^i_t)_{i\in I,\ t\in\mathbf{T}}.$$

Similarly,

$$u = (u^i)_{i\in I}$$

is a family of functions

$$u^i : \underline{\overline{\mathbf{X}}}_T \longrightarrow \mathbb{R}$$

indicates the ***terminal*** or ***final reward*** to player i. Intuitively: If the process terminates at ξ, then player i's payoff is $u^i(\xi)$.

Definition 3.1. *A* ***dynamic*** n***-person game (in extensive form)*** *is a tuple*

$$\Sigma = (\mathcal{X}, \mathcal{Y}; \boldsymbol{a}, f; u, T),$$

the data being explained as above.

Again we introduce the ***evaluation***. This is the function

$$
\begin{aligned}
\boldsymbol{C}^i : \underline{\overline{\mathbf{X}}} \times \underline{\overline{\mathbf{Y}}} &\longrightarrow \mathbb{R} \qquad (i \in I) \\
\boldsymbol{C}^i(\boldsymbol{x}, \boldsymbol{y}) &= \sum_{t=1}^{T} f_t^i(x_{t-1}, y_t) + u^i(x_T)
\end{aligned} \tag{2}
$$

which evaluates the payoffs for player $i \in I$ along a path $x \in \underline{\overline{\mathbf{X}}}$ and is, therefore, sometimes called an ***evaluation***.

So far we have defined the data defining the "extensive form", in this case the dynamic game. As yet, there is no clear description of how the process is jointly controlled or rather "how the game is being played". The description involves a characterization of "strategic behavior" for the players. During the play they choose successively actions given they observe the state of the process. They cannot observe the actions sumultaneously chosen by their opponents.

As for previous actions of the opponents, it is not necessary to take any information into account: at some time instant all that matters with respect to the next transition and the intermediate reward is the present state – in which the previous actions may have been incorporated by whatever influence.

It is however not sufficient to just perform actions while the game is being played. A strategy should consider in advance all possible situations in which players may be called upon for an action. That is, a strategy is a complete plan of responses for all situations.

Thus, the

2$^{\text{nd}}$STEP :

consists in the definition of strategic behavior as follows.

Definition 3.2. *1. Let Σ be a dynamic game. A (pure)* ***strategy*** *for player $i \in I$ is a family of mappings*

$$
\begin{aligned}
\boldsymbol{\alpha}^i &= (\boldsymbol{\alpha}_1^i, \ldots, \boldsymbol{\alpha}_T^i) \\
\boldsymbol{\alpha}_t^i &: \underline{\overline{\mathbf{X}}}_{t-1} \to \underline{\overline{\mathbf{Y}}}_t^i.
\end{aligned}
$$

2. $\mathfrak{S}^i$ denotes the set of stragies of player i.

3. A strategy n–tuple is written $\boldsymbol{\alpha} = (\boldsymbol{\alpha}^1, \ldots, \boldsymbol{\alpha}^n)$ and $\mathfrak{S} = \mathfrak{S}^1 \times \ldots \times \mathfrak{S}^n$ is the set of all such n–tuples.

Some further definitions should clear up the way players insert strategies, thus generating a process which in turn yields rewards.

Definition 3.3. *1. We define a mapping*

$$(X^{\bullet}_{\star}, Y^{\bullet}_{\star}) \ : \ \underline{\overline{\mathbf{X}}}_0 \times \mathfrak{S} \to \underline{\overline{\mathbf{X}}} \times \underline{\overline{\mathbf{Y}}}, \tag{3}$$

successively as follows: given $\xi \in \underline{\overline{\mathbf{X}}}_0$ *and* $\boldsymbol{\alpha} \in \mathfrak{S}$, $X^{\xi}_{\boldsymbol{\alpha}} = (X_0, X_1, \ldots X_T)$ *and* $Y^{\xi}_{\boldsymbol{\alpha}} = (Y_1, \ldots Y_T)$ *are the paths defined via*

$$\begin{array}{rclrcl} & & & X_0 & := & \xi, \\ Y_1 & := & \boldsymbol{\alpha}_1(X_0) \ ; & X_1 & := & \boldsymbol{a}_1(X_0, Y_1) \\ & & \ldots & & & \\ Y_T & := & \boldsymbol{\alpha}_1(X_{T-1}) \ ; & X_T & := & \boldsymbol{a}_T(X_{T-1}, Y_T). \end{array} \tag{4}$$

The mapping $X^{\bullet}_{\star}$ *is referred to as the* ***process*** *associated with* Σ*; yet, sometimes we may refer to the mapping* $X^{\xi}_{\star}$ *as the process or even call a path* X^{ξ}_{α} *generated by this mapping (i.e., a path associated to some initial state* ξ *and some strategy* n*–tuple* $\boldsymbol{\alpha}$*) the process generated by* ξ *and* $\boldsymbol{\alpha}$*. The second mapping* $Y^{\bullet}_{\star}$ *that generates the sequence of action* n*–tuples is treated in a similar way.*

2. There is also a process $F^{i\bullet}_{\star}$ *of intermediate rewards (for player* $i \in \boldsymbol{I}$*). For some initial state* ξ *and a strategy* n*–tuple* $\boldsymbol{\alpha}$*, this is the sequence of real numbers obtained by*

$$\begin{array}{rcl} F^{i\xi}_{\boldsymbol{\alpha}} & = & (F^i_1, \ldots, F^i_T), \\ F^i_t & := & f^i_t \left(X^{\xi}_{\boldsymbol{\alpha}t-1}, Y^{\xi}_{\boldsymbol{\alpha}t} \cdot \right) \end{array} \tag{5}$$

Obviously the notation gets involved quickly, so we will omit the generating ingredients $\xi, \boldsymbol{\alpha}$ *whenever it seems to be obvious that a process is involved and not just a path. E.g. we also use*

$$U^{i\xi}_{\boldsymbol{\alpha}T} \ := \ u^i(X_T). \tag{6}$$

3. Finally, the ***payoff*** *function for player* i *is the mapping obtained by composing the evaluation and the process, i.e., we have*

$$
\begin{aligned}
\boldsymbol{C}^{i\xi}_{\boldsymbol{\alpha}} &= U^{i\xi}_{\boldsymbol{\alpha}T} + \sum_{t=1}^{T} F^{i\xi}_{\boldsymbol{\alpha}t} \\
&= u^i(X_T) + \sum_{t=1}^{T} f^i_t(X_{t-1}, \boldsymbol{\alpha}_t(X_{t-1})) \\
&= u^i(X_T) + \sum_{t=1}^{T} f^i_t(X_{t-1}, Y_t)) \\
&= \boldsymbol{C}^i(X^{\xi}_{\boldsymbol{\alpha}}, Y^{\xi}_{\boldsymbol{\alpha}})
\end{aligned} \tag{7}
$$

3rdSTEP :

We may now introduce the n-person noncooperative game ("in normal form") induced by the dynamic game Σ via the introduction of strategies.

Definition 3.4. *Given some $\xi \in \Xi_0$, The **normal form** corresponding to Σ is the noncooperative n person game*

$$
\Gamma := \Gamma_\Sigma := \Gamma_{\Sigma,\xi} := (\mathfrak{S}^1, \dots, \mathfrak{S}^n \; ; \; \boldsymbol{C}^{1\xi}_{\bullet}, \dots, \boldsymbol{C}^{n\xi}_{\bullet}). \tag{8}
$$

The obvious generalization of (8) is obtained by letting a process start at some $\xi \in \underline{\overline{\mathbf{X}}}_t$. The payoff for player i is then computed when rewards f^i_t are counted from time $t+1$ only. I.e., given a process (X, Y) define

$$
C^{i\xi}_{\boldsymbol{\alpha}} := u^i(X_T) + \sum_{s=t+1}^{T} f^i_s(X_{s-1}, Y_s) \tag{9}
$$

This time $C^{i\bullet}_{\boldsymbol{\alpha}}$ is a function on $\underline{\overline{\mathbf{X}}}_t$ and should, therefore, carry an index t (so that the index 0 would be reserved for the function defined by (7)). However, this index shall be omitted in order not to overburden our notation.

Now we are in the position to perform the **3rdSTEP** Set up the normal form:

Definition 3.5. *For $\xi \in \underline{\overline{\mathbf{X}}}_t$*

$$
\Gamma_\xi := (\mathfrak{S}^1, \dots, \mathfrak{S}^n; C^{1\xi}, \dots, C^{n\xi}) \tag{10}
$$

*is the **subgame** with initial state $\xi \in \underline{\overline{\mathbf{X}}}_t$.*

We are now interested in conditions for the existence of Nash equilibria. There is a marked difference here compared to the model in SECTION 2. In a tree game, each player commands certain nodes and hence solves a maximizing problem whenever he is called upon for an action. In the present model the players influence is simultneously evaluated. Hence, backwards induction requires Nash equilibria in each recursively computed "one shot game".

Further steps will lead to the recursive definition of a certain type of Nash equilibria. We assume the reader to be familiar with the concept of Nash equilibrium as such from any textbook dealing with matrix/bimatrix games. Nevertheless we shall repeat the definition within our present context. Nash equilibrium is the solution concept for noncooperative $n-$person games in normal form. The normal form we are presently dealing with is the game Γ_Σ derived from an extensive form game Σ, it seems appropriate to recall the definition within this context.

We typically need some notation which refers to

1. the deviation of one player in equilibrium and
2. the time parameter running backwards.

As to the first convention, we denote for any n–tuple or vector, say $\xi = (\xi^1, \ldots, \xi^n)$ the vector obtained by canceling the i –coordinate by

$$\xi^{-i} = (\xi^1, \ldots, \xi^{i-1}, \xi^{i+1}, \ldots, \xi^n)$$

In particular, given $\boldsymbol{\alpha} \in \mathfrak{S}$,

$$\boldsymbol{\alpha}^{-i} \in \mathfrak{S}^{-i} = \prod_{j \neq i} \mathfrak{S}^j$$

is the $n-1$–tuple of strategies of the opponents of player i and with some abuse of notation we write sometimes

$$\boldsymbol{\alpha} = (\boldsymbol{\alpha}^i, \boldsymbol{\alpha}^{-i})$$

This way we can formulate

Definition 3.6. *A* ***Nash equilibrium*** *for* $\Gamma = \Gamma_\Sigma^\xi$ *(in pure strategies) is an* $n-$*tuple* $\bar{\boldsymbol{\alpha}} = (\bar{\boldsymbol{\alpha}}^1, \ldots, \bar{\boldsymbol{\alpha}}^n) \in \mathfrak{S}$ *such that for any* $i \in I$ *and any* $\boldsymbol{\alpha}^i \in \mathfrak{S}^i$ *the inequality*

$$C_{\bar{\boldsymbol{\alpha}}}^{i,\xi} \geq C_{\boldsymbol{\alpha}^i, \bar{\boldsymbol{\alpha}}^{-i}}^{i,\xi} \tag{11}$$

holds true.

There is second convention useful for the formulation of "backwards induction". When time is imagined to run backwards, we write $\boldsymbol{\alpha}_{-s}$ instead of $\boldsymbol{\alpha}_{T-s}$ or $f^i_{-(t-1)}$ for $f^i_{T-(t-1)}$ etc. This may simplify the common notation of quantities that run forward (the data of the extensive form) and backwards (the value function) in time.

Now, if $\boldsymbol{\alpha}$ is a strategy n–tuple, then a new law of motion or transition function emerges naturally by the composition with $\boldsymbol{a}$, a precise version of this is

$$(12)\qquad \begin{aligned} \boldsymbol{a}^{\boldsymbol{\alpha}}_s &:= \underline{\overline{\mathbf{X}}}_{s-1} \to \underline{\overline{\mathbf{X}}}_s \\ \boldsymbol{a}^{\boldsymbol{\alpha}}_s(\xi) &:= \boldsymbol{a}_s\left(\xi, \boldsymbol{\alpha}_s(\xi)\right). \end{aligned}$$

In particular, if player i deviates from $\boldsymbol{\alpha}^i(\xi)$ to some η^i we write

$$(13)\qquad \begin{aligned} \boldsymbol{a}^{\eta^i\boldsymbol{\alpha}}_s &:= \underline{\overline{\mathbf{X}}}_{s-1} \to \underline{\overline{\mathbf{X}}}_s \\ \boldsymbol{a}^{\eta^i\boldsymbol{\alpha}}_s(\xi) &:= \boldsymbol{a}_s\left(\xi, (\eta^i, \boldsymbol{\alpha}^{-i}_s(\xi))\right). \end{aligned}$$

The generic term that appears in the Optimality Equation (the one shot backwards induction equilibrium) looks like

$$w\left(\boldsymbol{a}^{\boldsymbol{\alpha}}_{-(s-1)}(\xi)\right) + f^i_{-(s-1)}\left(\xi, \boldsymbol{\alpha}_{-(s-1)}(\xi)\right);$$

here w is a function that will play the role of a value function defined recursively and f^i_{-s} is player i's reward *when there are still s steps to go.* We regard this change of argument in w and the action implied by the addition of f^i_{-s} as an *Operator* acting on the function w. This motivates the definition of ***Operators*** defined on the space of all functions w, which we denote tentatively by $\mathcal{F}(\bullet)$, i.e.,

$$(14)\qquad \begin{aligned} \mathcal{O}^{i\boldsymbol{\alpha}}_s &: \mathcal{F}(\underline{\overline{\mathbf{X}}}_{-(s-1)}) \to \mathcal{F}(\underline{\overline{\mathbf{X}}}_{-s}) \\ \left(\mathcal{O}^{i\boldsymbol{\alpha}}_s w\right)(\xi) &:= w\left(\boldsymbol{a}^{\boldsymbol{\alpha}}_{-(s-1)}(\xi)\right) + f^i_{-(s-1)}\left(\xi, \boldsymbol{\alpha}_{-(s-1)}(\xi)\right) \quad (\xi \in \underline{\overline{\mathbf{X}}}_{-s}) \end{aligned}$$

Indeed, this operator carries a function defined on the state space when there are just $t-1$ steps to go one step backwards into a function when there are still t steps to go. In other words, the operator performs ***backwards induction***. Note that, suitably, the time argument runs backwards. Also, the direction of action is reversed. When one looks at (12), the mapping points into the future, the operator given by (14) moves functions into the past. This is also expressed by saying that the operator acts "contravariant".

A second version is needed whenever player i deviates as above from $\boldsymbol{\alpha}^i(\xi)$ to some η^i, i.e., we also introduce

$$(15)\qquad \begin{aligned} \mathfrak{A}_s^{i\eta^i\boldsymbol{\alpha}} &: \mathcal{F}(\underline{\overline{\mathbf{X}}}_{-(s-1)}) \to \mathcal{F}(\underline{\overline{\mathbf{X}}}_{-s}) \\ \left(\mathfrak{A}_s^{i\eta^i\boldsymbol{\alpha}} w\right)(\xi) &:= w\left(\boldsymbol{a}_{-(s-1)}^{\eta^i\boldsymbol{\alpha}}(\xi)\right) + f_{-(s-1)}^i\left(\xi, (\eta^i, \boldsymbol{\alpha}_{-(s-1)}^{-i}(\xi))\right). \end{aligned}$$

Some further versions will be used, e.g., $\mathfrak{A}_s^{i\eta}$ is clearly derived from equation (14) for the case that the strategy n–tuple $\boldsymbol{\alpha}_{-(s-1)}$ inserted is constant end equals η.

Now we are in the position to formulate the idea of equilibrium by backwards induction. In game theoretical context it has become customary to speak of a "subgame perfect" equilibrium (Selten).

Definition 3.7. *Let Σ be a dynamic n person game in extensive form and let $\bar{\boldsymbol{\alpha}}$ be a strategy n–tuple. We shall say that $\bar{\boldsymbol{\alpha}}$ is obtained by* ***backwards induction*** *or is* ***subgame perfect*** *if there is a family of functions*

$$(16)\qquad v_t^i : \underline{\overline{\mathbf{X}}}_{-t} \to \mathbb{R} \quad (i \in \boldsymbol{I},\ t \in \mathbf{T}_0)$$

satisfying the following equations:

$$(17)\qquad \begin{aligned} v_t^i(\xi) &= \max_{\eta^i \in \underline{\overline{\mathbf{Y}}}_{-(t-1)}^i} \left\{\left(\mathfrak{A}_t^{i\eta^i\bar{\boldsymbol{\alpha}}} v_{t-1}^i\right)(\xi)\right\} \\ &= \left(\mathfrak{A}_t^{i\bar{\boldsymbol{\alpha}}} v_{t-1}^i\right)(\xi), \quad (\xi \in \underline{\overline{\mathbf{X}}}_{-t}, t \in \mathbf{T}) \end{aligned}$$

and

$$(18)\qquad v_0^i(\xi) = u^i(\xi) \quad (\xi \in \underline{\overline{\mathbf{X}}}_T).$$

Remark 3.8. Of course, equations (18) reflect a family of boundary conditions. As for (17), this is the generalization of the "optimality equation" but we cannot speak of an "optimalit principle" in the present context.

Rather, we are dealing with a "one shot equilibrium condition" or a "recursive equilibrium condition". For short we may say that we refer to the ***Isaacs condition*** in order to mention an author who certainly delt with this problem rather early.

Indeed, there is a "one shot recursive game" to be solved at each stage of the recursive compuation suggested by the above definition. This is indicated by the recursive family of equations 17.

Indeed, consider the action space

$$\underline{\overline{\mathbf{Y}}}_{-(t-1)} = \underline{\overline{\mathbf{Y}}}^1_{-(t-1)} \times \ldots \times \underline{\overline{\mathbf{Y}}}^n_{-(t-1)}$$

and define tentatively a family of functions, one for each player on this space, say

$$\begin{aligned} H^i &= H^{i\xi} \; : \; \underline{\overline{\mathbf{Y}}}_{-(t-1)} \to \mathbb{R}, \\ H^i(\eta) &= v^i_{t-1}\left(\boldsymbol{a}_{-(t-1)}(\xi,\eta)\right) + f^i_{-(t-1)}(\xi,\eta) \\ &= \left\{ \mathcal{O}\!\!\mathcal{I}^{i\eta}_t v^i_{t-1} \right\}(\xi) \quad (\eta \in \underline{\overline{\mathbf{Y}}}_{-(t-1)}). \end{aligned} \tag{19}$$

Now the $2n$–tuple

$$\Gamma^{\xi}_{-(t-1)} \;:=\; \left(\underline{\overline{\mathbf{Y}}}^1_{-(t-1)}, \ldots, \underline{\overline{\mathbf{Y}}}^n_{-(t-1)}; H^1, \ldots, H^n\right) \tag{20}$$

formally defines an n person noncooperative game. Inserting the variables of $\bar{\boldsymbol{\alpha}}_{-(t-1)}(\xi)$ and $(\eta^i, \bar{\boldsymbol{\alpha}}^{-i}_{-(t-1)}(\xi))$ instead of η shows that $\bar{\boldsymbol{\alpha}}_{-(t-1)}(\xi)$ is a Nash equilibrium of this game. Such a Nash equilibrium does not necessarily exist (as we are arguing in terms of pure stragies only). Therefore, it is by no means clear that subgame perfect strategies exist at all.

This is in marked difference to the previous section. If we have a game tree and a node is assigned to a single player, then, at each stage one has just to solve a maximization problem.

However, if the existence problem can be delt with, then Definition 3.7 says that

1. $\Gamma^{\xi}_{-(t-1)}$ has pure Nash equilibria,
2. $\bar{\boldsymbol{\alpha}}_{-(t-1)}(\xi)$ is such a pure Nash equilibrium,
3. $v^i_t(\xi)$ is the payoff to player i at this Nash equilibrium.

Still in other words, a subgame perfect equilibrium exists if, at each stage one is able to solve the one shot recursive game. Note that uniqueness in general will not prevail, at each stage there may occur several Nash eqilibria or non at all.

∘ ∼∼∼∼∼∼ ∘

Definition 3.9 (The Zero Sum Case). *Let $n = 2$ and assume that (with a slight abuse of notation)*

$$-f_t^2 = -f_t^1 =: f_t, \quad u^2 = -u^1 =: u \tag{21}$$

*that is, the sum of rewards at each stage equals zero. Then Σ (and any subgame) is called a **zero–sum** game.*

Remark 3.10. 1. Then obviously

$$\boldsymbol{C}_{\boldsymbol{\alpha}}^{1\xi} + \boldsymbol{C}_{\boldsymbol{\alpha}}^{2\xi} = 0 \tag{22}$$

holds true for all arguments $\boldsymbol{\alpha}$ and ξ. Hence, every normalform Γ_Σ^ξ is a ***zero–sum*** two person game in normal form or a ***matrix game*** in pure strategies.

2. Also, it is not hard to see (by induction) that each one of the one shot recursive games discussed in the previous remarks is a zero–sum two person game. If such a game has an equilibrium in pure strategies, it has a value and optimal strategies.

3. Then the rather complicated situation expressed by the Isaacs condition is at once greatly simplified. To this end we introduce the value of the two person zero sum game

$$\Gamma_\star^\xi := \left(\underline{\overline{\mathbf{Y}}}_\star^1, \underline{\overline{\mathbf{Y}}}_\star^2; H\right)$$

for any function H defined on the cartesian product

$$\underline{\overline{\mathbf{Y}}}_\star := \underline{\overline{\mathbf{Y}}}_\star^1 \times \underline{\overline{\mathbf{Y}}}_\star^2.$$

This value is the quantity

(23)

$$v_{\Gamma_\star^\xi} := \underset{\eta\in\underline{\overline{\mathbf{Y}}}_\star}{\text{val}}\, H(\eta) := \max_{\eta^1\in\underline{\overline{\mathbf{Y}}}_\star^1}\min_{\eta^2\in\underline{\overline{\mathbf{Y}}}_\star^2} H(\eta^1,\eta^2) := \min_{\eta^2\in\underline{\overline{\mathbf{Y}}}_\star^2}\max_{\eta^1\in\underline{\overline{\mathbf{Y}}}_\star^1} H(\eta^1,\eta^2)$$

provided it exists (otherwise the operation *val* is not defined).

◦ ∼∼∼∼∼∼∼ ◦

Now we have

Theorem 3.11. *Let Σ be a dynamic two person zero–sum game. A pair of strategies $\bar{\boldsymbol{\alpha}}$ is subgame perfect if and only if there exists a family of functions*

$$v_t \ : \ \underline{\overline{\mathbf{X}}}_{-t} \to \mathbb{R} \quad (t \in \mathbf{T}_0) \tag{24}$$

with the following properties:

1.

$$\begin{aligned} v_t(\xi) &= \underset{\underline{\overline{\mathbf{Y}}}_{-(t-1)}}{\mathrm{val}} \left\{ v_{t-1}\left(\boldsymbol{a}_{-(t-1)}(\xi,\eta)\right) + f_{-(t-1)}(\xi,\eta) \right\} \\ &= \underset{\underline{\overline{\mathbf{Y}}}_{-(t-1)}}{\mathrm{val}} \left\{ \mathfrak{A}_t^{\eta} v_{t-1} \right\}(\xi) \quad (t \in \mathbf{T}) \quad , \end{aligned} \tag{25}$$

2.

$$v_0(\xi) = u(\xi). \tag{26}$$

3. $\bar{\boldsymbol{\alpha}}_t(\xi) \in \underline{\overline{\mathbf{Y}}}_{-(t-1)}$ *is a pair of optimal strategies of the two person zero–sum game*

$$\Gamma_{-(t-1)} \ = \ (\underline{\overline{\mathbf{Y}}}^1_{-(t-1)}, \underline{\overline{\mathbf{Y}}}^2_{-(t-1)}, \left\{ \mathfrak{A}_t^{\bullet} v_{t-1} \right\}(\xi) \tag{27}$$

Proof: From Remark 3.8 we know that $v_t(\xi)$ is the payoff to player 1 in an equilibrium (of $\Gamma_{-(t-1)}$. Since we have a zero–sum game, the payoff in any equilibrium is the value of the game and any equilibrium strategy is an optimal strategy.

Hence, the equations (25) and (26) are just a simultaneous reformulation of the two equations (17) and (18) in view of the zero sum property.

q.e.d.

We note a suggestive relation between the payoff functions defined at varying time instants; this will eventually lead to the techniques of backwards induction. Recall the definitions concerning the subgames given in the context of Definition 3.5. The payoff functions $\boldsymbol{C}_{\star}^{i\bullet}$ actually constitute a family defined on all sheafs of the state space. The following is formulated "in forward time" but will be used "backwards" as well.

Lemma 3.12. *Let Σ be a dynamic game and $\boldsymbol{\alpha}$ a strategy n–tuple. Then, for all $\xi \in \underline{\overline{\mathbf{X}}}_{t-1}$ it follows that*

$$\boldsymbol{C}_{\boldsymbol{\alpha}}^{i\xi} = \boldsymbol{C}_{\boldsymbol{\alpha}}^{i\boldsymbol{a}_t^{\boldsymbol{\alpha}}(\xi)} + f_t^i(\xi, \boldsymbol{\alpha}_t(\xi)) \tag{28}$$

Note that the left side regards a function defined on $\underline{\overline{\mathbf{X}}}_{t-1}$ while the right side deals with the version on $\underline{\overline{\mathbf{X}}}_t$. Now to the simple

Proof: Let $(X, Y) = (X^{\boldsymbol{\alpha}}, Y^{\boldsymbol{\alpha}})$ be the path or process generated by $\boldsymbol{\alpha}$ starting at $\xi \in \underline{\overline{\mathbf{X}}}_{t-1}$. Then we have

$$
\begin{aligned}
\boldsymbol{C}^{i\xi}_{\boldsymbol{\alpha}} &= \sum_{s=t}^{T} f^i_s(X_{s-1}, Y_s) + u^i(X_T) \\
&= \sum_{s=t+1}^{T} f^i_s(X_{s-1}, Y_s) + f^i_t(\xi, \boldsymbol{\alpha}_t(\xi)) + u^i(X_T)
\end{aligned}
\tag{29}
$$

as $X_{t-1} = \xi$ and $Y_t = \boldsymbol{\alpha}_t(X_{t-1})$. Because of $X_t = \boldsymbol{a}_t(X_{t-1}, Y_t) = \boldsymbol{a}_t(\xi, \boldsymbol{\alpha}_t(\xi)) = \boldsymbol{a}^{\boldsymbol{\alpha}}_t(\xi)$, we have indeed on the right side the term

$$
\sum_{s=t+1}^{T} f^i_s(X_{s-1}, Y_s) + u^i(X_T) = \boldsymbol{C}^{i\boldsymbol{a}^{\boldsymbol{\alpha}}_t(\xi)}_{\boldsymbol{\alpha}}
$$

q.e.d.

Theorem 3.13 (Zermelo–von Neumann–Kuhn). *Let Σ be a dynamic game and let $\left(\Gamma^{\xi}_{\Sigma}\right)_{\xi \in \underline{\overline{\mathbf{X}}}_t, t \in \mathbf{T}_0}$ be the family of derived normal forms. Let the strategy n-tuple $\bar{\boldsymbol{\alpha}}$ admit of backwards induction and let $(v^i_t)_{i \in \boldsymbol{I}, t \in \boldsymbol{T}}$ be the corresponding family of functions (cf. Definition 3.7). Then, for all $t \in \boldsymbol{T}$ and all $\xi \in \underline{\overline{\mathbf{X}}}_t$ it follows that $\bar{\boldsymbol{\alpha}}$ (or rather the restriction to Σ^{ξ}) is a Nash equilibrium in Γ^{ξ}_{Σ}. Moreover, the family $(v^i_t)_{i \in \boldsymbol{I}, t \in \boldsymbol{T}}$ yields the equilibrium payoff, i.e.,*

$$
\boldsymbol{C}^{i,\xi}_{\boldsymbol{\alpha}} = v^i_{-t}(\xi) \ (i \in I, \xi \in \underline{\overline{\mathbf{X}}}_t).
\tag{30}
$$

Proof:

By backwards induction rather analogously to the proof in Section 2. For the terminal points $\xi \in \underline{\overline{\mathbf{X}}}_T$ nothing has to be proved.

1stSTEP : Let $t \geq 1$ and let $\xi \in \underline{\overline{\mathbf{X}}}_{-t}$. Then we have

$$
\begin{aligned}
v^i_t(\xi) &= v^i_{t-1}\left(\boldsymbol{a}^{\bar{\boldsymbol{\alpha}}}_{-(t-1)}(\xi)\right) + f^i_{-(t-1)}\left(\xi, \bar{\boldsymbol{\alpha}}_{-(t-1)}(\xi)\right) \\
&\qquad \text{by (17),} \\
&= \boldsymbol{C}^{\boldsymbol{a}^{\bar{\boldsymbol{\alpha}}}_{-(t-1)}(\xi)}_{\bar{\boldsymbol{\alpha}}} + f^i_{-(t-1)}\left(\xi, \bar{\boldsymbol{\alpha}}_{-(t-1)}(\xi)\right) \\
&\qquad \text{by induction as } \boldsymbol{C}^{i\star}_{\bar{\boldsymbol{\alpha}}} \text{ equals } v^i_{\bullet}(\star) \text{ on } \underline{\overline{\mathbf{X}}}_{-(t-1)}, \\
&= \boldsymbol{C}^{i\xi}_{\bar{\boldsymbol{\alpha}}} \\
&\qquad \text{by Lemma 3.12.}
\end{aligned}
\tag{31}
$$

This way we have veryfied that $\boldsymbol{C}^{i\star}_{\bar{\boldsymbol{\alpha}}}$ equals $v^i_{\bullet}(\star)$ on $\underline{\overline{\mathbf{X}}}_{-t}$,

2^nd STEP : Now, suppose that player i deviates from his strategy $\bar{\boldsymbol{\alpha}}^i$ and chooses some $\widehat{\boldsymbol{\alpha}}^i$.

Then we have the following equations and inequalities.

$$
\begin{aligned}
\boldsymbol{C}^{i\xi}_{\bar{\boldsymbol{\alpha}}} &= v^i_t(\xi)) \\
&\qquad \text{as shown in the previous step} \\
&= v^i_{t-1}\left(\boldsymbol{a}^{\bar{\boldsymbol{\alpha}}}_{-(t-1)}(\xi)\right) + f^i_{-(t-1)}\left(\xi, \bar{\boldsymbol{\alpha}}_{-(t-1)}(\xi)\right) \\
&\qquad \text{as in the previous step by (17),} \\
&\geq v^i_{t-1}\left(\boldsymbol{a}^{\widehat{\boldsymbol{\alpha}}}_{-(t-1)}(\xi)\right) + f^i_{-(t-1)}\left(\xi, \widehat{\boldsymbol{\alpha}}_{-(t-1)}(\xi)\right) \\
&\qquad \text{by (17) again, as deviation in the one shot} \\
&\qquad \text{recursive game is not profitable for player } i, \\
&= \boldsymbol{C}^{\boldsymbol{a}^{\widehat{\boldsymbol{\alpha}}}_{-(t-1)}}_{\bar{\boldsymbol{\alpha}}}(\xi) + f^i_{-(t-1)}\left(\xi, \widehat{\boldsymbol{\alpha}}_{-(t-1)}(\xi)\right) \qquad (32) \\
&\qquad \text{by the first step, as } \boldsymbol{C}^{i\star}_{\bar{\boldsymbol{\alpha}}} \text{ equals } v^i_{\bullet}(\star) \text{ on } \underline{\overline{\mathbf{X}}}_{-(t-1)}, \\
&\geq \boldsymbol{C}^{\boldsymbol{a}^{\widehat{\boldsymbol{\alpha}}}_{-(t-1)}}_{\widehat{\boldsymbol{\alpha}}}(\xi) + f^i_{-(t-1)}\left(\xi, \widehat{\boldsymbol{\alpha}}_{-(t-1)}(\xi)\right) \\
&\qquad \text{by induction, as } \bar{\boldsymbol{\alpha}} \text{ is a Nash equilibrium} \\
&\qquad \text{for the start in } \underline{\overline{\mathbf{X}}}_{-(t-1)} \\
&= \boldsymbol{C}^{i\xi}_{\widehat{\boldsymbol{\alpha}}} \\
&\qquad \text{By Lemma 3.12}
\end{aligned}
$$

q.e.d.

Corollary 3.14. *Let Σ be a 2–person zero sum dynamic game and let $\left(\Gamma^{\xi}_{\Sigma}\right)_{\xi\in\underline{\overline{\mathbf{X}}}_t, t\in\mathbf{T}_0}$ be the family of derived zero sum normal form games . Assume that every one shot recursive game $\Gamma^{\xi}_{-(t-1)}$ has a value and optimal strategies. Then, for all $t \in \boldsymbol{T}$ and all $\xi \in \underline{\overline{\mathbf{X}}}_t$ it follows that Γ^{ξ}_{Σ} has a value and optimal strategies. The family $(v_t)_{t\in\boldsymbol{T}}$ yields the value of each of these games, i.e.,*

$$
v_t(\xi) = \operatorname*{val}_{\boldsymbol{\alpha}\in\mathfrak{S}^1\times\mathfrak{S}^2} \boldsymbol{C}^{\xi}_{\boldsymbol{\alpha}^1,\boldsymbol{\alpha}^2}. \tag{33}
$$

Moreover, the optimal strategies on each stage obtained by solving (33) *constitute the optimal strategies in each Γ^{ξ}_{Σ}.*

We now discuss the “one deviation principle”.

Definition 3.15. *A strategy n–tuple $\bar{\boldsymbol{\alpha}}$ satisfies the one deviation principle (**ODP**) if, for every $i \in \boldsymbol{I}$, every $t \in \mathbf{T}$, every $\xi \in \underline{\overline{\mathbf{X}}}_t$ and any strategy $\widehat{\boldsymbol{\alpha}}^i$ of player i satisfying*

$$\widehat{\boldsymbol{\alpha}}^i_s(\xi') \;=\; \bar{\boldsymbol{\alpha}}^i_s(\xi') \quad ((s,\xi') \neq (t,\xi)), \tag{34}$$

it follows that

$$\boldsymbol{C}_{\bar{\boldsymbol{\alpha}}}{}^{i}\xi \geq \boldsymbol{C}^{i\xi}_{\widehat{\boldsymbol{\alpha}}^i,\bar{\boldsymbol{\alpha}}^{-i}} \tag{35}$$

holds true.

That is, it is sufficient to test, at any subgame, for one deviation at the beginning of this subgame in order to ensure subgame perfectness.

Theorem 3.16. *A strategy n–tuple $\bar{\boldsymbol{\alpha}}$ is subgame perfect if and only if it satisfies ODP.*

Proof: Obviously a subgame perfect equilibrium satisfies ODP as it is restistan against *any* deviation of a player.

So we prove the converse direction by induction. For $\xi \in \underline{\overline{\mathbf{X}}}_T$ nothing has to be proved. Assume by induction that $\bar{\boldsymbol{\alpha}}$ (or rather the restriction to $\Sigma^{\xi'}$) is subgame perfect in any subgame $\Gamma^{\xi'}_{\Sigma}$ for all $\xi' \in \underline{\overline{\mathbf{X}}}_1$. We then have to show that $\bar{\boldsymbol{\alpha}}$ is a Nash equilibrium in Γ^{ξ}_{Σ} for all $\xi \in \underline{\overline{\mathbf{X}}}_0$.

Now, let $\widehat{\boldsymbol{\alpha}}^i$ be a deviation of player i from $\bar{\boldsymbol{\alpha}}^i$. Then, in any subgame we have by the induction hypothesis

$$\boldsymbol{C}^{i\xi'}_{\bar{\boldsymbol{\alpha}}} \geq \boldsymbol{C}^{i\xi'}_{\widehat{\boldsymbol{\alpha}}^i,\bar{\boldsymbol{\alpha}}^{-i}} \quad (\xi' \in \underline{\overline{\mathbf{X}}}_1). \tag{36}$$

Pick $\xi \in \underline{\overline{\mathbf{X}}}_0$ and define a strategy $\widetilde{\boldsymbol{\alpha}}^i$ for player i by

$$\begin{aligned} \widetilde{\boldsymbol{\alpha}}^i_s(\nu) &= \bar{\boldsymbol{\alpha}}^i_s(\nu) \quad (s,\nu) \neq (1,\xi) \\ \widetilde{\boldsymbol{\alpha}}^i_s(\xi) &= \widehat{\boldsymbol{\alpha}}^i_s(\xi) \end{aligned} \tag{37}$$

Then $\widetilde{\boldsymbol{\alpha}}^i$ shows just one deviation at $(1,\xi)$. We put $\widetilde{\boldsymbol{\alpha}} := (\widetilde{\boldsymbol{\alpha}}^i, \bar{\boldsymbol{\alpha}}^{-i})$ and $\widehat{\boldsymbol{\alpha}} := (\widehat{\boldsymbol{\alpha}}^i, \bar{\boldsymbol{\alpha}}^{-i})$ Then we obtain

$$\begin{aligned} \boldsymbol{C}^{i\xi}_{\bar{\boldsymbol{\alpha}}} &\geq \boldsymbol{C}^{i\xi}_{\widetilde{\boldsymbol{\alpha}}^i,\bar{\boldsymbol{\alpha}}^{-i}} \\ &= f^i(\xi, \widetilde{\boldsymbol{\alpha}}_1(\xi)) + \boldsymbol{C}^{i\widetilde{\boldsymbol{\alpha}}_1(\xi)}_{\widetilde{\boldsymbol{\alpha}}^i,\bar{\boldsymbol{\alpha}}^{-i}} \\ &= f^i(\xi, \widehat{\boldsymbol{\alpha}}_1(\xi)) + \boldsymbol{C}^{i\widehat{\boldsymbol{\alpha}}_1(\xi)}_{\bar{\boldsymbol{\alpha}}} \\ &\geq f^i(\xi, \widehat{\boldsymbol{\alpha}}_1(\xi)) + \boldsymbol{C}^{i\widehat{\boldsymbol{\alpha}}_1(\xi)}_{\widehat{\boldsymbol{\alpha}}^i,\bar{\boldsymbol{\alpha}}^{-i}} \\ &= \boldsymbol{C}^{i\xi}_{\widehat{\boldsymbol{\alpha}}^i,\bar{\boldsymbol{\alpha}}^{-i}}\ ; \end{aligned} \tag{38}$$

the last inequality follows from (36),

q.e.d.

4 Examples

We consider some examples, sometimes slightly augmenting the model.

Example 4.1. Consider the following zero–sum game. There are two states (payoff matrices) representing a "right of control" for each player respectively. The player in charge may either stick to this state or transfer the right of control to his opponent. Thus, the state spaces are

$$\underline{\overline{\mathbf{X}}}_0 \ = \ \underline{\overline{\mathbf{X}}}_1 \ = \ \ldots \ = \ \underline{\overline{\mathbf{X}}}_T \ = \ \{1,2\}$$

her $i \in \underline{\overline{\mathbf{X}}}_0$ reflects the state in which player i is in control. Next, the action spaces are

$$\underline{\overline{\mathbf{Y}}}_0^1 \ = \ \underline{\overline{\mathbf{Y}}}_0^2 \ = \ \underline{\overline{\mathbf{Y}}}_1^1 \ = \ \underline{\overline{\mathbf{Y}}}_1^2 \ = \ \ldots \ = \ \underline{\overline{\mathbf{Y}}}_T^2 \ = \ \{+,-\}.$$

We interprete + to mean (from the viespoint of a player) "I agree to switch the state" while − isregarded as "I disagree to switch the state".

Now we write down the transitions. If we fix a state $\xi = 1,2$ then $\boldsymbol{a}_0(1,\bullet,\bullet)$ is a function in (η_1,η_2), hence a 2×2 matrix. These matrices reflect the power of the player in charge, hence we put

$$\boldsymbol{a}_0(1,\bullet,\bullet) \ = \ \begin{array}{c} \\ + \\ - \end{array} \begin{array}{c} \begin{array}{cc} + & - \end{array} \\ \begin{pmatrix} 2 & 2 \\ 1 & 1 \end{pmatrix} \end{array}$$

(1)

$$\boldsymbol{a}_0(2,\bullet,\bullet) \ = \ \begin{array}{c} \\ + \\ - \end{array} \begin{array}{c} \begin{array}{cc} + & - \end{array} \\ \begin{pmatrix} 1 & 2 \\ 1 & 2 \end{pmatrix} \end{array} \quad .$$

That is, in state $\xi = 1$ (when player 1 is in control), the switch to player 2's control takes place whenever player 1 wishes to do so etc. Controlling the switch (and insisting to keept he state) may be costly. The intermediate rewards are given by $f_t \ = \ f_1 \ \ (t \in \mathbf{T})$ and

$$f_0(1,\bullet,\bullet) \ = \ \begin{array}{c} \\ + \\ - \end{array} \begin{array}{c} \begin{array}{cc} + & - \end{array} \\ \begin{pmatrix} 0 & 1 \\ -1 & 2 \end{pmatrix} \end{array}$$

(2)

$$f_0(2,\bullet,\bullet) \ = \ \begin{array}{c} \\ + \\ - \end{array} \begin{array}{c} \begin{array}{cc} + & - \end{array} \\ \begin{pmatrix} 0 & -1 \\ 2 & 1 \end{pmatrix} \end{array} \quad .$$

Finally, it is an advantage for player 1 to be in control at terminal time T. We put

$$u(1) := 6, \quad u(2) := 0. \tag{3}$$

This way we have defined a dynamic two person zero–sum game $\Sigma = (\underline{\overline{\mathbf{X}}}, \underline{\overline{\mathbf{Y}}}; \boldsymbol{a}; f, u; T)$ for arbitrary $T \in \mathbb{N}$.

∘ ∼∼∼∼∼∼∼ ∘

Chapter 2

Variants: Stopping, Discounting

Within this chapter we discuss variants of the standard modell devellopped so far. The Horizon may either be varriable inasmuch as it depends on the path evolving: we introduce (a naive version of) stopping times. Or else, the horizon of a controlled process is admitted to be infinite. In order to ensure convergence of the payoff, we then assume that payoffs in the future are discounted.

At this stage it is not clear that both versions are closely related; this can rreasonably be explained only in the presence of stochastic influence.

1 Stopping Times

We consider the case that the duration of the game or horizon is not fixed but depends on the development of the process or path controlled by the players.

2 Infinite Horizon, Discounted Payoffs

Within this section the dynamic process or controllable system does not terminate in finite time, the horizon T is now thought to be infinite.

To simplify matters, we also restrict ourselves to the *stationary* case, that is, state spaces are given by

$$\begin{array}{c} \underline{\overline{\mathbf{X}}}_0 = \underline{\overline{\mathbf{X}}}_1 = \underline{\overline{\mathbf{X}}}_2 = \dots \\ \underline{\overline{\mathbf{Y}}}_0 = \underline{\overline{\mathbf{Y}}}_1 = \underline{\overline{\mathbf{Y}}}_2 = \dots \end{array} \tag{1}$$

The same is true for transition and reward functions, i.e., we have

$$\begin{array}{c} \boldsymbol{a}_1 = \boldsymbol{a}_2 = \dots \;\; =: \;\boldsymbol{a} \\ f_1 = f_2 = \dots \;\; =: \; f \;, \end{array} \tag{2}$$

we omit the time index and use now$\boldsymbol{a}$ for the transition function instead for the full family etc. A final reward is not specified in the model with infinite horizon. Yet, for some intuitive considerations, we will also consider the finite horizon model with terminal payoffs.

We restrict our discussion to zero–sum games, thus omitting the upper index for the players. Simultaneously, we deal with optimization problems – all formalisms can be carried over by replacing, if suitable, the *val* operator by *max* or *min*.

The recursive equation for the value of the game is then given by

$$v_t(\xi) \;=\; \operatorname*{val}_{\eta \in \underline{\overline{\mathbf{Y}}}_1} \{v_{t-1}\left(a(\xi,\eta)\right) + f(\xi,\eta)\} \tag{3}$$

as has been discussed extensively in the previous section. This assumes of course an additional boundary or initial condition connecting the value function with the terminal or boundary payoff.

Chapter 3

Stochastic Influence, Incomplete Information and Behavioral Strategies

Within this chapter we introduce chance as a major source of influence on n–person games in extensive form.

1 Distributions on a Graph

Definition 1.1. *Let $(\mathcal{X}, \prec)$ be a graph without loops and circles.*

1. *The set of roots, i.e., nodes without predecessors is denoted by*

$$\mathcal{R} := \{\xi \in \mathcal{X} \mid \mathbf{V}(\xi) = \emptyset\} \tag{1}$$

2. *For $\xi \in \mathcal{X}$ let*

$$\underline{\overline{\mathbf{Z}}}_\xi = \{0, \ldots, a_\xi\} \neq \emptyset \tag{2}$$

   be a finite set. Let

$$\mathcal{Z} := \left(\underline{\overline{\mathbf{Z}}}_\xi\right)_{\xi \in \mathcal{X}}. \tag{3}$$

   Then

$$\Sigma_0 = (\mathcal{X}, \prec, \mathcal{Z}) \tag{4}$$

   *is called a **state enhanced network** or just a **network** with states.*

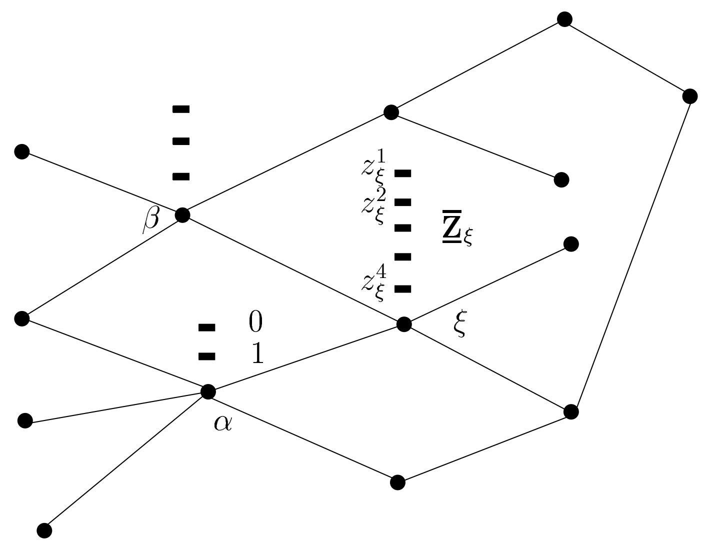

Figure 1.1: A network with states

Figure 1.1 suggests a network, the states are attached to the various nodes.

Example 1.2. We imagine that $\underline{\overline{\mathbf{Z}}}_\xi$ represents the possible states of a node $\xi \in \mathfrak{X}$. For instance, a signal or current may pass through the network and a node ξ may act as a switch which is either open or closed. Then we would take $\underline{\overline{\mathbf{Z}}}_\xi = \{0, 1\}$ in order to represent the possible two states of this switch.

Another possibility is that we consider

$$\underline{\overline{\mathbf{Z}}}_\xi = \{0, \ldots, a_\xi\}$$

adding the interpretation that node ξ has varying faulty states like

$$\begin{array}{rl} 0 & \text{stop – complete failure} \\ 1 & \text{failure of } 1^{st} \text{ kind} \\ 2 & \text{failure of } 2^{nd} \text{ kind} \\ \ldots & \ldots \\ a_\xi - 1 & \text{failure of } (a_\xi - 1)^{th} \text{ kind} \\ a_\xi & \text{failure of } a_\xi^{th} \text{ kind} \end{array} \tag{5}$$

∘ ∼∼∼∼∼∼∼ ∘

Recall that $\mathbf{V}(\xi) = \{\eta \mid \eta \prec \xi\}$ is the set of predecessors of $\xi \in \mathfrak{X}$.

Definition 1.3. *The Cartesian Product*

$$\underline{\overline{\mathbf{Z}}} := \prod_{\xi \in \mathfrak{X}} \underline{\overline{\mathbf{Z}}}_\xi \tag{6}$$

is the (global) ***state space*** *of the state network Σ_0. Similarly, for some set of predecessors $\mathbf{V} = \mathbf{V}(\xi)$ the set*

$$\underline{\overline{\mathbf{Z}}}_{\mathbf{V}} = \underline{\overline{\mathbf{Z}}}_{\mathbf{V}(\xi)} = \prod_{\eta \in \mathbf{V}} \underline{\overline{\mathbf{Z}}}_\eta \tag{7}$$

is the (local) ***parental state space*** *of $\xi \in \mathfrak{X}$. The elements of $\underline{\overline{\mathbf{Z}}}$ are denoted by $\boldsymbol{z} = (z_\xi)_{\xi \in \mathfrak{X}}$ and those of $\underline{\overline{\mathbf{Z}}}_{\mathbf{V}}$ are $\boldsymbol{z} = (z_\eta)_{\eta \in \mathbf{V}}$.*

Hence, a global state is a complete description of states for all nodes, a jointly determined state of the system. A parental state fixes the states of all predecessors of a node, it is a joint description of the relevant situation preceding the node.

Now we wish to model the situation, that there is stochastic influence issued from the parental states (the states of all predecessors) towards the state of some node. The idea is to specify, for each node ξ which is not a root a Markovian kernel which will be interpreted as the conditional probabilities, given a parental state for the states in $\underline{\overline{\mathbf{Z}}}_\xi$.

Definition 1.4. *Let* $\Sigma_0 = (\mathcal{X}, \prec, \mathcal{Z})$ *be a state network. Suppose that*

$$\mathbf{P}_\bullet = (\mathbf{P}_\xi)_{\xi \in \mathcal{X} - \mathcal{R}} \tag{8}$$

is a family of Markovian kernels

$$\mathbf{P}_\xi \mid \underline{\overline{\mathbf{Z}}}_{\mathbf{V}(\xi)} \Rightarrow \underline{\overline{\mathbf{Z}}}_\xi \ . \tag{9}$$

Then

$$\Sigma := (\mathcal{X}, \prec; \mathcal{Z}; \mathbf{P}_\bullet) \tag{10}$$

is called a ***causal network****.*

Example 1.5 (Burglary or Earthquake). The following sketch suggests a state network and some interpretations (Figure 1.2). We assume that each node ξ admits of two states, i.e.

$$\underline{\overline{\mathbf{Z}}}_\xi = \{0, 1\} \quad (\xi \in \mathcal{X}) \ .$$

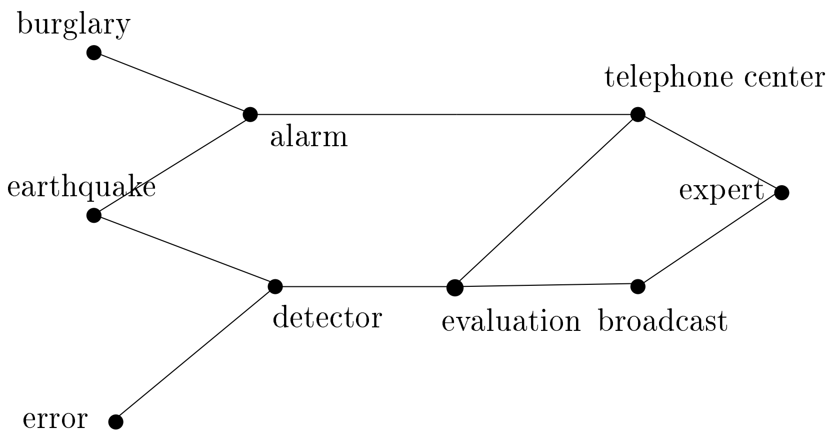

Figure 1.2: Burglary or earthquake ?

∘ ∼∼∼∼∼∼ ∘

Theorem 1.6. *Let* $(\mathcal{X}, \prec)$ *be a graph without loops and circles. Then there is* $T \in \mathbb{N}$ *and a bijective mapping*

$$\begin{array}{rcl} \{0, 1 \ldots, T\} & \to & \mathcal{X} \\ s & \to & \xi_s \end{array} \tag{11}$$

such that $\xi_s \prec \xi_t$ *implies* $s < t \ \ (s, t \in \mathbf{T}_0)$*. That is, we can find an enumeration such that no node is listed before any predecessor.*

Proof: The proof is obvious and works by induction. For $|\mathcal{X}| = 1$ nothing has to be done. For $|\mathcal{X}| > 1$ pick some $\xi \in \partial\mathcal{X}$ and take an enumeration of the graph obtained by cancelling ξ, i.e., of

$$\left(\mathcal{X} - \{\xi\}, \prec \big|\ _{\mathcal{X}-\{\xi\}}\right)$$

This constitutes a bijective mapping say

$$\{0, 1, \ldots, T-1\} \to \mathcal{X} - \{\xi\} \ .$$

Now put $\xi = \xi_T$. Then all predecessors of ξ do have lower indices in the enumeration, **q.e.d.**

We call an enumeration defined above ***compatible*** (with $\prec$). It is clear that, at any compatible enumertion, the set $\mathcal{R}$ of roots obtains the first indices; we write this

$$\text{(12)} \qquad \mathcal{R} \ = \ \{\xi \in \mathcal{X} \,|\, \mathbf{V}(\xi) \ = \ \emptyset\} \ = \ \{\xi_1, \ldots \xi_r\}$$

where $r \ := \ |\mathcal{R}|$ is the number of roots.

Definition 1.7. *Let*

$$\Sigma \ := \ (\mathcal{X}, \prec; \mathcal{Z}; \mathbf{P}_\bullet)$$

be a causal network and let $(\boldsymbol{\mu}^\xi)_{\xi \in \mathcal{R}}$ be a family of probabilities, each one defined on the set $\overline{\underline{\mathbf{Z}}}_\xi$ for some $\xi \in \mathcal{R}$. Let $\mathbf{T}_0 \to \mathcal{X}$ be a compatible enumeration.
*The **Markovian measure** generated by $\boldsymbol{\mu}^\bullet$ (and $\mathbf{P}_\bullet$) is the ($\sigma-$)additive set function $\boldsymbol{m} \ = \ \boldsymbol{m}^{\mathbf{P}_\bullet}_{\boldsymbol{\mu}^\bullet}$ defined on $\overline{\underline{\mathbf{Z}}}$ via*

$$\text{(13)} \qquad \begin{aligned} \boldsymbol{m}_{\boldsymbol{z}} \ &= \ \boldsymbol{m}(\{\boldsymbol{z}\}) \ = \ \boldsymbol{m}\left(\left\{(z_\xi)_{\xi \in \mathcal{X}}\right\}\right) \\ &:= \ \boldsymbol{\mu}^{\xi_0}_{z_{\xi_0}} \boldsymbol{\mu}^{\xi_1}_{z_{\xi_1}} \ldots \boldsymbol{\mu}^{\xi_r}_{z_{\xi_r}} \mathbf{P}_{\xi_{r+1}}(\boldsymbol{z}_{\mathbf{V}(\xi_{r+1})}, z_{\xi_{r+1}}) \cdots \mathbf{P}_{\xi_T}(\boldsymbol{z}_{\mathbf{V}(\xi_T)}, z_{\xi_T}) \\ &\qquad\qquad (\boldsymbol{z} \in \overline{\underline{\mathbf{Z}}}) \ . \end{aligned}$$

*$\boldsymbol{\mu}^\bullet$ is called the **initial distribution**.*

Theorem 1.8. *The Markovian measure is a probability and it is independent on the compatible enumeration.*

The **Proof** is omitted as it is the same as for standard Markovian chains. One just has to perform the summation successively, observing that every kernel is a probability in the second coordinate.

Definition 1.9. *Let* $(\Omega, \underline{\underline{\mathbf{F}}}, \mathbb{P})$ *be a probability space. Also, let* Σ *be a causal network and let* $\boldsymbol{\mu}^{\bullet}$ *be a probability on the roots of* $(\mathcal{X}, \prec)$. *A random variable* $Z : \Omega \to \overline{\underline{\mathbf{Z}}}$ *is called a* ***causal process*** *if its distribution is the Markovian measure generated by* $\boldsymbol{\mu}^{\bullet}$ *and* $\mathbf{P}_{\bullet}$, *i.e. if the measure obtained by transformation with* Z *is*

$$Z\,\mathbb{P} \;=\; \boldsymbol{m}^{\mathbf{P}_{\bullet}}_{\boldsymbol{\mu}^{\bullet}}. \tag{14}$$

Remark 1.10. Note that a causal process can also be seen as a collection of random variables, i.e.,

$$Z \;=\; (Z_{\xi})_{\xi\in\mathcal{X}} \quad Z_{\xi} \;:\; \Omega \to \overline{\underline{\mathbf{Z}}}_{\xi}\ .$$

The point is of course that the probability $\boldsymbol{m}^{\mathbf{P}_{\bullet}}_{\boldsymbol{\mu}^{\bullet}}$ specifies the *joint* distribution of all the random variables involved.

∘ ~~~~~~ ∘

Corollary 1.11. *Let* Z *be a causal process. Then we have*

$$\mathbb{P}(Z_{\xi} = z_{\xi}) = \boldsymbol{\mu}^{\xi}(\{z_{\xi}\}) = \mu^{\xi}_{\{z_{\xi}\}} \quad (\xi \in \mathcal{R}) \tag{15}$$

and

$$\mathbb{P}\left(Z_{\xi} = z_{\xi} \mid Z_{\mathbf{V}} = z_{\mathbf{V}}\right) = \mathbf{P}_{\xi}\left(z_{\mathbf{V}}, z_{\xi}\right) \quad (\xi \notin \mathcal{R},\ \mathbf{V} = \mathbf{V}(\xi)) \tag{16}$$

The ***Proof*** *is the same as for Markovian processes.*

Thus, by specifying conditional probabilities and initial distributions we can construct Markovian measures and processes.

Example 1.12 (Generator and Inhibitor). Let $\mathcal{X} = \{\alpha, \beta, \xi\}$ be such that α and β are the predecessors of ξ. The state spaces are all $\{0, 1\}$ and can be interpreted as switches or clean and faulty states. (Figure 1.3)

We abbreviate the random variables involved by

$$\begin{array}{rcll} U & := & Z_{\alpha} & \text{the generator} \\ I & := & Z_{\beta} & \text{the inhibitor} \\ X & := & Z_{\xi} & \text{the indicator} \end{array} \tag{17}$$

and we wish to model the situtation that *without any disturbance* the generator directly influences the state of the indicator (the observed instrument) but may be inhibited by some hidden disturbance, the inhibitor.

The tacit assumption is that there are random variables X_1 and X_2 such that

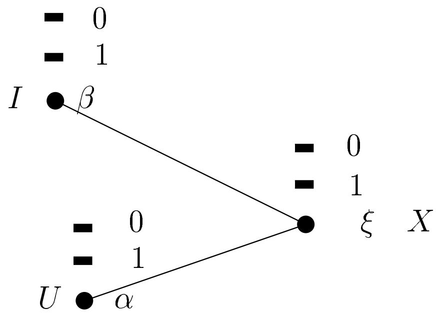

Figure 1.3: Generator and Inhibitor

1. (X_1, U) and (X_2, I) are stochastically independent,
2. $X = X_1 X_2$

holds true. Then $\mathbb{P}(X \in F|X_2 = 1) = \mathbb{P}(X_1 \in F|X_2 = 1) = \mathbb{P}(X_1 \in F)$ would reflect the distribution of X in the situation of no disturbance present.

This means we would like to require that "without disturbance"

$$\mathbb{P}(X_1 = 1|U = 1) = 1, \quad \mathbb{P}(X_1 = 1|U = 0) = 0 \tag{18}$$

is true for the probability $\mathbb{P}$ to be constructed. The complementary probabilities are then obvious. This kind of property of $\mathbb{P}$ reflects the fact that the generator and the indicator are in a direct and strong correlation when there is no disturbance in the model.

As for the inhibitor, we would like to have for some $0 < q < 1$

$$\mathbb{P}(X_2 = 1|I = 1) = q, \quad \mathbb{P}(X_2 = 1|I = 0) = 1 \quad , \tag{19}$$

again the complementary probabilities can be computed immediately. That is, given there is a disturbance, the probability that the (dependent component of the) indicator shows the disturbance is $1-q$. As we assume that both influences are stochastically independent, *item* 1 of our above requirement would then imply

$$\mathbb{P}(X = 1|U, I) = \mathbb{P}(X_1 = 1, X_2 = 1|U, I) = \mathbb{P}(X_1 = 1|U)\,\mathbb{P}(X_2 = 1|V). \tag{20}$$

Then, the conditional probabilities given the states of all predecessors of ξ are given by

$$\begin{aligned} \mathbb{P}(X = 1|U = 1, I = 1) &= 1 \cdot q = q \\ \mathbb{P}(X = 1|U = 1, I = 0) &= 1 \cdot 1 = 1 \\ \mathbb{P}(X = 1|U = 0) &= 0 \ . \end{aligned} \tag{21}$$

The problem is the construction of $\mathbb{P}$. It is common to most applications of Probability Theory that it suffices to specify the *distributions* of the random variables involved.

To this end we define Markovian kernels

$$\mathbf{P}_{\alpha\xi}|\underline{\overline{\mathbf{Z}}}_{\alpha} \Rightarrow \underline{\overline{\mathbf{Z}}}_{\xi}, \quad \mathbf{P}_{\beta\xi}|\underline{\overline{\mathbf{Z}}}_{\beta} \Rightarrow \underline{\overline{\mathbf{Z}}}_{\xi}, \tag{22}$$

by

$$\mathbf{P}_{\alpha\xi}(1,\bullet) = (1,0), \quad \mathbf{P}_{\alpha\xi}(0,\bullet) = (0,1), \tag{23}$$

(which reflects (18)) and

$$\mathbf{P}_{\beta\xi}(1,\bullet) = (q,1-q), \quad \mathbf{P}_{\beta\xi}(0,\bullet) = (1,0) \tag{24}$$

(which reflects (19)). The Markovian kernel $\mathbf{P}_{\xi}|\underline{\overline{\mathbf{Z}}}_{\mathbf{V}} \Rightarrow \underline{\overline{\mathbf{Z}}}_{\xi}$ on the set of predecessors (reflecting (20)) is then defined via

$$\mathbf{P}_{\xi}(z_{\bullet},1) = \mathbf{P}_{\alpha\xi}(z_{\alpha},1) \cdot \mathbf{P}_{\beta\xi}(z_{\beta},1). \tag{25}$$

As the complemantary probabilities are resulting, this is now a *definition* of a Markovian kernel

$$\mathbf{P}_{\xi}|\underline{\overline{\mathbf{Z}}}_{\mathbf{V}} \Rightarrow \underline{\overline{\mathbf{Z}}}_{\xi} \; .$$

Hence, given any $\boldsymbol{\mu}$ on the roots α and β we can construct the distribution $\boldsymbol{m}^{\mathbf{P}}_{\boldsymbol{\mu}}\!\boldsymbol{:}$ according to Theorem 1.8 and set up a random variable Z with this distribution (actually, the concrete definition of $(\Omega, \underline{\underline{\mathbf{F}}}, \mathbb{P})$ is not required as all probabilities of interest can be formulated in terms of the distribution).

It is then not hard to see that $\mathbf{P}_{\xi}$ is indeed a Markovian kernel and that $\boldsymbol{m}^{\mathbf{P}}_{\boldsymbol{\mu}}\!\boldsymbol{:}$ and $\mathbb{P}$ satisfy indeed the desired equations like (21), etc.

Once the probabilities are established, one can compute further data. E.g. the main problem in our simple example is to solve the following question: Given the the indicator shows no fault, what is the probability of a disturbance?

This amounts the computation of $\mathbb{P}(I = 1|X = 1)$. The procedure to base this computation entirely on the given data (the initial distribution and the kernels) is known as "Bayes Theorem". In the present context, assuming that we know the probabilities for the states of the roots (in particular, the one for the inhibitor, i.e., $(\boldsymbol{\mu}^{\beta}_{z_{\beta}})_{z_{\beta}\in\underline{\overline{\mathbf{Z}}}_{\beta}}$) then we obtain

(26)

$$
\begin{aligned}
\mathbb{P}(I=1|X=1) &= \mathbb{P}(U=1, I=1|X=1) \\
&= \frac{\mathbb{P}(U=1, I=1, X=1)}{\mathbb{P}(X=1)} \\
&= \frac{\mathbb{P}(X=1, U=1, I=1)}{\mathbb{P}(U=1, I=1)} \frac{\mathbb{P}(U=1, I=1)}{\mathbb{P}(X=1)} \\
&= \mathbb{P}(X=1|U=1, I=1)\frac{\mathbb{P}(U=1, I=1)}{\mathbb{P}(X=1)} \\
&= \frac{q\boldsymbol{\mu}_1^\alpha \boldsymbol{\mu}_1^\beta}{\sum_{(z_\alpha, z_b)\in \underline{\overline{\mathbf{Z}}}_\alpha \times \underline{\overline{\mathbf{Z}}}_\beta} \mathbb{P}(X=1|(U,I)=(z_\alpha, z_\beta))\mathbb{P}((U,I)=(z_\alpha, z_\beta))} \\
&= \frac{q\boldsymbol{\mu}_1^\alpha \boldsymbol{\mu}_1^\beta}{q\boldsymbol{\mu}_1^\alpha \boldsymbol{\mu}_1^\beta + 1\cdot \boldsymbol{\mu}_1^\alpha \boldsymbol{\mu}_0^\beta} \\
&= \frac{q\boldsymbol{\mu}_1^\beta}{q\boldsymbol{\mu}_1^\beta + 1\cdot \boldsymbol{\mu}_0^\beta} \,.
\end{aligned}
$$

∘ ∼∼∼∼∼∼∼ ∘

The above example motivates the following definition:

Definition 1.13. *Let* $\Sigma := (\mathcal{X}, \prec; \mathcal{Z}; \mathbf{P}_\bullet)$ *be a causal network and let, for all* $\xi \in \mathcal{X}$ *be the states given by* $\underline{\overline{\mathbf{Z}}}_\xi = \{0,1\}$. *A node* $\xi \notin \mathcal{R}$ *is said to constitute **conjunctive interaction** (or just a to be an **and–gate**) if there exists a family of Markovian kernels*

$$\mathbf{P}_{\eta\xi}|\underline{\overline{\mathbf{Z}}}_\eta \Rightarrow \underline{\overline{\mathbf{Z}}}_\xi \quad (\eta \in \mathbf{V}(\xi)) \tag{27}$$

such that

$$\mathbf{P}_\xi(z_\mathbf{V}, 1) = \prod_{\eta\in\mathbf{V}(\xi)} \mathbf{P}_{\eta\xi}(z_\eta, 1) \tag{28}$$

holds true for any joint state $z_\mathbf{V} = (z_\eta)_{\eta\in\mathbf{V}} \in \underline{\overline{\mathbf{Z}}}_\mathbf{V}$ *of* $\mathbf{V} = \mathbf{V}(\xi)$ *and any state* $\alpha \in \underline{\overline{\mathbf{Z}}}_\xi$.

The interpretation is obvious: if Z is a causal process with distribution $\boldsymbol{m}_{\boldsymbol{\mu}}^{\mathbf{P}_\bullet}$, then the conditional probability

$$\mathbb{P}\left(Z_\xi = 1 \,\middle|\, Z_\mathbf{V} = z_\mathbf{V}\right) = \prod_{\eta\in\mathbf{V}(\xi)} \mathbb{P}\left(Z_{\eta\xi} = \alpha \mid Z_\eta = z_\eta\right) \tag{29}$$

reflects the stochastic independence of the influence of some disturbances $Z_{\eta\xi}$ of the predecessors regarding the "gate" ξ jointly constituting the influence of the inhibitor.

In particular, there may be a exceptional state say $\bar{u}$ which reflects the case that all predecessors agree to letting a signal pass. Then the independent consent of all predecessors opens the gate, formally:

Theorem 1.14. *Let $\Sigma := (\mathcal{X}, \prec; \mathcal{Z}; \mathbf{P}_\bullet)$ be a causal network and let ξ be an end–gate. Suppose $\bar{u} = (\bar{u}_\eta)_{\eta \in \mathbf{V}}$ is a state of $\mathbf{V} = \mathbf{V}(\xi)$ satisfying*

$$\mathbf{P}_{\eta\xi}(\bar{u}_\eta, 1) = 1 \quad , \quad \mathbf{P}_{\eta\xi}(z_\eta, 1) = 0 \quad (\eta \in \mathbf{V},\ z_\eta \neq \bar{u}_\eta) \tag{30}$$

for $\underline{\overline{\mathbf{Z}}}_\xi = \{0, 1\}$. Then, for any causal process with corresponding distribution we have

$$\mathbb{P}\left(Z_\xi = 1 \,|\, Z_{\mathbf{V}} = z_{\mathbf{V}}\right) \;=\; 1 \tag{31}$$

if and only if $z_{\mathbf{V}} = \bar{u}$ holds true, i.e., if all predecessors consent.

Proof: We know that

$$\begin{aligned} \mathbb{P}\left(Z_\xi = 1 \,|\, Z_{\mathbf{V}} = z_{\mathbf{V}}\right) \;&=\; \mathbf{P}_\xi(z_{\mathbf{V}}, 1) \\ = \prod_{\eta \in \mathbf{V}(\xi)} \mathbf{P}_{\eta\xi}(z_\eta, 1) \;&=\; \begin{cases} 1 & \text{if all } z_\eta = u_\eta \\ 0 & \text{if one } z_\eta \neq u_\eta \end{cases} \end{aligned}$$

holds true.

q.e.d.

The following definition again is attempted only for the case that all state spaces satisfy $\underline{\overline{\mathbf{Z}}}_\xi = \{0, 1\}$. Then we have

Definition 1.15. *Let $\Sigma := (\mathcal{X}, \prec; \mathcal{Z}; \mathbf{P}_\bullet)$ be a causal network and let, for all $\xi \in \mathcal{X}$ be the states given by $\underline{\overline{\mathbf{Z}}}_\xi = \{0, 1\}$. A node $\xi \notin \mathcal{R}$ is said to constitute* ***disjunctive interaction*** *(or just a to be an* ***or–gate****) if there exists a family of Markovian kernels*

$$\mathbf{P}_{\eta\xi} | \underline{\overline{\mathbf{Z}}}_\eta \Rightarrow \underline{\overline{\mathbf{Z}}}_\xi \quad (\eta \in \mathbf{V}(\xi)) \tag{32}$$

such that

$$\mathbf{P}_\xi(z_{\mathbf{V}}, 1) \;=\; 1 - \prod_{\eta \in \mathbf{V}(\xi)} \mathbf{P}_{\eta\xi}(1 - z_\eta, 1) \tag{33}$$

holds true for any joint state $z_{\mathbf{V}} = (z_\eta)_{\eta \in \mathbf{V}} \in \underline{\overline{\mathbf{Z}}}_{\mathbf{V}}$ of $\mathbf{V} = \mathbf{V}(\xi)$ and state $1 \in \underline{\overline{\mathbf{Z}}}_\xi$.

Lemma 1.16. *Suppose we have, for some node ξ with disjunctive interaction*

(34) $$\mathbf{P}_{\eta\xi}(0,1) = 0, \quad \mathbf{P}_{\eta\xi}(1,1) = 1$$

(with the complementary probabilities resulting). Then, for any process Z with distribution accordingly, we find $\mathbb{P}(Z_\eta = 1|Z_{\mathbf{V}} = u) = 1$ if and only if for at least one η we have $z_\eta = 1$. That is, the gate is open whenever at least one predecessor is open.

Proof: Clearly, as we have (writing $X := Z_\xi,\ U = Z_{\mathbf{V}}$)

(35) $$\begin{aligned} \mathbb{P}(X = 1|U = u) &= 1 - \prod_{\eta\in\mathbf{V}} \mathbf{P}_{\eta\xi}(1 - z_\eta, 1) \\ &= 1 - \prod_{\eta\in\mathbf{V},\mathbf{z}_\eta=\mathbf{0}} \mathbf{P}_{\eta\xi}(1,1) \cdot \prod_{\eta\in\mathbf{V},\mathbf{z}_\eta=\mathbf{1}} \mathbf{P}_{\eta\xi}(0,1) \\ . \\ &= 1 - 1\cdot 0 = 1 \end{aligned}$$

whenever $\{\eta \in \mathbf{V}, z_\eta = 1\} \neq \emptyset$ holds true.

q.e.d.

We may now continue by combining certain types of gates in various manner, thus describe the probabilistic nature of propagation of flows or information. E.g., the following example describes a simple combinatin of various *AND*–gates with an *OR* gate.

Example 1.17 (The noisy OR–gate). Imagine that the inputs to an *OR*–gate are of the nature that the direct influence is disturbed by an inhibitor as described in Example 1.12.

We describe the final gate at node ξ by a random variable X, the various *OR*–gates by nodes $\{\eta|\eta \in \mathbf{V}\}$ with random Variables Z_η $(\eta \in \mathbf{V})$. Each node η has two predecessors, say α_η, β_η and the random variable are U_η, I_η, each playing the role of a generator and an inhibitor (cf. Example 1.12). Of course, all state spaces are $\{0, 1\}$

We are now interested in the joint influence of the joint variables $U = (U_\eta)_{\eta\in\mathbf{V}}$ and $I = (I_\eta)_{\eta\in\mathbf{V}}$ on the states of the final gate.

Thus, we obtain for the condition probabilities that the gate is open:

(36)
$$\mathbb{P}(X = 1|U = u, I = v) = \sum_{z\in\underline{\overline{\mathbf{Z}}}_{\mathbf{V}}} \mathbb{P}(X = 1|Z = z)\mathbb{P}(Z = z|U = u, I = v)\ .$$

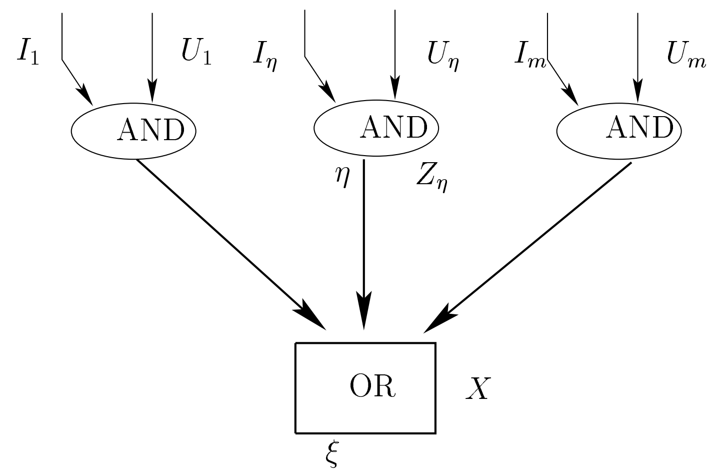

Figure 1.4: The noisy OR–gate

The first expression is in view of the character of the *OR*–gate given by

$$\begin{aligned}
&\mathbb{P}(X = 1|Z = z)\\
= \;& 1 - \prod_{\eta\in\mathbf{V}} \mathbb{P}(X = 1|Z_\eta \neq z_\eta)\\
= \;& 1 - \prod_{\eta\in\mathbf{V},z_\eta=1} \mathbb{P}(X = 1|Z_\eta = 0) \prod_{\eta\in\mathbf{V},z_\eta=0} \mathbb{P}(X = 1|Z_\eta = 1)\\
= \;& 1 \text{ iff at least one } z_\eta = 1\ .
\end{aligned} \tag{37}$$

Of course this ist just a reformulation of equation (35). Hence, returning to equation (36), we continue

$$\begin{aligned}
&\mathbb{P}(X = 1|U = u, I = v)\\
= \;& \sum_{\substack{z \in \overline{\mathbf{Z}}_{\mathbf{V}},\\ \text{at least one } z_h = 1}} \mathbb{P}(Z = z|U = u, I = v)\\
= \;& 1 - \mathbb{P}(Z = 0|X = u, I = v)\\
= \;& 1 - \prod_{\eta\in\mathbf{V}} \mathbb{P}(Z_\eta = 0|X_\eta = u_\eta, I_\eta = v_\eta)\\
= \;& 1 - \prod_{\eta\in\mathbf{V},u_\eta=0} 1 \prod_{\eta\in\mathbf{V},u_\eta=1,v_\eta=0} 0 \prod_{\eta\in\mathbf{V},u_\eta=1,v_\eta=1} q_\eta\\
= \;& \begin{cases} 0 & \text{all } u_\eta = 0\\ 1 & \exists\eta : u_\eta = 1,\ v_\eta = 0\\ 1 - \prod_{\eta\in\mathbf{V},u_\eta=1,v_\eta=1} q_\eta & \text{otherwise}\ . \end{cases}
\end{aligned} \tag{38}$$

The interpretation is obvious: The final gate is certainly open if one of the originators is open and uninhibited.

Otherwise the probability the gate being closed is the (independent) product probability that all inhibitors close all those gates, at which the originator decides to open its gate. The probability for the gate beeing open is the complementary one.

◦ ∼∼∼∼∼∼∼ ◦

2 Processes on a tree

Within this section we assume that the basic graph $(\mathcal{X}, \prec)$ is a tree.

Definition 2.1. *Let* $\Sigma_0 = (\mathcal{X}, \prec, \mathcal{Z})$ *be a state network such that* $(\mathcal{X}, \prec)$ *is a tree and* $\underline{\overline{\mathbf{Z}}}_\xi = \{0,1\}$ *holds true for all* $\xi \in \mathcal{X}$. *Let* $\underline{\overline{\mathbf{Z}}} := \{0,1\}^{\mathcal{X}}$ *be the state space and let* $(\Omega, \underline{\underline{\mathbf{F}}}, \mathbb{P})$ *be a probability space. A random variable* $Z : \Omega \to \underline{\overline{\mathbf{Z}}}$ *is a* ***node seizure*** *if it satisfies*

1. $\mathbb{P}(Z_\eta = 1 | Z_\xi = 1) = 1 \quad (\eta \prec \xi, \xi \in \mathcal{X} - \mathcal{R})$
2. $\sum_{\eta \in \mathbf{N}(\xi)} \mathbb{P}(Z_\xi = 1 | Z_\eta = 1) = 1 \quad (\xi \in \mathcal{X} - \partial\mathcal{X})$

Thus, given that a node ξ is "seized", then the predecessor necessarily has been seized. On the other hand, given that a node is seized, it follows that one of the successors is seized with probability 1.

Lemma 2.2. *Let* $\boldsymbol{x} = (x_1, \ldots, x_T)$ *be a play (i.e., a path connecting the root and the boundary) and let* Z *be a node seizure. Then it follows that*

1. $\mathbb{P}(Z_{x_T} = 1) = \mathbb{P}(Z_{x_0} = 1, \ldots, Z_{x_T} = 1)$
2. $\mathbb{P}(Z_{x_T} = 1 | Z_{x_{T-1}} = 1) = \mathbb{P}(Z_{x_T} = 1 | Z_{x_0} = 1, \ldots, Z_{x_{T-1}} = 1)$

That is the last node is seized with the same probability as the complete path. Moreover, node seizures behave "Markovian".

Proof:

1^st^STEP : Generally, for additive set functions we can verify:

If $\boldsymbol{m}(A \cap B) = \boldsymbol{m}(A)$, then for all F $\boldsymbol{m}(A \cap B \cap F) = \boldsymbol{m}(A \cap F)$.
This follows from

$$\boldsymbol{m}(A \cap F) = \boldsymbol{m}(A \cap B \cap F) + \underbrace{\boldsymbol{m}((A - B) \cap F)}_{0}$$

2^nd^STEP : In the following we assume that all sets involved have positive probability – in order to make sure that all fractions are well defined. However, the proof can easily be adapted to cases in which this assumption is violated. Thus, we observe

$$1 = \mathbb{P}(Z_{x_{t-1}} = 1 | Z_{x_t} = 1) = \frac{\mathbb{P}(Z_{x_{t-1}} = 1, Z_{x_t} = 1)}{\mathbb{P}(Z_{x_{t-1}})}$$

hence

$$\begin{aligned} \mathbb{P}(Z_{x_t}=1) &= \mathbb{P}(Z_{x_{t-1}}=1, Z_{x_t}=1) \\ \mathbb{P}(Z_{x_{t-1}}=1) &= \mathbb{P}(Z_{x_{t-1}}=1, Z_{x_{t-2}}=1). \end{aligned} \tag{1}$$

The second line of (1) (form the intersection with $\{Z_{x_t}=1\}$) implies in view of the first step

$$\mathbb{P}(Z_{x_{t-1}}=1, Z_{x_t}=1) \ = \ \mathbb{P}(Z_{x_{t-1}}=1, Z_{x_{t-2}}=1, Z_{x_t}=1).$$

and in view of the first line of 1 we conclude

$$\mathbb{P}(Z_{x_t}=1) \ = \ \mathbb{P}(Z_{x_t}=1, Z_{x_{t-1}}=1, Z_{x_{t-2}}=1).$$

This way we may continue untill the first statement of the Lemma is verified.

3rdSTEP : In order to verify the second statement, we can now employ the first result thus we have

$$\begin{aligned} 1 \ &= \ \mathbb{P}(Z_{x_T}=1|Z_{x_{T-1}}=1) \\ &= \ \frac{\mathbb{P}(Z_{x_T}=1, Z_{x_{T-1}}=1)}{\mathbb{P}(Z_{x_{T-1}}=1)} \\ &= \ \frac{\mathbb{P}(Z_{x_T}=1,\ldots,Z_{x_0}=1)}{\mathbb{P}(Z_{x_{T-1}}=1,\ldots,Z_{x_0}=1)} \\ &= \ \mathbb{P}(Z_{x_T}=1|Z_{x_{T-1}}=1,\ldots,Z_{x_0}=1) \end{aligned} \tag{2}$$

q.e.d.

Theorem 2.3. *Let Z be a node seizure. Suppose Z satisfies*

$$\mathbb{P}(Z_{\xi_0}) = 1 \tag{3}$$

where ξ_0 is the root of the tree. Then

$$\sum_{\boldsymbol{x}\in\overline{\mathbf{X}}} \mathbb{P}(Z_{x_0}=1,\ldots,Z_{x_T}=1) = 1 \ . \tag{4}$$

That is, if the initial node is surely seized, then a play will be seized with probability 1.

Proof: The proof runs by induction according to the maximal length of a play. It is clear that for length 0 we have

$$\sum_{x_0} \mathbb{P}(Z_{x_0}=1) = \mathbb{P}(Z_{\xi_0}=1) = 1 \ .$$

For length 1 we obtain

$$\begin{aligned}\sum_{\{x_1 | x_0 \prec x_1\}} \mathbb{P}(Z_{x_0} = 1, Z_{x_1} = 1) &= \sum_{\{x_1 | x_0 \prec x_1\}} \frac{\mathbb{P}(Z_{x_0} = 1, Z_{x_1} = 1)}{\mathbb{P}(Z_{x_0} = 1)} \mathbb{P}(Z_{x_0} = 1) \\ &= \underbrace{\sum_{\{x_1 | x_0 \prec x_1\}} \mathbb{P}(Z_{x_0} = 1 | Z_{x_1} = 1)}_{=1,\text{ item 2 of Definition 2.1}} \underbrace{\mathbb{P}(Z_{x_0} = 1)}_{=1}.\end{aligned}$$

This provides the clue for the induction. We have

(5)
$$\begin{aligned}&\sum_{x_0 \prec x_1 \ldots \prec x_T} \mathbb{P}(Z_{x_0} = 1, \ldots, Z_{x_T} = 1) \\ &= \sum_{x_0 \prec x_1 \ldots \prec x_T} \mathbb{P}(Z_{x_T} = 1 | Z_{x_0} = 1, \ldots, Z_{x_{T-1}} = 1) \mathbb{P}(Z_{x_0} = 1, \ldots, Z_{x_{T-1}} = 1) \\ &= \sum_{x_0 \prec x_1 \ldots \prec x_T} \mathbb{P}(Z_{x_T} = 1 | Z_{x_{T-1}} = 1) \mathbb{P}(Z_{x_0} = 1, \ldots, Z_{x_{T-1}} = 1) \\ &= \sum_{x_0 \prec x_1 \ldots \prec x_{T-1}} \mathbb{P}(Z_{x_0} = 1, \ldots, Z_{x_{T-1}} = 1) \sum_{\{x_T | x_{T-1} \prec x_T\}} \mathbb{P}(Z_{x_T} = 1 | Z_{x_{T-1}} = 1).\end{aligned}$$

Here, the second sum equals 1 again in view of *item* 2 of Definition 2.1 and, thereafter, the first sum equals 1 by induction hypothesis.

q.e.d.

Next we would like the node seizure to generate a unique path. But the process is supposed to be stochastic, hence depends on ω. Thus we come up with

Definition 2.4. *Let $Z \to \underline{\overline{\mathbf{Z}}}$ be a node seizure. A sample ω* ***generates a path*** $\boldsymbol{x}$ $= (x_0, \ldots, x_T)$ *(a play) if*

$$\{\xi | Z_\xi(\omega) = 1\} = \{x_0, \ldots, x_T\} \tag{6}$$

holds true.

Theorem 2.5. *Let Z be a node seizure and suppose that, in addition, we have the two conditions*

$$\mathbb{P}(Z_{\xi_0} = 1) = 1 \tag{7}$$

and

$$\begin{gathered}\text{For all } \xi \in \mathfrak{X} - \partial\mathfrak{X} \text{ and } \eta, \eta' \in \mathbf{N}(\xi) \\ \mathbb{P}(Z_\eta = 1, Z_{\eta'} = 1) = 0\ ,\end{gathered} \tag{8}$$

holds true. Then, there exists $\bar{\Omega} \in \Omega$ such that

1. *Every* $\omega \in \Omega$ *generates a path,*
2. $\mathbb{P}(\bar{\Omega}) = 1$

holds true.

That is, if the root is surely seized and no branching takes place, then plays are generated with probability 1.

Proof: For any $\boldsymbol{x} = (x_0, \dots x_T) \in \underline{\overline{\mathbf{X}}}$ define

$$\Omega_{\boldsymbol{x}} \ := \ \{\omega \,|\, \{\xi | Z_\xi(\omega) = 1\} = \{x_0, \dots, x_T\}\} \ .$$

Then, by Theorem 2.3, a path is generated with probability 1 i.e., we have

$$\mathbb{P}\left(\bigcup_{\boldsymbol{x} \in \underline{\overline{\mathbf{X}}}} \Omega_{\boldsymbol{x}}\right) \ = \ 1 \ . \tag{9}$$

It follows from equation (8) at once, that

$$\mathbb{P}\left(\Omega_{\boldsymbol{x}} \cap \Omega_{\boldsymbol{x}'}\right) = 0 \tag{10}$$

for any $\boldsymbol{x} \neq \boldsymbol{x}'$. Therefore, if we put

$$\bar{\Omega}_{\boldsymbol{x}} \ := \ \Omega_{\boldsymbol{x}} - \bigcup_{\{\boldsymbol{x}' | \boldsymbol{x}' \neq \boldsymbol{x}\}} \Omega_{\boldsymbol{x}'}, \tag{11}$$

then we obtain that

$$\mathbb{P}(\bar{\Omega}_{\boldsymbol{x}}) \ = \ \mathbb{P}(\Omega_{\boldsymbol{x}}) - \mathbb{P}(\Omega_{\boldsymbol{x}} \cap \bigcup_{\{\boldsymbol{x}' | \boldsymbol{x}' \neq \boldsymbol{x}\}} \Omega_{\boldsymbol{x}'}) = \mathbb{P}(\Omega_{\boldsymbol{x}})$$

holds true. Therefore, putting

$$\bar{\Omega} \ := \ \bigcup_{\boldsymbol{x} \in \underline{\overline{\mathbf{X}}}} \bar{\Omega}_{\boldsymbol{x}}$$

we obtain

$$\mathbb{P}(\bar{\Omega}) = 1, \quad \mathbb{P}(\bar{\Omega}_{\boldsymbol{x}}) \cap \mathbb{P}(\bar{\Omega}_{\boldsymbol{x}'}) \ = \ \emptyset \ \ (\boldsymbol{x} \neq \boldsymbol{x}'). \tag{12}$$

Obviously, for any $\omega \in \Omega$ there exists a *unique* $\boldsymbol{X} \in \underline{\overline{\mathbf{X}}}$ satisfying

$$\{\xi | Z_\xi(\omega) = 1\} = \{x_0, \dots, x_T\}$$

q.e.d.

Remark 2.6. Given the situation as in the previous Theorem, the last sentence can be interpreted as to say: there is a well defined mapping

$$X : \bar{\Omega} \to \underline{\overline{\mathbf{X}}}$$

such that

$$\{\xi | Z_\xi(\omega) = 1\} = \{X_0(\omega), \ldots, X_T(\omega)\}$$

holds true. This motivates the following definition.

∘ ∼∼∼∼∼∼∼ ∘

Definition 2.7. *A* ***flow*** *is a node seizure Z that satisfies conditions* (7) *and* (8)*. A* ***process*** *is a mapping $X : \bar{\Omega} \to \underline{\overline{\mathbf{X}}}$. A node seizure Z* ***induces*** *a process X whenever*

$$\{\xi | Z_\xi(\omega) = 1\} = \{X_0(\omega), \ldots, X_T(\omega)\}$$

is the case.

Theorem 2.8. *Any flow induces uniquely a process. For any proces there exists a unique flow inducing it.*

Proof: The first statement is actually a consequence of the proof for the previous theorem as been made clear in Remark 2.6. On the other hand, if then define Z by

$$Z_\xi(\omega) = 1 \Leftrightarrow \xi \in \{X_0(\omega), \ldots, X_T(\omega)\} \quad (\omega \in \Omega), \tag{13}$$

then it is not hard to check that Z is a flow.

q.e.d.

Theorem 2.9. *Let $(\mathcal{X}, \prec)$ be a tree. Suppose that, for every $\xi \in \mathcal{X} - \partial\mathcal{X}$, there is a probability $\boldsymbol{\mu}^\xi_\bullet$ defined on $\mathbf{N}(\xi)$. Then there is a probability space $(\Omega, \underline{\underline{\mathbf{F}}}, \mathbb{P})$, and a node seizure Z with the following properties:*

1. $\mathbb{P}(Z_\xi = 1 | Z_\eta = 1) = \boldsymbol{\mu}^\xi_\eta = \boldsymbol{\mu}^\xi(\{\eta\}) \quad \eta \prec \xi$
2. $\mathbb{P}$ *satisfies* (7) *and* (8)*, hence induces a process X.*
3. *The process X satisfies*
   (a) $\mathbb{P}(\{X_0 = \xi_0\}) = 1$ *where ξ_0 is the root of $(\mathcal{X}, \prec)$,*

(b) for any play $\boldsymbol{x}$ *it follows that*

$$\text{(14)} \qquad \mathbb{P}(\{X_0 = x_0, \ldots, X_T = x_T) = 1 \cdot \boldsymbol{\mu}_{x_1}^{x_0} \cdot \boldsymbol{\mu}_{x_2}^{x_1} \cdot \ldots \cdot \boldsymbol{\mu}_{x_T}^{x_{T-1}}.$$

Proof: Take $\Omega := \underline{\overline{\mathbf{X}}}$ and *define* $\boldsymbol{m}$ by

$$\text{(15)} \qquad \boldsymbol{m}(\{x_0, \ldots, x_T) = 1 \cdot \boldsymbol{\mu}_{x_1}^{x_0} \cdot \boldsymbol{\mu}_{x_2}^{x_1} \cdot \ldots \cdot \boldsymbol{\mu}_{x_T}^{x_{T-1}}.$$

Define X to be the identity and Z by Theorem 2.8. All claims follow by standard computations.

q.e.d.

Definition 2.10. *1. Let* $\Sigma = (\mathfrak{X}, \prec; \boldsymbol{\iota}; f, u)$ *be a game tree. A* ***behavioral strategy*** *for player* i *is a family of probabilities*

$$\text{(16)} \qquad \boldsymbol{A}^i(\xi, \bullet) \quad (\xi \in \mathfrak{X}^i = \boldsymbol{\iota}^{-1}(\{i\})$$

each one defined on $\mathbf{N}(\xi)$ $(\xi \in \mathfrak{X}^i)$.

We write $\boldsymbol{A}$ *for an* n*–tuple of behavioral strategies.*

2. Given an n*–tuple* $\boldsymbol{A}$ *of behavioral strategies, let* $\boldsymbol{m}^{\boldsymbol{A}}$ *denote the probability on* $\underline{\overline{\mathbf{X}}}$ *defined by* (15)*, i.e., by*

$$\text{(17)} \qquad \begin{aligned} &\boldsymbol{m}^{\boldsymbol{A}}(\{x_0, \ldots, x_T) \\ &= 1 \cdot \boldsymbol{A}^{\boldsymbol{\iota}(x_0)}(x_0, x_1) \cdot \boldsymbol{A}^{\boldsymbol{\iota}(x_1)}(x_1, x_2) \cdot \ldots \cdot \boldsymbol{A}^{\boldsymbol{\iota}(x_{T-1})}(x_{T-1}, x_T) \\ &\qquad (\boldsymbol{x} \in \underline{\overline{\mathbf{X}}}) \ . \end{aligned}$$

clearly, a process $X : \Omega \to \underline{\overline{\mathbf{X}}}$ has distribution $\boldsymbol{m}^{\boldsymbol{A}}$ if it satisfies *item* 3. of Theorem 2.9 with the obvious modification ($\boldsymbol{\mu}$ given via $\boldsymbol{A}$). Thus we may continue

Definition 2.11. *Let* Σ *be a game tree and* $\boldsymbol{A}$ *an* n*–tuple of behavioral strategies. Also, let*

$$\text{(18)} \qquad \begin{aligned} \boldsymbol{C}^i &= \boldsymbol{C}^{i\xi_0} : \underline{\overline{\mathbf{X}}} \to \mathbb{R} \\ \boldsymbol{C}^i(\boldsymbol{x}) &:= \sum_{t=0}^{T} f^i(x_{t-1}, x_t) + u^i(x_T) \\ &=: F^i(\boldsymbol{x}) + u^i(x_T) \quad (\boldsymbol{x} \in \underline{\overline{\mathbf{X}}}) \end{aligned}$$

denote the evaluation (see Definition 2.7). The ***payoff*** *to player* $i \in \boldsymbol{I}$ *at* $\boldsymbol{A}$ *is*

$$
\begin{aligned}
\boldsymbol{C}^i_{\boldsymbol{A}} &:= \mathbb{E}\sum_{t=0}^{T} f^i(X_{t-1}, X_t) + u^i(X_T) \\
&= \int_{\Omega} C^i \circ X d\mathbb{P} \\
&= \int_{\underline{\mathbf{X}}} C^i d\boldsymbol{m}^{\boldsymbol{A}}
\end{aligned}
\tag{19}
$$

where X *is a process with distribution* $\boldsymbol{m}^{\boldsymbol{A}}$.

Let $\mathfrak{B}^i$ *denote the set of behavioral strategies of player* $i \in \boldsymbol{I}$. *The noncooperative* n*–person game*

$$
\Gamma^{\star}_{\Sigma} := (\mathfrak{B}^1, \ldots, \mathfrak{B}^n; \boldsymbol{C}^1, \ldots \boldsymbol{C}^n)
\tag{20}
$$

is the ***normal form in behavioral strategies*** *generated by* Σ.

We can now start talking about backwards induction and subgame perfectness in behavioral strategies. However, we consider the last exposition only as a supplement to SECTION *1. Indeed, the above defined extended normal form does not yield anything new for tree games, as Nash equilibria in such games exist already in pure strategies. The situation changes when we take imperfect or incomplete information into account. This is the subject of the next section.*

3 Tree Games with Imperfect Information

We consider a model with *imperfect information* of the players concerning the state of the game. Thus, may not be informed about the current state of the process or the current node the play has been moved into. Nevertheless, the player in charge has to make a decision concerning the next node.

A consistent and precise model containing the exact notion of imperfect or incomplete information turns out to be somewhat tedious device. To say that a player cannot distinguish between several nodes of his command is to say that he receives a signal depending on the nodes a process may acquire. The signal may be the same for everal nodes. Equivalently, we may describe a set of nodes resulting in the same signal, called an "information set".

In what follows we include a player named "chance " into the model. Hence, the player set is always meant to be $\boldsymbol{I}_0 := \boldsymbol{I} \cup \{0\}$. Our intuition is that player "chance" does not act strategically but ads to randomness: at some node it chooses a successor by a random device represented by probability attached to that node.

The further details are specified successively by a sequence of definitions. The first task is to descibe the extensive form.

Definition 3.1. *A **game tree with i.i.** is a tuple*

$$\Sigma = (\mathcal{X}, \prec; \boldsymbol{\iota}, \boldsymbol{\mu}^{\bullet}; \boldsymbol{\kappa}, \boldsymbol{\sigma}; f, u) \tag{1}$$

the ingredients of which are defined as follows:

$(\mathcal{X}, \prec; \boldsymbol{\iota}; f, u)$ is a game tree in the sense of Definition 2.1. In particular $\boldsymbol{\iota} : \mathcal{X} - \partial\mathcal{X} \to \boldsymbol{I}_0$ assigns every node $\xi \in \mathcal{X} - \partial\mathcal{X}$ to either some player $i \in \boldsymbol{I}$ or chance ($i = 0$). As previously, we write

$$\mathcal{X}^i := \boldsymbol{\iota}^{-1}(\{i\}) = \{\xi \mid \boldsymbol{\iota}(\xi) = i\} \tag{2}$$

for the set of nodes that are controlled by player $i \in \boldsymbol{I}$. In particular, $\mathcal{X}^0$ is the set of ***chance nodes***. Recall that $(\mathcal{X}^i)_{i \in I_0}$ constitutes a partition which is referred to as the ***player partition***.

$\boldsymbol{\mu}^{\bullet} = (\boldsymbol{\mu}^{\xi})_{\xi \in \mathcal{X}^0}$ is a family of probabilities. To each chance node $\xi \in \mathcal{X}^0$ there is attached a probability $\boldsymbol{\mu}^{\xi}$ defined on the set $\mathbf{N}(\xi) = \{\eta \mid \xi \prec \eta\}$ of successors of ξ.

Intuitively, if the process of playing the game reaches ξ, then a random device represented by $\boldsymbol{\mu}^{\xi}$ chooses the successor of ξ, the probability for $\eta \in \mathbf{N}(\xi)$ is denoted by $\boldsymbol{\mu}^{\xi}_{\eta}$.

Next, $\boldsymbol{\kappa} = (\boldsymbol{\kappa}^i)_{i \in \boldsymbol{I}}$ is a family of mappings

$$(3) \qquad \boldsymbol{\kappa}^i \ : \ \mathcal{X}^i \to \mathcal{Q}^i.$$

Here $\mathcal{Q}^i$ is a finite set, called the set of player i's ***signals***. Clearly, he system

$$(4) \qquad \underline{\underline{\mathbf{Q}}}^i \ := \ \left\{ (\boldsymbol{\kappa}^i)^{-1}(q) \,\middle|\, q \in \mathcal{Q}^i \right\}$$

constitute a partition of $\mathcal{X}^i$, this system is called the **information partition** of player i. An element $Q \in \underline{\underline{\mathbf{Q}}}^i$ is an **information set** of player i.

Intuitively, whenever some $\xi \in \mathcal{X}^i$ occurs, player i is only informed about the signal $\boldsymbol{\kappa}^i(\xi)$ and not about the actual state. Thus, player i cannot distinguish between $\xi, \eta \in Q = (\boldsymbol{\kappa}^i)^{-1}(q)$ as he receives the same signal q.

The mapping $\boldsymbol{\kappa}^0$ is assumed to be the identity, – "chance has full information".

Now, that we observe the information available to player i, we have to specify the procedure according to which this player chooses a successor, even though he does not know at which state exactly the process is located. Clearly, if all states compatible with his information would have the same number of successors, then a player could just draw a number in order to specify the next node. This idea is captured by the following exposition concerning the quantitiy $\boldsymbol{\sigma} \ = \ (\boldsymbol{\sigma}_{\xi})_{\xi \in \mathcal{X} - \partial \mathcal{X}}$.

We require that, for every $q \in \mathcal{Q}^i$ and every $\xi \in (\boldsymbol{\kappa}^i)^{-1}(q)$ there is a finite set S^q (the ***choices*** of i when he observes q) and a *bijective* mapping

$$(5) \qquad \boldsymbol{\sigma}_{\xi} \ : \ \mathbf{N}(\xi) \to S^q.$$

Intuitively, if player i observes q, then he may opt for a choice $s \in S^q$. Assuming that the process or play is actually located in $\xi \in Q = (\boldsymbol{\kappa}^i)^{-1}(q)$ (which player i does not observe), the image $\boldsymbol{\sigma}_{\xi}(s)$specifies exactly one successor $\eta \in \mathbf{N}(\xi)$, which is thought of to be the next state resulting from player i's action.

Each $\boldsymbol{\sigma}_\xi$ maps the successors of ξ bijectively onto the set of choices S^q, which is the same for all nodes which yield the same signal. Therefore, all $\xi \in Q = (\boldsymbol{\kappa}^i)^{-1}(q)$ must necessarily have the number of successors, this number is

$$\mid \mathbf{N}(\xi) \mid = \mid S^q \mid \quad (\xi \in (\boldsymbol{\kappa}^i)^{-1}(q)).$$

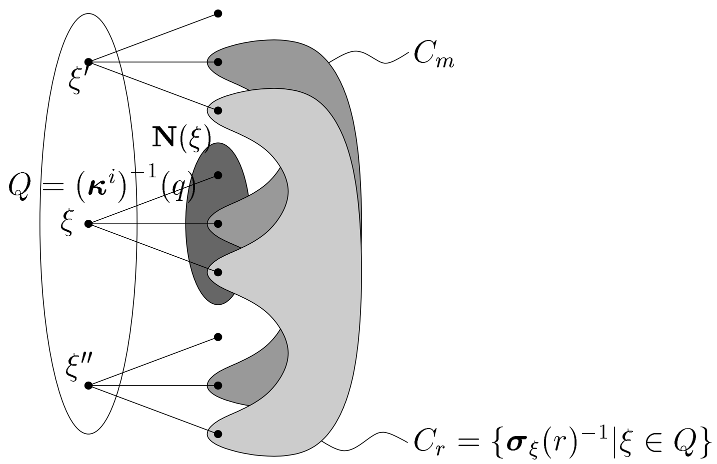

Figure 3.1: An information set and choices

Figure 3.1 depicts some essentials. It is common to represent choices by attaching the code or letter s to those successor which are identified by their common value s of the choice function $\boldsymbol{\sigma}$. I.e., we write r to all elements of $C_r = \{\boldsymbol{\sigma}_\xi^{-1}(r) | \xi \in Q\}$ indicating that these successors will result from their predecessor if the player in charge chooses "right"...

◦ ∼∼∼∼∼∼ ◦

We will now describe various types of **strategic behavior** , i.e., classes of ***strategies*** for a player. Accordingly, there will be the corresponding "normal forms" depending on which kind of strategies we focus upon.

Player i is not capable of observing the states $\xi \in \mathfrak{X}$. Indeed, his observation is restricted to the elements of $\mathfrak{Q}^i$. We postulate that by therefore, player i is forced to restrict his behavior on strategies depending on $\mathfrak{Q}^i$ at most.

Definition 3.2. *Let Σ be a tree game. A* ***pure strategy*** *for player i is a mapping*

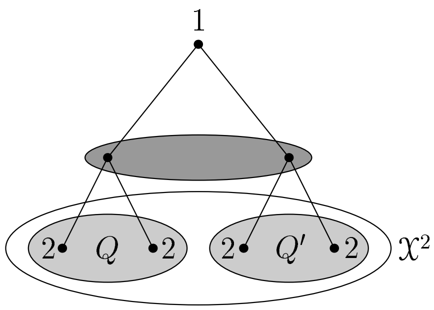

Figure 3.2:

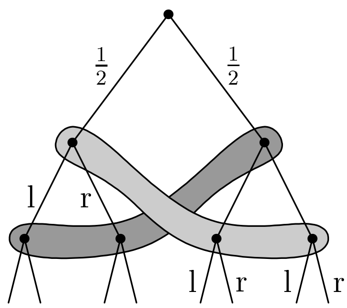

Figure 3.3:

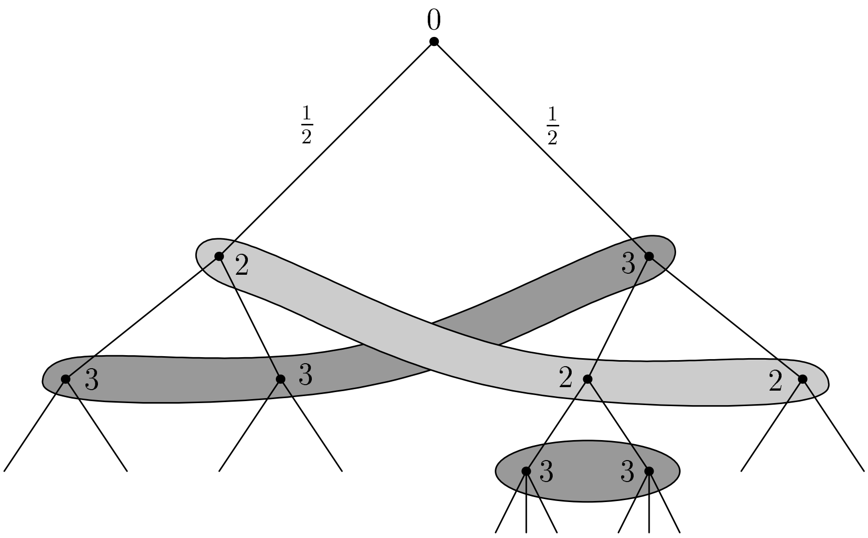

Figure 3.4:

(6) $$\boldsymbol{\alpha}^i : \mathfrak{Q}^i \to \bigcup_{q \in \mathfrak{Q}^i} S^q$$

such that

(7) $$\boldsymbol{\alpha}^i(q) \in S^q \quad (q \in \mathfrak{Q}^i)$$

holds true. We denote by

(8) $$\mathfrak{S}^i := \{\boldsymbol{\alpha}^i \mid \boldsymbol{\alpha}^i \text{ satisfies } (6) \text{ and } (7)\}$$

the set of pure strategies for player i.

The mapping $\boldsymbol{\alpha}^i$ may also be written $\boldsymbol{\alpha}^i = (\boldsymbol{\alpha}^i(q))_{q \in \mathfrak{Q}^i}$; we prefer this version within the next definition.

Definition 3.3. *Let Σ be a tree game. A* ***behavioral strategy (b.s.)*** *for player i is a family*

$$\boldsymbol{A}^i = (\boldsymbol{A}^i(q, \bullet))_{q \in \mathfrak{Q}^i}$$

such that, for every $q \in \mathfrak{Q}^i$

(9) $$\boldsymbol{A}^i(q, \bullet) \text{ is a probability on } S^q.$$

That is, at each information q player i selects a random device over the choices available to him.

(10) $$\mathcal{A}^i := \{\boldsymbol{A}^i \mid \boldsymbol{A}^i \text{ satisfies } (9)\}$$

is the set of behavioral strategies of player i.

There is finally the concept of a mixed strategy which as usual incorporates the idea of randomizing over the pure strategies.

Definition 3.4. *Let Σ be a tree game and let $\mathfrak{S}^i$ denote the pure strategies of player i. A* ***mixed strategy*** *for player i is a probability $\boldsymbol{M}^i$ on $\mathfrak{S}^i$.*

(11) $$\mathfrak{M}^i = \{\boldsymbol{M}^i \mid \boldsymbol{M}^i \text{ is a mixed strategy for player } i\}$$

denotes the corresponding set.

Our standard notation assumes that the Cartesian product is taken if we omit the upper index i, thus

$$\begin{array}{rcl} \mathfrak{S} & = & \mathfrak{S}^1 \times \cdots \times \mathfrak{S}^n \\ \mathcal{A} & = & \mathcal{A}^1 \times \cdots \times \mathcal{A}^n \\ \mathfrak{M} & = & \mathfrak{M}^1 \times \cdots \times \mathfrak{M}^n \end{array}$$

denotes sets of strategy n-tuples.

We call any "coherent" sequence $x = (x_1, \ldots, x_T)$ of nodes (i.e., $x_1 \prec x_2 \prec \ldots \prec x_T$) a ***path***. A path $x = (x_0, \ldots, x_T)$ starting with the root x_0 and ending at the boundary (i.e., $\xi_T \in \partial\mathfrak{X}$) is sometimes called a it **play.** $\underline{\overline{\mathbf{X}}}$ denotes the set of plays.

Remark 3.5.

1. Obviously, if a player i has just one information set, then mixed strategies and behavioral strategies coincide, that is we have $\mathfrak{M}^i = \mathcal{A}^i$.

2. As usual we may regard pure strategies as a particular type of behavioral ones as well as of mixed ones; in other words, pure strategies can be embedded into behavioral strategies (as well as into mixed strategies).

   For, if $\boldsymbol{\alpha}^i$ is a pure strategy then the corresponding behavioral strategy $\boldsymbol{A}^{i,\boldsymbol{\alpha}^i}$ is defined by
   $$\boldsymbol{A}^i(q, \bullet) = \delta_{\boldsymbol{\alpha}^i(q)}(\bullet).$$
   This is a behavioral strategy which puts probability 1 on the choice provided by $\boldsymbol{\alpha}^i(q)$. In this sense we can now state that $\mathfrak{S}^i \subseteq \mathcal{A}^i$ or $\mathfrak{S} \subseteq \mathcal{A}$ holds true.

   The embedding of pure strategies into mixed strategies is also the canonical one; we do skip the discussion of the details.

3. Therefore, if we discuss the distribution induced by behavioral strategies on the paths of a tree, then the distribution induced by pure strategies is as well defined. Note that the latter distribution is not necessarily $\delta-$shaped, as chance moves may be incorporated in the structure of the game tree.

4. As it turns out mixed strategies eventually can be disposed of; it is sufficient to consider behavioral strategies. The background is provided by KUHN's Theorem (cf. Theorem **??**), which however, in the present context requires the additional assumption of perfect recall. To develop the details will be the aim of this section.

5. Finally we would like to note that the choices made by chance (by player 0) can also be viewed as a behavioral strategy (of player 0). More precisely, if we consider the family of distributions provided by $\mu^{\bullet}$ and keep in mind that the information partition of $\mathcal{X}^0$ is the one into singletons, i.e.

$$\mathcal{X}^0 = \sum_{\xi \in \mathcal{X}^0} \{\xi\},$$

then for $q = \xi \in \mathcal{X}^0$ the notation

$$\boldsymbol{A}^0(q, \bullet) = \mu^{\xi}_{\bullet}$$

obviously defines a behavioral strategy for the chance player.

∘ ∼∼∼∼∼∼∼ ∘

Definition 3.6. *Let Σ be a tree game and let*

$$\boldsymbol{A} = (\boldsymbol{A}^i)_{i \in I}$$

be an n–tuple of behavioral strategies . Then a probability $\boldsymbol{m}^{\boldsymbol{A}}_{\boldsymbol{\mu}}$ in $\overline{\mathbf{X}}$ (involving the chance distributions $\boldsymbol{\mu}^{\bullet}$) is provided as follows.

1. *Define transition probabilities via*

$$\boldsymbol{q}^{\xi}_{\eta} := \boldsymbol{A}^{\boldsymbol{\iota}(\xi)}(\boldsymbol{\kappa}^{\boldsymbol{\iota}(\xi)}(\xi), \boldsymbol{\sigma}_{\xi}(\eta)) \quad (\xi \in \mathring{\mathcal{X}}, \eta \in \mathbf{N}(\xi)) \tag{12}$$

   in particular, we have $\boldsymbol{q}^{\xi}_{\eta} = \boldsymbol{\mu}^{\xi}_{\eta} = \boldsymbol{A}^0(\{\xi\}, \{\eta\})$ if $\xi \in \mathcal{X}^0$ is a chance move.

2. *Define the distribution of a node seizure Z via*

$$\mathbb{P}(Z_{\eta} = 1 | Z_{\xi} = 1) = \boldsymbol{q}^{\xi}_{\eta} \quad (\xi \prec \eta) \tag{13}$$

3. *Equivalently, for any play $x = (x_0, \ldots, x_T)$ let*

$$\boldsymbol{m}^{\boldsymbol{A}}_{\boldsymbol{\mu}}(\{x\}) := \boldsymbol{q}^{x_0}_{x_1} \quad \boldsymbol{q}^{x_1}_{x_2} \ldots \boldsymbol{q}^{x_{T-1}}_{x_T}, \tag{14}$$

   *and define a process $X : \Omega \to \overline{\mathbf{X}}$ to be **controlled** by $\boldsymbol{A}$ if X has distribution $\boldsymbol{m}^{\boldsymbol{A}}_{\boldsymbol{\mu}}$.*

4. *The **payoff** to player i given **A** is*

$$C_{\boldsymbol{A}}^{i\boldsymbol{\mu}} := \int_{\underline{\overline{\mathbf{X}}}} \left(\sum_{t=1}^{\tau_0} f^i(x_{t-1} x_t) + u^i(x_{\tau_0}) \right) d\boldsymbol{m}_{\boldsymbol{\mu}}^{\boldsymbol{A}}(x). \tag{15}$$

here $\tau_0 : \underline{\overline{\mathbf{X}}} \to \mathbb{N}$ is the length of the path, i.e., defined by $\tau_0(x) = T$ if $x = (x_0, \ldots, x_T)$.

The payoff to player i resulting from A can also computed in terms of a process X with distribution $\boldsymbol{m}_{\boldsymbol{\mu}}^{\boldsymbol{A}}$. We can write

$$\begin{aligned} C_{\boldsymbol{A}}^{i\boldsymbol{\mu}} &:= \mathbb{E} \left(\sum_{t=1}^{\boldsymbol{\tau}} f^i(X_{t-1}, X_t) + u^i(X_{\tau}) \right) \\ &= \int \left(\sum_{t=1}^{\tau} f^i(x_{t-1}, x_t) + u^i(x_{\tau}) \right) d\boldsymbol{m}_{\boldsymbol{\mu}}^{\boldsymbol{A}}(x) \end{aligned} \tag{16}$$

Here we have used $\boldsymbol{\tau} = \tau_0 \circ X$ to indicate the *random* length of a play. Not unusual, the argument is provided by introducing the ***evaluation***

$$\begin{gathered} C^i : \underline{\overline{\mathbf{X}}} \to \mathbb{R} \\ C^i(x) = \sum_{t=1}^{\tau(x)} f^i(x_{t-1}, x_t) + u^i(x_{\tau_0(x)}) \end{gathered} \tag{17}$$

Since $\tau = \tau_0 \circ X$ it is seen that (16) defines $C_{\boldsymbol{A}}^{i\boldsymbol{\mu}}$ via $C_{\boldsymbol{A}}^{i\boldsymbol{\mu}} = E\ C^i \circ X$ while (15) claims that $C_{\boldsymbol{A}}^{i\boldsymbol{\mu}} = \int_{\underline{\overline{\mathbf{X}}}} C^i(x) d\boldsymbol{m}_{\boldsymbol{\mu}}^{\boldsymbol{A}}(x)$ holds true. The claim is correct by the formula for "transformation of variables" since $\boldsymbol{m}_{\boldsymbol{\mu}}^{\boldsymbol{A}}$ is the distribution of X and hence

$$\mathbb{E}\ C^i \circ X = \int_{\Omega} (C^i \circ X) d\mathbb{P}$$

$$= \int_{\underline{\overline{\mathbf{X}}}} C^i d(X\mathbb{P}) = \int C^i d\boldsymbol{m}_{\boldsymbol{\mu}}^{\boldsymbol{A}}$$

Remark 3.7.

1. As we have already remarked a pure strategy is a special behavioral strategy. Therefore the induced distribution called m^{α}_{μ} is a well defined quantity.

2. It should be clear that for a game tree without chance moves the distribution induced by a pure strategy α is a point measure concentrated on some path $x^{\alpha} \in \underline{\underline{\mathbf{X}}}$.

∘ ~~~~~~~ ∘

We are now in the position to define the normal form corresponding to behavioral strategies.

Definition 3.8. *Let Σ be a tree game and let $\mathcal{A}^1, \ldots, \mathcal{A}^n$ denote the sets of behavioral strategies of the players. Define*

$$C^{i\boldsymbol{\mu}}_{\bullet} : \mathcal{A}^1 \times \cdots \times \mathcal{A}^n \to \mathbb{R}$$

by either (16) or (15). The **normal form** *game induced by Σ (in behavioral strategies) is the (noncooperative) n-person game*

$$\Gamma = \Gamma_{\Sigma\boldsymbol{\mu}} := (\mathcal{A}^1, \ldots, \mathcal{A}^n;\ C^{1\boldsymbol{\mu}}_{\bullet}, \ldots, C^{n\boldsymbol{\mu}}_{\bullet})$$

Pure strategies can be seen as a particular case of behavioral strategies via the embedding

$$\alpha \to \boldsymbol{A}^{\alpha}$$

that is described in (3.5). Hence, if $\mathfrak{S}^i$ is the space of pure strategies of player i and if $C^{i\boldsymbol{\mu}}_{\alpha}$ is given by $C^{i\boldsymbol{\mu}}_{\boldsymbol{A}^{\alpha}}$, then we have an immediate definition of the normal form in pure strategies induced by Σ.

Definition 3.9. *Let Σ be a tree game and let $\mathfrak{S}^1, \ldots, \mathfrak{S}^n$ denote the sets of pure strategies of the players. Define*

$$C^{i\boldsymbol{\mu}}_{\bullet} : \mathfrak{S}^1 \times \cdots \times \mathfrak{S}^n \to \mathbb{R}$$

by $C^{i\boldsymbol{\mu}}_{\alpha} := C^{i\boldsymbol{\mu}}_{\boldsymbol{A}^{\alpha}}$ $(\alpha \in \mathfrak{S})$. Then the **normal form** *game induced by Σ (in pure strategies) is the (noncooperative) n-person game*

$$\tilde{\Gamma} = \tilde{\Gamma}_{\Sigma\boldsymbol{\mu}} = (\mathfrak{S}^1, \ldots, \mathfrak{S}^n; C^{1\boldsymbol{\mu}}_{\bullet}, \ldots, C^{n\boldsymbol{\mu}}_{\bullet}). \tag{18}$$

The first approach for the game in mixed strategies is traditionally based on the concept of the mixed extension: a mixed strategy $\boldsymbol{M}^i$ for player i is a probability on $\mathfrak{S}^i \ (i \in I)$. If

$$\boldsymbol{M} = \boldsymbol{M}^1 \otimes \cdots \otimes \boldsymbol{M}^n \tag{19}$$

denotes the stochastically independent mixture of all players' mixed strategies, then the naive concept for the payoff results in the corresponding expectation. That is, we have

$$C^{i\boldsymbol{\mu}}_{\boldsymbol{M}} := \int\limits_{\mathfrak{S}} C^{i\boldsymbol{\mu}}_{\alpha} d\boldsymbol{M}\ (\alpha). \tag{20}$$

On the other hand, there should be the concept of a process resulting from the application of $\boldsymbol{M}$ or controlled by $\boldsymbol{M}$.

In order to formalize this idea, the first task is to define the distribution induced by a mixed strategy. To this end we naturally attempt to define the mixture of the distributions that result from pure strategies.

We have seen in Remark (3.7) that a pure strategy α results in a distribution $\boldsymbol{m}^{\alpha}_{\boldsymbol{\mu}}$. Mixing these distributions is formally described as follows.

Definition 3.10. *Let Σ be a tree game (with chance distributions $\boldsymbol{\mu}^{\bullet}$). Define a transition kernel*

$$\boldsymbol{K}^{\boldsymbol{\mu}} \mid \mathfrak{S} \Rightarrow \underline{\overline{\mathbf{X}}}$$

by

$$\boldsymbol{K}^{\boldsymbol{\mu}}(\alpha, \bullet) = \boldsymbol{m}^{\alpha}_{\boldsymbol{\mu}} \quad (\alpha \in \mathfrak{S}). \tag{21}$$

Also, let $\boldsymbol{M} \in \mathfrak{M}$ be a mixed strategy and $X : \Omega \to \underline{\overline{\mathbf{X}}}$ a process. Then X is **controlled** *by $\boldsymbol{M}$ (or results from the application of $\boldsymbol{M}$) if the distribution of X is given by*

$$X\mathbb{P} = \boldsymbol{K}^{\boldsymbol{\mu}}\boldsymbol{M} = \int\limits_{\mathfrak{S}} \boldsymbol{m}^{\alpha}_{\boldsymbol{\mu}}(\bullet)\boldsymbol{M}(d\alpha). \tag{22}$$

If we adopt a process distributed according to $\boldsymbol{K}^{\boldsymbol{\mu}}\boldsymbol{M}$ as to "be controlled by $\boldsymbol{M}$", then corresponding payoff should be given by

$$
\begin{aligned}
C^{i\boldsymbol{\mu}}_{\boldsymbol{M}} =& \mathbb{E}\left(\sum_{t=1}^{\tau} f^i(X_{t-1}, X_t) + u^i(X_\tau)\right) \\
=& \int_{\underline{\overline{\mathbf{X}}}} C^i(x)(\boldsymbol{K}^{\boldsymbol{\mu}}M)(dx) \\
=& \int_{\underline{\overline{\mathbf{X}}}} C^i d(\boldsymbol{K}^{\boldsymbol{\mu}}\boldsymbol{M}),
\end{aligned} \tag{23}
$$

where C^i is the "evaluation" defined in (17); compare the corresponding formula for the behavioral strategies which is (16). We hasten to show that (23) does not contradict (20), for we prove

Theorem 3.11. *If Σ is a tree game and $\boldsymbol{M}$ a mixed strategy, then*

$$
\int_{\mathfrak{S}} C^{i\boldsymbol{\mu}}_{\alpha} \boldsymbol{M}(d\alpha) = \int_{\underline{\overline{\mathbf{X}}}} C^i(x)(\boldsymbol{K}^{\boldsymbol{\mu}}\boldsymbol{M})(dx)
$$

Proof: The first step is to prove that

$$
\begin{aligned}
C^{i\boldsymbol{\mu}}_{\alpha} =& \int_{\underline{\overline{\mathbf{X}}}} C^i(x)\boldsymbol{m}^{\alpha}_{\boldsymbol{\mu}}(dx) \\
=& \int_{\underline{\overline{\mathbf{X}}}} C^i(x)\boldsymbol{K}^{\boldsymbol{\mu}}(\alpha, dx) \\
=& (\boldsymbol{K}^{\boldsymbol{\mu}}C^i)\ (\alpha)\ ;
\end{aligned}
$$

holds true. The claim follows immediately in a second step which shows that

$$
\begin{aligned}
C^{i\boldsymbol{\mu}}_{\boldsymbol{M}} =& \int_{\mathfrak{S}} C^{i\boldsymbol{\mu}}_{\alpha}\boldsymbol{M}(d\alpha) \\
=& \int_{\mathfrak{S}} (\boldsymbol{K}^{\boldsymbol{\mu}}C^i)(\alpha)\boldsymbol{M}(d\alpha) \\
=& \int_{\underline{\overline{\mathbf{X}}}} C^i(x)(\boldsymbol{K}^{\boldsymbol{\mu}}\boldsymbol{M})(dx) \quad .
\end{aligned}
$$

The last equation is due to the familiar formula for the transformation of variables. **q.e.d.**

We have now obtained two consistent intuitive ways of regarding a mixed strategy; we may either regard the mixture over the pure strategies and the corresponding expectations or we may consider the mixture over the distributions resulting from pure strategies and take the payoff from a process distributed accordingly. Both versions lead to the same payoffs. Hence we have a non-contradictory definition of the normal form in mixed strategies.

Definition 3.12. *Let Σ be a tree game and let $\mathfrak{M}^1, \ldots, \mathfrak{M}^n$ denote the sets of mixed strategies of the players. Define*

$$C_{\bullet}^{i\boldsymbol{\mu}} : \mathfrak{M}^1 \times \cdots \times \mathfrak{M}^n \to \mathbb{R}$$

by (20) or (23) Then the **normal form** *game induced by Σ (in mixed strategies) is the (noncooperative) n-person game*

$$\bar{\Gamma} = \bar{\Gamma}_{\Sigma\boldsymbol{\mu}} = (\mathfrak{M}^1, \ldots, \mathfrak{M}^n; C_{\bullet}^{1\boldsymbol{\mu}}, \ldots, C_{\bullet}^{n\boldsymbol{\mu}}), \tag{24}$$

which is also the mixed extension of $\tilde{\Gamma}_{\Sigma\boldsymbol{\mu}}$.

Before we compare behavioral and mixed strategies we have to consider the relation between a fixed path or play and a pure strategy; clearly not all pure strategies are compatible with a given path.

Definition 3.13. *Let $\bar{\alpha}^i \in \mathfrak{S}^i$ be a pure strategy for player i and let $\bar{x} = (x_0, \ldots, x_T)$ be a path. We shall say that $\bar{\alpha}^i$* ***supports*** *$\bar{x}$ if, for all $t = 1, \ldots, T$ and $\bar{x}_{t-1} \in \mathfrak{X}^i$, it follows that*

$$\sigma_{\bar{x}_{t-1}} \circ \bar{\alpha}^i \circ \boldsymbol{\kappa}^i(\bar{x}_{t-1}) \;=\; \sigma_{\bar{x}_{t-1}} \left(\bar{\alpha}^i(\boldsymbol{\kappa}^i(\bar{x}_{t-1}))\right) = \bar{x}_t \; . \tag{25}$$

That is, if player i at $\bar{x}_{t-1}$ observes $q = \boldsymbol{\kappa}^i(\bar{x}_{t-1})$ and decides for some $s = \bar{\alpha}^i(q)$, then $\bar{x}_t$ is the uniquely defined next node following from this decision.

Remark 3.14. Generally the distribution induced by a pure strategy on the paths of the tree can now be represented as follows:

$$m_{\mu}^{\alpha}(\{x\}) := \left\{ \begin{array}{ll} \prod_{\{t|x_{t-1}\in\mathfrak{X}^0\}} \mu_{x_t}^{x_{t-1}} & \text{if } \alpha \text{ supports } x \\ 0 & \text{otherwise} \end{array} \right\} . \tag{26}$$

∘ ∼∼∼∼∼∼ ∘

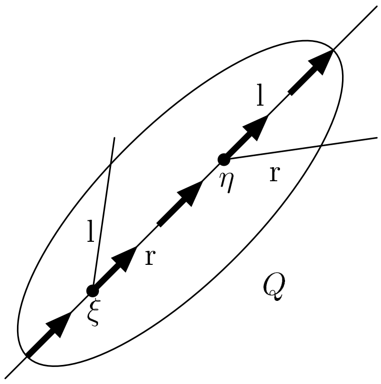

Figure 3.5:

Now it turns out that there occurs a serious problem with our model if a path visits an information more than once. For instance, we can easily construct a path that is not supported by any pure strategy.

Also it is not clear how the specification of a behavioral strategy at q, i.e., $\boldsymbol{A}^i(q, \bullet)$ is interpreted: is the random choice performed twice independently? Or should the result of one random choice be applied both times (in which case some nodes could be superfluous).

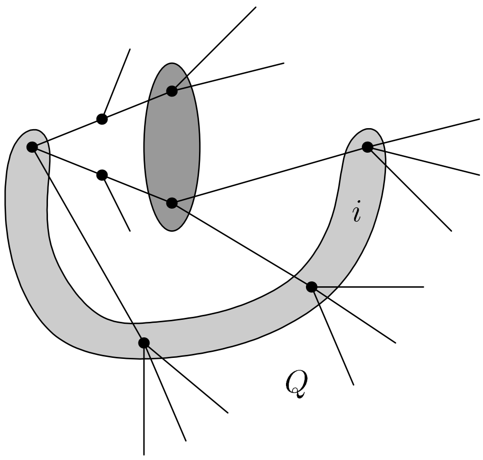

Figure 3.6:

We would like to avoid some problems resulting from this possibility.

So far we have not made an assumption whether player i observes a signal twice because the process reenters an information set or stays therein for several steps.

Definition 3.15. *To every path* $\boldsymbol{x} = (x_0, \ldots, x_t)$ *we associate player* i*'s* ***history*** $\mathcal{K}^i(\boldsymbol{x}) = (q_0, \ldots, q_s)$ *with* $s \leq t$ *and* $q_r \in \mathcal{Q}^i \cup \{\emptyset\}$ $(r = 0, \ldots, s)$

defined recursively as follows.
(27)

$$\mathcal{K}^i(x_0) = \begin{cases} (\emptyset) & x_0 \notin \mathcal{X}^i \\ (\boldsymbol{\kappa}^i(x_0)) & x_0 \in \mathcal{X}^i \end{cases}$$

$$\mathcal{K}^i(x_0, \ldots, x_t) = \begin{cases} (\boldsymbol{\kappa}^i(x_t)) & \mathcal{K}^i(x_0, \ldots, x_{t-1}) = (\emptyset), \ x_t \in \mathcal{X}^i \\ (\mathcal{K}^i(x_0, \ldots, x_{t-1}), \boldsymbol{\kappa}^i(x_t)) & \mathcal{K}^i(x_0, \ldots, x_{t-1}) \neq (\emptyset), \ x_t \in \mathcal{X}^i \\ \mathcal{K}^i(x_0, \ldots, x_{t-1}) & x_t \notin \mathcal{X}^i \end{cases}$$

We shall say that Σ *admits of* ***consistent signals*** *if, for any two paths*

$$\boldsymbol{x} = (x_0, \ldots, x_t), \quad \boldsymbol{y} = (y_0, \ldots, y_{t'})$$

satisfying

$$\boldsymbol{\kappa}^i(x_t) = \boldsymbol{\kappa}^i(y_{t'})$$

it follows that

$$\mathcal{K}^i(x) = \mathcal{K}^i(y)$$

holds true.

Thus, to any path we associate the corresponding sequence of signals player i observes. Feasibly, he observes nothing. However, if he keeps track of his observations, then this will not change his state of information when he observes a new signal. For, if he observes some $q = \boldsymbol{\kappa}^i(\xi)$, then he can compute all possible path that could lead to this signal in the sense, that the endpoint ξ_t satisfies $\boldsymbol{\kappa}^i(x_t) = \boldsymbol{\kappa}^i(\xi)$. All these path yield the same equence of signals. Hence, player i can at once determine the past history he observed based on his last signal – both informations are consistent.

Our first observation is as follows.

Theorem 3.16. *Let* Σ *be a tree game,* $\boldsymbol{A}$ *a behavioral strategy and* $\boldsymbol{M}$ *a mixed strategy. The processes controlled by* $\boldsymbol{A}$ *and* $\boldsymbol{M}$ *have the same distribution if*

$$\text{(28)} \qquad \boldsymbol{M}^i(\{\boldsymbol{\alpha}^i \mid \boldsymbol{\alpha}^i \text{ supports } \bar{x}\}) = \prod_{\{t \mid \bar{x}_{t-1} \in \mathcal{X}^i\}} \boldsymbol{A}^i(\boldsymbol{\kappa}^i(\bar{x}_{t-1}), \boldsymbol{\sigma}_{\bar{x}_{t-1}}(\bar{x}_t))$$

holds true for any play $\bar{x} = (\bar{x}_0, \ldots, \bar{x}_T)$ *and for any* $i \in I$.

Proof: The distribution resulting from the application of $\boldsymbol{M}$ puts on any $x \in \underline{\overline{\mathbf{X}}}$ the weight

$$
\begin{aligned}
\boldsymbol{K}^{\boldsymbol{\mu}}\boldsymbol{M}(\{x\}) &= \sum_{\alpha \in \mathfrak{S}} \boldsymbol{m}^{\alpha}_{\boldsymbol{\mu}}(\{x\})\boldsymbol{M}(\{\alpha\}) \\
&= \sum_{\{\alpha \in \mathfrak{S} \mid \alpha \text{ supports } x\}} \left(\prod_{\{t \mid x_{t-1} \in \mathfrak{X}^0\}} \boldsymbol{\mu}^{x_{t-1}}_{x_t} \right) \boldsymbol{M}(\{\alpha\})
\end{aligned} \tag{29}
$$

in view of Remark (3.14). We continue (29) by

$$
\begin{aligned}
\cdots &= \prod_{\{t \mid x_{t-1} \in \mathfrak{X}^0\}} \boldsymbol{\mu}^{x_{t-1}}_{x_t} \ \boldsymbol{M}(\{\alpha \mid \alpha \ \text{ supports } \ x\}) \\
&= \prod_{\{t \mid x_{t-1} \in \mathfrak{X}^0\}} \boldsymbol{\mu}^{x_{t-1}}_{x_t} \ \prod_{i \neq 0} \boldsymbol{M}^i(\{\boldsymbol{\alpha}^i \mid \boldsymbol{\alpha}^i \ \text{ supports } \ x\})
\end{aligned} \tag{30}
$$

this is so since the players' mixed strategies are combined to the product measure and the event involved has a product character, i.e.,

$$
\{\alpha \mid \alpha \ \text{ supports } \ x\} = \bigcap_{\substack{i \in I \\ i \neq 0}} \{\alpha \mid \boldsymbol{\alpha}^i \ \text{ supports } \ x\}
$$

According to the assumption (28) we may continue formula (30) by

$$
\cdots = \prod_{\{t \mid x_{t-1} \in \mathfrak{X}^0\}} \boldsymbol{\mu}^{x_{t-1}}_{x_t} \ \prod_{i \neq 0} \prod_{\{t \mid x_{t-1} \in \mathfrak{X}^i\}} \boldsymbol{A}^i(\boldsymbol{\kappa}^i(\bar{x}_{t-1}), \boldsymbol{\sigma}_{\bar{x}_{t-1}}(\bar{x}_t))
$$

which, by Definition (14) is indeed $\boldsymbol{m}^{\boldsymbol{A}}_{\boldsymbol{\mu}}(\{x\})$ if we compare Remark **??**, **q.e.d.**

Next let us focus on the imitation of mixed strategies by behavioral strategies.

Definition 3.17. *Let $\boldsymbol{A}$ be a b.s. Then a mixed strategy $\boldsymbol{M}^{\boldsymbol{A}} \in \mathfrak{M}$ is given by*

$$
\boldsymbol{M}^{i\boldsymbol{A}^i}(\{\alpha\}) = \prod_{q \in \mathfrak{Q}^i} \boldsymbol{A}^i(q, \boldsymbol{\alpha}^i(q)) \quad (i \in I) \quad .
$$

That is, player i considers his random devices $\boldsymbol{A}^i(q, \bullet)$ as stochastically independent and this way computes the probability that $\boldsymbol{A}^i$ decides as $\boldsymbol{\alpha}^i$ would do.

Lemma 3.18. *Assume that Σ admits of consistent signals. Let $\boldsymbol{A}$ be a b.s. and $\bar{x} = (x_1, \ldots, x_T)$ a path. Then, for the corresponding mixed strategy $\boldsymbol{M}^{\boldsymbol{A}}$*

$$
\boldsymbol{M}^{i\boldsymbol{A}^i}(\{\boldsymbol{\alpha}^i \mid \boldsymbol{\alpha}^i \ \text{ supports } \ \bar{x}\}) = \prod_{\{t \mid \bar{x}_{t-1} \in \mathfrak{X}^i\}} \boldsymbol{A}^i(\boldsymbol{\kappa}^i(\bar{x}_{t-1}), \boldsymbol{\sigma}_{\bar{x}_{t-1}}(\bar{x}_t))
$$

holds true for $i \in I$.

Proof: By definition of $\boldsymbol{M}^{\boldsymbol{A}}$ we have

$$\begin{aligned}\boldsymbol{M}^{i\boldsymbol{A}^i}(\{\boldsymbol{\alpha}^i \mid \text{ supports } \bar{x}\}) \\ &= \sum_{\{\boldsymbol{\alpha}^i \mid \boldsymbol{\alpha}^i \text{ supports } \bar{x}\}} \prod_{q \in \mathfrak{Q}^i} \boldsymbol{A}^i \ (q, \boldsymbol{\alpha}^i(q)) \\ &= \cdots\end{aligned}$$

With respect to the last product, consider some $\bar{q} \in \mathfrak{Q}^i$ such that *no* $\bar{x}_t$ yields the signal $\bar{q}$. Separate this factor and split the sum according to whether some $\boldsymbol{\alpha}^i$ decides for some $s \in S_{\bar{q}}$ (every $\boldsymbol{\alpha}^i$ decides for some unique s since we have consistent signals!). This yields

$$\begin{aligned}\cdots &= \sum_{\{\boldsymbol{\alpha}^i \mid \boldsymbol{\alpha}^i \text{ supports } \bar{x}\}} \boldsymbol{A}^i(\bar{q}, \boldsymbol{\alpha}^i(\bar{q})) \prod_{\{q \in \mathfrak{Q}^i \mid q \neq \bar{q}\}} \boldsymbol{A}^i(q, \boldsymbol{\alpha}^i(q)) \\ &= \sum_{s \in S_{\bar{q}}} \sum_{\{\boldsymbol{\alpha}^i \mid \boldsymbol{\alpha}^i \text{ supports } \bar{x}, \boldsymbol{\alpha}^i(\bar{q}) = s\}} \boldsymbol{A}^i(\bar{q}, \boldsymbol{\alpha}^i(\bar{q})) \prod_{\{q \in \mathfrak{Q}^i \mid q \neq \bar{q}\}} \boldsymbol{A}^i(q, \boldsymbol{\alpha}^i(q)) \\ &= \sum_{s \in S_{\bar{q}}} \boldsymbol{A}^i(\bar{q}, s) \left(\sum_{\{\boldsymbol{\alpha}^i \mid \boldsymbol{\alpha}^i \text{ supports } \bar{x}, \alpha^i(\bar{q}) = s\}} \prod_{\{q \in \mathfrak{Q}^i \mid q \neq \bar{q}\}} \boldsymbol{A}^i(q, \boldsymbol{\alpha}^i(q)) \right)\end{aligned}$$

Consider the last sum, in which the product does not depend on s. To every $\boldsymbol{\alpha}^i$ supporting $\bar{x}$ which decides for s, i.e., yields $\boldsymbol{\alpha}^i(q) = s$, there corresponds an $\boldsymbol{\alpha}'^i$ which decides $\boldsymbol{\alpha}'^i(q) = s'$ and otherwise looks exactly like $\boldsymbol{\alpha}^i$. This is so far all $s' \neq s$ and hence the sum yields the same value for each $s \in S_{\bar{q}}$. Actually, the common value can be written

$$\sum_{\{\hat{\boldsymbol{\alpha}}^i \in \mathfrak{S}^{i0}\}} \prod_{\{q \in \mathfrak{Q}^i \mid q \neq \bar{q}\}} \boldsymbol{A}^i(q, \boldsymbol{\alpha}^i(q))$$

where we have introduced a set

$$\mathfrak{S}^{i0} := \{\hat{\boldsymbol{\alpha}}^i \in \mathfrak{S}^i \mid \ \hat{\boldsymbol{\alpha}}^i \text{ is a mapping on } \mathfrak{Q}^i - \{\bar{q}\}, \ \hat{\boldsymbol{\alpha}}^i \text{ supports } \bar{x}\}$$

for convenience.

As this common value does not depend on s, we can put it as a common factor in front of the sum

$$\sum_{s \in S_{\bar{q}}} \boldsymbol{A}^i(\bar{q}, s) \ \ ,$$

which, however, equals 1.

Now, consider the case that $\boldsymbol{\kappa}^i(\bar{x}_{t-1}) = \bar{q}$ for $\bar{q} \in \mathcal{Q}^i$, i.e., the path yields (once!) the signal $\bar{q}$. Since $\bar{x}_{t-1}$ yields $\bar{q}$ we know that $\bar{x}_t$ does not yield $\bar{q}$ and there is exactly one $\bar{s}$ such that $\bar{s} = \boldsymbol{\sigma}_{\bar{x}_{t-1}}(\bar{x}_t)$ (we have assumed consistent signals). Hence, any pure strategy supporting $\bar{x}$ (and hence satisfying $\boldsymbol{\alpha}^i(\bar{q}) = \bar{s}$) can be split into $(\boldsymbol{\alpha}^i(\bar{Q}), \hat{\boldsymbol{\alpha}}^i)$ such that $\hat{\boldsymbol{\alpha}}^i$ is a mapping on $\mathcal{Q}^i - \{\bar{q}\}$ (a "restricted strategy"). Vice versa, any mapping $\hat{\boldsymbol{\alpha}}^i$ on $\mathcal{Q}^i - \{\bar{q}\}$ induces a strategy indicated by $(\bar{q}, \hat{\boldsymbol{\alpha}}^i)$.

This way we have a partition

$$\begin{aligned} \{\boldsymbol{\alpha}^i \mid \boldsymbol{\alpha}^i \text{ supports } \bar{x}\} & \\ &= \{(\bar{s}, \hat{\boldsymbol{\alpha}}^i) \mid \hat{\boldsymbol{\alpha}}^i \in \mathfrak{S}^{i0}\} \\ &= \sum \{\hat{\boldsymbol{\alpha}}^i \mid \hat{\boldsymbol{\alpha}}^i \in \mathfrak{S}^{i0}\} \times \{\bar{s}\} \end{aligned}$$

The remaining arguments are the same, so the factor that appears this time is not 1 but

$$\boldsymbol{A}^i(\bar{q}, \bar{s}) \;=\; \boldsymbol{A}^i\left(\bar{q}, \boldsymbol{\sigma}_{\bar{x}_{t-1}}(\bar{x}_t)\right) \;=\; \boldsymbol{A}^i\left(\boldsymbol{\kappa}^i(\bar{x}_{t-1}), \boldsymbol{\sigma}_{\bar{x}_{t-1}}(\bar{x}_t)\right)$$

The further reduction of the quantity follows along the same path and yields the desired product, **q.e.d.**

The reverse direction requires more structure: usually we cannot expect that behavioral strategies are sufficient to imitate mixed ones unless we impose additional requirements on Σ.

Definition 3.19. *Let Σ be a game tree admitting consistent signals. Σ allows for **consistent choices** if the following holds true.*

Let

$$\bar{\boldsymbol{x}} = (\bar{x}_0, \ldots, \bar{x}_t), \bar{\boldsymbol{y}} = (\bar{y}_0, \ldots, \bar{y}_{t'})$$

be two paths satisfying

$$\bar{x}_t, \bar{y}_{t'} \in \mathcal{X}^i, \quad \boldsymbol{\kappa}^i(\bar{x}_t) = \boldsymbol{\kappa}^i(\bar{y}_{t'}).$$

Also let

$$\mathcal{K}^i(\bar{\boldsymbol{x}}) = \mathcal{K}^i(\bar{\boldsymbol{y}}) = (q_1, \ldots, q_l, q)$$

be the sequence of signals observed. (see (27) for the definition of $\mathcal{K}(\bullet)$). Suppose for some h, t, t', u, r that

$$\boldsymbol{\kappa}^i(\bar{x}_{u-1}) = q_h = \boldsymbol{\kappa}^i(\bar{y}_{r-1})$$

holds true. Then it follows that

$$\boldsymbol{\sigma}_{\bar{x}_{u-1}}(\bar{x}_u) = \boldsymbol{\sigma}_{\bar{y}_{r-1}}(\bar{y}_r) \tag{31}$$

holds true.

We offer two interpretations. The first one isprovided by a direct observation: player i finds himself in the position of observing q. Intermediately, he observed some q_h.If he could apply different choices and still arrive at observing q, then he might be able to distinguish between having passed $\bar{x}_u$ and $\bar{y}_r$ – which would contradict our general tendency to provide consistent observation in the model.

Another interpretation is provided by the following Lemma, according to which a strategy can be employed to precisely point out those paths that it is caompatible with.

Lemma 3.20. *Let Σ be a tree game admitting of consistent choices. Let $\bar{\boldsymbol{x}} = (\bar{x}_0, \ldots, \bar{x}_T)$ be a path with $\boldsymbol{\kappa}^i(\bar{x}_T) = q$. Then, if some $\boldsymbol{\alpha}^i$ supports $\bar{\boldsymbol{x}}$, it supports any $\boldsymbol{y} = (y_0, \ldots, y_s)$ with $\boldsymbol{\kappa}^i(y_S) = q$.*

Proof: Since Σ admits of consistent signals, the paths $\bar{\boldsymbol{x}}$ and $\boldsymbol{y}$ generate the same sequence of signals, say

$$\mathcal{K}^i(\bar{\boldsymbol{x}}) = \mathcal{K}^i(\boldsymbol{y}) = (q_1, \ldots, q_l, q)$$

Let $y_{r-1} \in \mathcal{X}^i$ and assume that $\boldsymbol{\kappa}^i(y_{r-1}) = q_h$. Then there is x_{u-1} such that $\boldsymbol{\kappa}^i(_{u-1}) = q_h$ holds true as well. Suppose the strategy $\boldsymbol{\alpha}^i$ decides for $\boldsymbol{\alpha}^i(q_h) = s$. As $\boldsymbol{\alpha}^i$ supports $\bar{\boldsymbol{x}}$ we know that $\boldsymbol{\sigma}_{\bar{x}_{t-1}}(\bar{x}_t) = s$ and because of (31) we have $\boldsymbol{\sigma}_{y_{t-1}}(y_t) = s$ and $\boldsymbol{\sigma}^{-1}_{y_{t-1}}(s) = y_t$. Thus, $\boldsymbol{\alpha}^i$ supports $\boldsymbol{y}$ as well.

q.e.d.

Definition 3.21. *A tree game Σ allows for* ***perfect recall*** *if the following holds true:*

1. *Σ admits of consistent signals.*
2. *Σ admits of consistent choices.*

That is, the history player i observes is consistent with the information set this history leads to as well as with any strategy player i has chosen to generate this history.

Accordingly, we may consider the set of strategies that are consistent with a certain signal. This is the set of all strategies supporting any path that leads to that signal, a set which is now unambiguously specified by the signal.

Definition 3.22. *Let Σ be a tree game with perfect recall.*

1. *Define for $i \in \boldsymbol{I}$*

   $$\mathfrak{S}^i(q) := \{\boldsymbol{\alpha}^i \in \mathfrak{S}^i \mid \boldsymbol{\alpha}^i \text{supports } x = (x_0, \ldots, x_t), x_t \in q \in \mathfrak{Q}^i\}. \tag{32}$$

   to be the set of strategies that ***support a signal*** *$q \in \mathfrak{Q}^i$.*

2. *Let $\boldsymbol{M}^i$ be a mixed strategy for player $i \in \boldsymbol{I}$. The corresponding behavioral strategy for player i is denoted by $\boldsymbol{A}^i = \boldsymbol{A}^{i\boldsymbol{M}^i}$ and defined by*

   $$\boldsymbol{A}^i(q, s) = \frac{\boldsymbol{M}^i(\{\boldsymbol{\alpha}^i \mid \boldsymbol{\alpha}^i \in \mathfrak{S}^i(q) \ \boldsymbol{\alpha}^i(q) = s\})}{\boldsymbol{M}^i(\mathfrak{S}^i(q))}, \tag{33}$$

   provided the denominator is positive. Otherwise the definition is arbitrary.

Theorem 3.23. *Let Σ be a tree game with perfect recall, and let $\boldsymbol{M}^i$ a mixed strategy for player $i \in \boldsymbol{I}$. Let $\boldsymbol{A}^i = \boldsymbol{A}^{i\boldsymbol{M}^i}$ correspond to $\boldsymbol{M}^i$ via Definition 3.22. Then, for any path $\bar{x}$*

$$\boldsymbol{M}^i(\{\boldsymbol{\alpha}^i \mid \boldsymbol{\alpha}^i \text{supports } \bar{x}\}) = \prod_{\{t \in \{1,\ldots,T\} \mid \bar{x}_{t-1} \in \mathfrak{X}^i\}} \boldsymbol{A}^i(\boldsymbol{\kappa}^i(\bar{x}_{t-1}), \boldsymbol{\sigma}_{\bar{x}_{t-1}}(\bar{x}_t))$$

holds true.

Proof: Consider the path $(\bar{x}_0, \ldots, \bar{x}_t)$ and let

$$\{\bar{x}_0, \ldots, \bar{x}_t\} \cap \mathfrak{X}^i = \{\bar{x}_{r_1}, \ldots, \bar{x}_{r_s}\}$$

with $\bar{x}_{r_1}$ preceding $\bar{x}_{r_2,\ldots}$ etc. Correspondingly, there is a unique sequence of information sets or signals

$$\mathcal{K}^i(\bar{x}_{r_1}, \ldots, \bar{x}_{r_{s-1}}) = (q_1, \ldots, q_{s-1}), \quad \boldsymbol{\kappa}^i(\bar{x}_{r_l}) = q_l \ (l = 1, \ldots, s-1)$$

as well as choices generated

$$s_2 := \boldsymbol{\sigma}_{\bar{x}_{r_1}}(\bar{x}_{r_2}), \ldots, s_s := \boldsymbol{\sigma}_{x_{r_{s-1}}}(x_{r_s}).$$

Therefore, we have

$$
\begin{aligned}
&= \prod_{\{r \mid \bar{x}_{r-1} \in \mathcal{X}^i\}} \boldsymbol{A}^i(\boldsymbol{\kappa}^i(\bar{x}_{r-1}), \boldsymbol{\sigma}_{\bar{x}_{t-1}}(\bar{x}_r)) \\
&= \boldsymbol{A}^i(\boldsymbol{\kappa}^i(\bar{x}_{r_1}), \boldsymbol{\sigma}_{\bar{x}_{r_1}}(\bar{x}_{r_1+1})) \cdot \boldsymbol{A}^i(\boldsymbol{\kappa}^i(\bar{x}_{r_2}), \boldsymbol{\sigma}_{\bar{x}_{r_2+1}}(\bar{x}_{r_3})) \cdot \ldots \\
&\ldots \cdot \boldsymbol{A}^i(\boldsymbol{\kappa}^i(\bar{x}_{r_{s-1}}), \boldsymbol{\sigma}^i_{x_{r_{s-1}}}(\bar{x}_{r_{s-1}+1})) \\
&= \boldsymbol{A}^i(q_1, s_2) \cdot \boldsymbol{A}^i(q_2, s_3) \cdot \ldots \cdot \boldsymbol{A}^i(q_{s-1}, s_s) \\
&= \frac{\boldsymbol{M}^i(\{\boldsymbol{\alpha}^i \mid \boldsymbol{\alpha}^i \in \mathfrak{S}^i(q_1),\ \boldsymbol{\alpha}^i(q_1) = s_2\})}{\boldsymbol{M}^i(\mathfrak{S}^i(q_1))} \cdot \\
&\quad \frac{\boldsymbol{M}^i(\{\boldsymbol{\alpha}^i \mid \boldsymbol{\alpha}^i \in \mathfrak{S}^i(q_2),\ \boldsymbol{\alpha}^i(q_2) = s_3\})}{\boldsymbol{M}^i(\mathfrak{S}^i(q_2))} \cdot \\
&\ldots \cdot \frac{\boldsymbol{M}^i(\{\boldsymbol{\alpha}^i \mid \boldsymbol{\alpha}^i \in \mathfrak{S}^i(q_{s-1}),\ \boldsymbol{\alpha}^i(q_{s-1}) = s_s\})}{\boldsymbol{M}^i(\mathfrak{S}^i(q_{s-1}))}
\end{aligned} \tag{34}
$$

Consider the second denominator. We obtain

$$
\begin{aligned}
\boldsymbol{M}^i(\mathfrak{S}^i(q_2)) &= \boldsymbol{M}^i(\{\boldsymbol{\alpha}^i \mid \boldsymbol{\alpha}^i(q_1) = s_1\}) \\
&= \boldsymbol{M}^i(\{\boldsymbol{\alpha}^i \mid \boldsymbol{\alpha}^i \in \mathfrak{S}^i(q_1), \alpha^i(q_1) = s_1\})
\end{aligned}
$$

as every $\boldsymbol{\alpha}^i$ is compatible with observing q_1: it is the first time player 1 observes something and he has not made any decision so far; hence actually $\mathfrak{S}^i(q_1) = \mathfrak{S}^i$ and the first denominator in (34) equals 1. It follows that the second denominator equals the first enumerator etc., so after cancelling we obtain the last enumerator, i.e.,

$$
\begin{aligned}
\ldots &= \boldsymbol{M}^i(\{\boldsymbol{\alpha}^i \mid \boldsymbol{\alpha}^i \in \mathfrak{S}^i(q_{s-1}), \boldsymbol{\alpha}^i(q_{s-1}) = s_s\}) \\
&= \boldsymbol{M}^i(\{\boldsymbol{\alpha}^i \mid \boldsymbol{\alpha}^i(q_1) = s_1, \ldots, \alpha^i(q_{s-1}) = s_s\}) \\
&= \boldsymbol{M}^i(\{\boldsymbol{\alpha}^i \mid \boldsymbol{\alpha}^i \text{supports } \bar{x}\}),
\end{aligned}
$$

q.e.d.

It is well known that the theorem is wrong without the assumption of perfect recall. The standard counterexample can be constructed using a game with one player only:

Example 3.24. Consider the game represented by the following sketch

There is just one player involved and the set of his pure strategies can be written to be

$$
\mathfrak{S}^1 = \{lL, lR, rL, rR\}. \tag{35}
$$

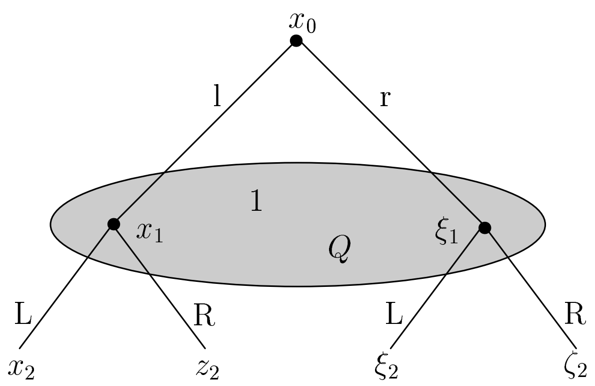

Figure 3.7:

We consider the mixed strategy $\boldsymbol{M} = (\frac{1}{2}, 0, 0, \frac{1}{2})$ and compute the distribution $\boldsymbol{KM}$ induced by $\boldsymbol{M}$ on the paths of our tree. To this end, observe that the distributions $\boldsymbol{m}^{\alpha}$ induced by pure strategies $\alpha \in \mathfrak{S}^1$ are

(36)
$$\begin{array}{rcll} \boldsymbol{m}^{lL} & = & \boldsymbol{\delta}_{\{x_0,x_1,x_2\}} & \text{for} \quad \alpha = lL \\ \boldsymbol{m}^{rR} & = & \boldsymbol{\delta}_{\{x_0,\xi_1,\zeta_2\}} & \text{for} \quad \alpha = rR. \end{array}$$

From this the distribution induced by $\boldsymbol{M}$ is computed to be

(37)
$$\boldsymbol{KM} \;=\; \int_{\mathfrak{S}^1} \boldsymbol{m}^{\alpha}(\bullet)\boldsymbol{M}(d\alpha) \;=\; \frac{1}{2}\left(\boldsymbol{\delta}_{\{x_0,x_1,x_2\}} + \boldsymbol{\delta}_{\{x_0,\xi_1,\zeta_2\}}\right).$$

We now claim, that this distribution cannot be generated by a behavioral strategy. Indeed, take an arbitrary behavioral strategy $\boldsymbol{A}$ which can be written

(38)
$$\begin{array}{rcl} \boldsymbol{A} & = & (A(x_0,\bullet), A(q,\bullet)) \\ & = & ((p, 1-p), (q, 1-q)) \\ & & \;\;\; l \quad\; r \qquad\; L \quad\; R \end{array}$$

The distribution generated by $\boldsymbol{A}$ may then be described as follows:

(39)
$$\begin{array}{rcl} \boldsymbol{m}^{\boldsymbol{A}}(\{x_0,x_1,x_2\}) & = & pq \\ \boldsymbol{m}^{\boldsymbol{A}}(\{x_0,x_1,z_2\}) & = & p(1-q) \\ \boldsymbol{m}^{\boldsymbol{A}}(\{x_0,\xi_1,\xi_2\}) & = & (1-p)q \\ \boldsymbol{m}^{\boldsymbol{A}}(\{x_0,\xi_1,\zeta_2\}) & = & (1-p)(1-q) \end{array}$$

Now, $\boldsymbol{KM}$ puts mass 0 on both the “inner” paths. Therefore, in order for $\boldsymbol{A}$ to imitate $\boldsymbol{M}$ with respect to the distribution, we must necessarily have

(40)
$$0 \;=\; p(1-q) \;=\; (1-p)q,$$

that is, we have either

$$p = 0 \text{ and } q = 0$$

or

$$p = 1 \text{ and } q = 1.$$

Accordingly it follows that either

$$\boldsymbol{m}^{\boldsymbol{A}} = \boldsymbol{\delta}_{\{x_0,\xi_1,\zeta_2\}}$$

or

$$\boldsymbol{m}^{\boldsymbol{A}} = \boldsymbol{\delta}_{\{x_0,x_1,x_2\}},$$

which is never $\boldsymbol{KM}$. Thus there is no behavioral strategy which can imitate $\boldsymbol{M}$ by generating the same distribution on the paths of our tree.

∘ ∼∼∼∼∼∼∼ ∘

Theorem 3.25. *Let Σ be a tree game with perfect recall. Let $\boldsymbol{A}$ a behavioral strategy and $\boldsymbol{M}$ a mixed strategy. Suppose that both $\boldsymbol{A}$ and $\boldsymbol{M}$ generate the same distribution, i.e., we have*

$$\boldsymbol{K}^{\mu}\boldsymbol{M} = \boldsymbol{m}^{\boldsymbol{A}}_{\boldsymbol{\mu}} \quad .$$

Then it follows that the payoffs to all players are the same, i.e.,

$$C^{i\boldsymbol{\mu}}_{\boldsymbol{M}} = C^{i\boldsymbol{\mu}}_{\boldsymbol{A}} \qquad (i \in \boldsymbol{I})$$

holds true.

Proof: once.

Recall the "evaluation", i.e., the mapping

$$\begin{aligned} &C^i : \underline{\overline{\mathbf{X}}} \longrightarrow \mathbb{R} \\ &C^i(x) = u^i(x_{\boldsymbol{\tau}}) + \sum_{t=1}^{\boldsymbol{\tau}} f^i(x_t), \end{aligned} \tag{41}$$

where $\boldsymbol{\tau}$ is the length of a play. This mapping evaluates the payoffs for player i along a path. Then we have

$$C^{i\boldsymbol{\mu}}_{\boldsymbol{A}} = E_{\boldsymbol{m}^{\boldsymbol{A}}_{\boldsymbol{\mu}}} C^i = \int C^i \, d\boldsymbol{m}^{\boldsymbol{A}}_{\boldsymbol{\mu}}. \tag{42}$$

Similarly, for every $\alpha \in \mathfrak{S}$

$$
\begin{aligned}
C^{i\boldsymbol{\mu}}_{\alpha} &= C^{i\boldsymbol{\mu}}_{\boldsymbol{A}^{\alpha}} = \int C^i(x)\, d\boldsymbol{m}^{\boldsymbol{A}^{\alpha}}_{\boldsymbol{\mu}}(x) \\
&= \int C^i(x)\boldsymbol{K}^{\boldsymbol{\mu}}(\alpha, dx) = \boldsymbol{K}^{\boldsymbol{\mu}} C^i(\alpha).
\end{aligned} \tag{43}
$$

By the general formula for the transformation of variables we find

$$
\begin{aligned}
C^{i\boldsymbol{\mu}}_{\boldsymbol{M}} &= \int C^{i\boldsymbol{\mu}}_{\alpha}\, d\boldsymbol{M}(\alpha) \\
&= \int (\boldsymbol{K}^{\boldsymbol{\mu}} C^i)\, d\boldsymbol{M} \\
&= \int C^i\, d(\boldsymbol{K}^{\boldsymbol{\mu}} \boldsymbol{M}) = \int C^i\, d\boldsymbol{m}^{\boldsymbol{A}}_{\boldsymbol{\mu}} \\
&= C^{i\boldsymbol{\mu}}_{\boldsymbol{A}},
\end{aligned} \tag{44}
$$

q.e.d.

Concluding we observe the following Theorem which shows that it suffices to consider behavioral strategies.

Theorem 3.26 (KUHN's Theorem). *Let Σ be a tree game with perfect recall and let $\boldsymbol{\mu}^{\bullet}$ be the corresponding family of distributions of the random moves.*

1. *Let $\bar{\boldsymbol{A}} \in \mathcal{A}$ be a Nash equilibrium in behavioral strategies (i.e., for $\Gamma_{\Sigma,\boldsymbol{\mu}}$). Then $\boldsymbol{M}^{\bar{\boldsymbol{A}}} \in \mathfrak{M}$ is a Nash equilibrium in mixed strategies (i.e. for $\bar{\Gamma}_{\Sigma,\boldsymbol{\mu}}$). In particular, if Σ is a zero-sum two person game then (so are $\bar{\Gamma}_{\Sigma,\boldsymbol{\mu}}$ and $\Gamma_{\Sigma,\boldsymbol{\mu}}$ and)*
$$v_{\Gamma_{\Sigma,\boldsymbol{\mu}}} = v_{\bar{\Gamma}_{\Sigma,\boldsymbol{\mu}}}$$
2. *If $\bar{\boldsymbol{M}}$ is a Nash equilibrium (for $\bar{\Gamma}_{\Sigma,\boldsymbol{\mu}}$) then $\boldsymbol{A}^{\bar{\boldsymbol{M}}}$ is a Nash equilibrium (for $\Gamma_{\Sigma,\boldsymbol{\mu}}$).*
3. *In any equilibrium a player cannot improve upon his payoff by deviating to the other type of strategy.*

Proof: Consider the first statement. Observe that we have defined the mappings

$$\alpha \to \boldsymbol{A}^{\alpha} \quad , \quad \boldsymbol{A} \to \boldsymbol{M}^{\boldsymbol{A}} \quad , \quad \alpha \to \boldsymbol{M}^{\alpha}$$

in such a way that

$$\boldsymbol{M}^{\boldsymbol{A}^{\alpha}} = \boldsymbol{M}^{\alpha} = \boldsymbol{\delta}_{\alpha}$$

is verified at once. We know by Theorem 3.16 and Theorem 3.25 that $\boldsymbol{A}$ and $\boldsymbol{M}^{\boldsymbol{A}}$ always yield the same payoff.

Next, let $\bar{\boldsymbol{A}}$ be an equilibrium in $\Gamma_{\Sigma,\boldsymbol{\mu}}$. We are to show that $\boldsymbol{M}^{\bar{\boldsymbol{A}}}$ is an equilibrium in $\bar{\Gamma}_{\Sigma,\boldsymbol{\mu}}$ (the statement concerning the values follows then immediately). For simplicity, let us assume $n = 2$, the general n is just as simple. Assume now that "player 1 deviates" (in $\bar{\Gamma}_{\Sigma,\boldsymbol{\mu}}$) i.e., he plays a mixed strategy $\hat{\boldsymbol{M}}^1$. We have to show that

$$C^{1\boldsymbol{\mu}}_{\boldsymbol{M}^{\bar{\boldsymbol{A}}}} \geq C^{1\boldsymbol{\mu}}_{\hat{\boldsymbol{M}}^1,\boldsymbol{M}^{2,\bar{\boldsymbol{A}}^2}} \tag{45}$$

(and a similar relation for player 2, however, the argument is just symmetric).

Now, as $\bar{\boldsymbol{A}}$ is an equilibrium,

$$C^{1\boldsymbol{\mu}}_{\boldsymbol{M}^{\bar{A}}} = C^{1\boldsymbol{\mu}}_{\bar{\boldsymbol{A}}} \geq C^{1\boldsymbol{\mu}}_{(\boldsymbol{A}^1,\bar{\boldsymbol{A}}^2)} \qquad \left(\boldsymbol{A}^1 \in \mathfrak{A}^1\right) \tag{46}$$

for any b.s. $\boldsymbol{A}^1$ of player 1. In particular, player 1 may "play pure", i.e., we may choose $\boldsymbol{A}^{1,\alpha^1}$ for any $\alpha^1 \in \mathfrak{S}^1$. Let us identify α^1, $\boldsymbol{A}^{1,\alpha^1}$, $\boldsymbol{M}^{1,\alpha^1}$ etc. and insert this in (46), thus obtaining

$$C^{1\boldsymbol{\mu}}_{\boldsymbol{M}^{\bar{\boldsymbol{A}}}} \geq C^{1\boldsymbol{\mu}}_{(\alpha^1,\bar{\boldsymbol{A}}^2)} \tag{47}$$

Define for the moment $\hat{\boldsymbol{A}} := (\alpha^1, \bar{\boldsymbol{A}}^2)$ such that (47) may be extended to

$$\begin{aligned}
C^{1\boldsymbol{\mu}}_{\boldsymbol{M}^{\bar{\boldsymbol{A}}}} \geq C^{1\boldsymbol{\mu}}_{\hat{\boldsymbol{A}}} \qquad & = C^{1\boldsymbol{\mu}}_{\boldsymbol{M}^{\hat{\boldsymbol{A}}}} \\
= \int C^{1\boldsymbol{\mu}}_{\alpha} \boldsymbol{M}^{\hat{\boldsymbol{A}}}(d\alpha) \quad & = E_{\boldsymbol{M}^{\hat{\boldsymbol{A}}}} C^{1\boldsymbol{\mu}}_{\bullet} \\
= E_{\boldsymbol{\delta}_{\alpha^1} \otimes \boldsymbol{M}^{2,\bar{\boldsymbol{A}}^2}} C^{1\boldsymbol{\mu}}_{\bullet} \quad & = E_{\boldsymbol{M}^{2,\bar{\boldsymbol{A}}^2}} \left(E_{\boldsymbol{\delta}_{\alpha^1}} C^{1\boldsymbol{\mu}}_{\bullet}\right) \\
& \quad (\text{"Fubini's theorem"}) \\
= E_{\boldsymbol{M}^{2,\bar{\boldsymbol{A}}^2}} C^{1\boldsymbol{\mu}}_{(\alpha^1,\bullet)}. &
\end{aligned} \tag{48}$$

As (48) holds true for all $\alpha^1 \in \mathfrak{A}_1$, we "integrate" with the mixed strategy $\hat{\boldsymbol{M}}^1$ player 1 is using to "deviate" (the left side is a constant w.r.t. α_1 and thus taking the integral does not change its value). We obtain

$$\begin{aligned}
C^{1\boldsymbol{\mu}}_{\boldsymbol{M}^{\bar{\boldsymbol{A}}}} & \geq \int E_{\boldsymbol{M}^{2,\bar{\boldsymbol{A}}^2}} C^{1\boldsymbol{\mu}}_{(\alpha^1,\bullet)} \hat{\boldsymbol{M}}^1(d\alpha^1) \\
& = E_{\hat{\boldsymbol{M}}^1 \times \boldsymbol{M}^{2,\bar{\boldsymbol{A}}^2}} C^{1\boldsymbol{\mu}}_{\bullet} \\
& = C^{1\boldsymbol{\mu}}_{\hat{\boldsymbol{M}}^1,\boldsymbol{M}^{2,\bar{\boldsymbol{A}}^2}}
\end{aligned}$$

Thus, we have proved (45) and we are finished with the first statement of the theorem.

Statements 2 and 3 however, follow just as easily by applying Theorem 3.25 accordingly. **q.e.d.**

Corollary 3.27. *Let* Σ *be a tree game with perfect recall. Then* $\Gamma_{\Sigma\boldsymbol{\mu}}$ *has Nash equilibria.*

Proof: Since the game in pure strategies $\tilde{\Gamma}_{\Sigma\boldsymbol{\mu}}$ has finitely many strategies only, its mixed extension $\bar{\Gamma}_{\Sigma\boldsymbol{\mu}}$ is essentially a generalized matrix game, its strategy sets are simplices in Euclidean spaces. This game has Nash equilibria according to NASH's Theorem. According to KUHN's Theorem these Nash equilibria in mixed strategies induce Nash equilibria in behavioral strategies, **q.e.d.**

Proof: A second proof can be enginiered by means of the one deviation priciple (see CHAPTER I, SECTION 3). The version for an imperfect information setup puts an agent at each information set. Accordingly, one has to prove that a Nash equilibrium of the agent normal form yields a Nash equilibrium of the original normal form. This is exactly the statement of ODV. Now, if a player sits at each information set, then mixed strategies and behavioral ones are equivalent and the existence for behavioral strategies follows again from Nash's theorem. **q.e.d.**