Microeconomics IV (Game Theory)

Lecture 4 – Equilibria of Extensive-Form Games

Lasha Chochua

TA: Luka Makharadze

2026

ISET

Introduction

Overview

- We study how players make optimal decisions in dynamic settings.
- Focus is on equilibrium concepts specific to extensive-form games.
- Topics we will cover:
    - Mixed and behavioral strategies
    - Equilibria in behavioral strategies
    - Sequential rationality and Eq'um
    - Computing sequential equilibria
    - Subgame-perfect equilibria
    - Forward induction
- These tools help us analyze rational behavior, beliefs, and credibility in dynamic games.

Mixed and Behavioral Strategies

Mixed Strategies in Extensive-Form Games

- Since Selten (1975) and Kreps & Wilson (1982), equilibria in extensive-form games have gained attention.
- Studying **behavioral strategies** and **sequential equilibria** offers more insight than only strategic-form analysis.
- Before defining those, we must first define **randomized strategies** for extensive-form games.

Notation and Setup

- Let Γ^e be an extensive-form game.
- The set of players is N.
- For each player $i \in N$, let S_i be the set of all possible information states for player i.
- States are disjoint across players: $S_i \cap S_j = \emptyset$ if $i \neq j$.
- Let

$$S^* = \bigcup_{i \in N} S_i$$

denote the set of all information states.

Nodes and Moves

- For player i and state $s \in S_i$:
  - Let Y_s be the set of nodes where i has information state s.
  - Let D_s be the set of moves available to i at state s.
- So, D_s collects all possible actions after nodes in Y_s.

Pure Strategies

- A pure strategy for player i assigns an action at each $s \in S_i$.
- Let C_i be the set of all such strategies:

$$C_i = \times_{s \in S_i} D_s$$

Mixed vs Behavioral Strategies

- Two ways to define randomized strategies in Γ^e:
  1. Normal representation $\rightarrow$ mixed strategies
  2. Multiagent representation $\rightarrow$ behavioral strategies

Mixed-Strategy Profiles

- A mixed-strategy profile of Γ^e is a profile over C_i:

$$\times_{i \in N} \Delta(C_i)$$

Behavioral-Strategy Profiles

- A behavioral-strategy profile of Γ^e randomizes at each information state:

$$\times_{s \in S^*} \Delta(D_s) = \times_{i \in N} \times_{s \in S_i} \Delta(D_s)$$

- This defines a distribution over actions at each information set.

Behavioral-Strategy Profile

- A mixed strategy assigns probabilities over full plans of action.
- A behavioral strategy assigns probabilities to actions at each decision point.

- A behavioral-strategy profile specifies, for each information state, a probability over possible moves.
- It is also called a **scenario**, as it defines behavior at every decision node in the game.

Notation for Behavioral Strategies

- If σ is a behavioral-strategy profile, then:

$$\sigma = (\sigma_i)_{i \in N} = (\sigma_{i,s})_{s \in S_i,\ i \in N}$$

- For player i:

$$\sigma_i = (\sigma_{i,s})_{s \in S_i} \in \times_{s \in S_i} \Delta(D_s)$$

- For each information state s:

$$\sigma_{i,s} = (\sigma_{i,s}(d_s))_{d_s \in D_s} \in \Delta(D_s)$$

Interpretation

- $\sigma_{i,s}(d_s)$ is the **local move probability** assigned to move d_s at information state s.
    - Formally, $\sigma_{i,s}$ is a probability distribution over the available moves D_s.
- A behavioral strategy for player i is any:

$$\sigma_i \in \times_{s \in S_i} \Delta(D_s)$$

Mixed vs Behavioral Strategy

- A **mixed strategy** for player i is a distribution over complete plans:

$$\tau_i \in \Delta(C_i)$$

- $\tau_i(c_i)$ gives the probability that player i follows pure strategy $c_i \in C_i$.
- Mixed strategies randomize over plans; behavioral strategies randomize at each node.

Example – Extensive Game (1/2)

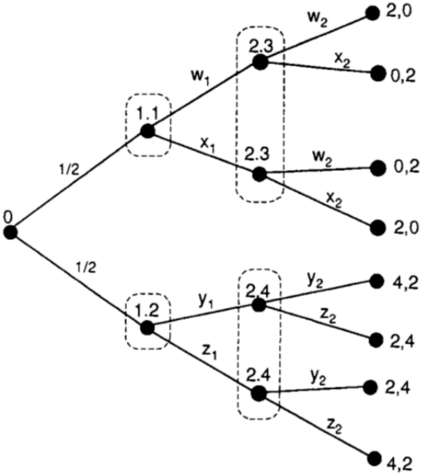

Figure 1: Game 1

Example – Normal representation of the Extensive Game (2/2)

	C_2			
C_1	w_2y_2	w_2z_2	x_2y_2	x_2z_2
w_1y_1	3,1	2,2	2,2	1,3
w_1z_1	2,2	3,1	1,3	2,2
x_1y_1	2,2	1,3	3,1	2,2
x_1z_1	1,3	2,2	2,2	3,1

Multiagent Representation (1/2)

- Temporary agents:

$$N^m = \{1.1, 1.2, 2.3, 2.4\}$$

- Action sets:

$$D_{1.1} = \{w_1, x_1\}, \qquad D_{1.2} = \{y_1, z_1\}$$
$$D_{2.3} = \{w_2, x_2\}, \qquad D_{2.4} = \{y_2, z_2\}$$

- Payoffs are ordered as:

$$(v_{1.1}, v_{1.2}, v_{2.3}, v_{2.4}) = (u_1, u_1, u_2, u_2).$$

Multiagent Representation (2/2)

- If agent 1.2 chooses y_1:

	y_2		z_2	
	w_2	x_2	w_2	x_2
w_1	$3,3,1,1$	$2,2,2,2$	$2,2,2,2$	$1,1,3,3$
x_1	$2,2,2,2$	$3,3,1,1$	$1,1,3,3$	$2,2,2,2$

- If agent 1.2 chooses z_1:

	y_2		z_2	
	w_2	x_2	w_2	x_2
w_1	$2,2,2,2$	$1,1,3,3$	$3,3,1,1$	$2,2,2,2$
x_1	$1,1,3,3$	$2,2,2,2$	$2,2,2,2$	$3,3,1,1$

Mixed vs Behavioral Strategies: Key Distinction

- A **mixed strategy** is a probability distribution over complete plans.
  It is defined on the **normal-form representation** of the extensive game.
- A **behavioral strategy** assigns a probability to each action at each decision node.
  It is defined on the **multiagent representation** (of extensive form).

Example: Equivalence via Local Behavior

- For any numbers α and β between 0 and $1/2$, the mixed-strategy profile

$$(\alpha[w_1y_1] + \alpha[x_1z_1] + (0.5-\alpha)[w_1z_1] + (0.5-\alpha)[x_1y_1]),$$
$$(\beta[w_2y_2] + \beta[x_2z_2] + (0.5-\beta)[w_2z_2] + (0.5-\beta)[x_2y_2])$$

- Although the **plans differ**, each move has **probability 0.5** at its decision point.

Why This Matters

- All these mixed strategies **look different** in the normal form.
- But in the extensive form, they lead to the **same behavior**:

$$(0.5[w_1] + 0.5[x_1],\ 0.5[y_1] + 0.5[z_1],\ 0.5[w_2] + 0.5[x_2],\ 0.5[y_2] + 0.5[z_2])$$

- This is the behavioral strategy: randomizing at each information state.

Marginal Probabilities at Player 1's Information Sets

- At information state **1.1** (node 1.1), what is the chance of playing w_1?
- Strategies involving w_1:
  $w_1 y_1$ and $w_1 z_1$ $\rightarrow$ weight $= \alpha + (0.5 - \alpha) = 0.5$
- Strategies involving x_1:
  $x_1 y_1$ and $x_1 z_1$ $\rightarrow$ weight $= 0.5$
- $\rightarrow$ Therefore:

$$P(w_1) = 0.5, \quad P(x_1) = 0.5$$

Marginal Probabilities at Information State 1.2

- Strategies involving y_1:
  $w_1 y_1$ and $x_1 y_1 \rightarrow$ weight $= \alpha + (0.5 - \alpha) = 0.5$
- Strategies involving z_1:
  $w_1 z_1$ and $x_1 z_1 \rightarrow$ weight $= 0.5$

- $\rightarrow$ Therefore:

$$P(y_1) = 0.5, \quad P(z_1) = 0.5$$

Misleading Equivalence Between Mixed and Behavioral Strategies

- Consider the game (extensive form with perfect recall).

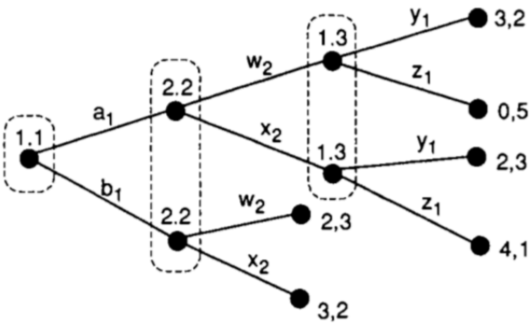

Figure 2: Game 2

Misleading Equivalence Between Mixed and Behavioral Strategies

- Let Player 1 use the mixed strategy:

$$0.5[a_1 y_1] + 0.5[b_1 z_1]$$

- At first, this may seem to correspond to the behavioral strategy:

$$(0.5[a_1] + 0.5[b_1],\ 0.5[y_1] + 0.5[z_1])$$

- But this is **incorrect**. Let's see why.

Why It's Wrong: Conditional Play at 1.3

- Under the mixed strategy $0.5[a_1 y_1] + 0.5[b_1 z_1]$, the path reaches Player 1's node 1.3 only if a_1 is played.
- So:
  - If the game reaches 1.3, the pure strategy must be $a_1 y_1$
  - Therefore, Player 1 chooses y_1 with **probability 1** at 1.3
  - z_1 is never chosen at 1.3, even though it has weight 0.5 in the plan $b_1 z_1$

The Correct Behavioral Strategy

- To convert the mixed strategy $0.5[a_1 y_1] + 0.5[b_1 z_1]$ to a behavioral strategy:
  - At node 1.1:
    $0.5[a_1] + 0.5[b_1]$
  - At node 1.3:
    Only y_1 is used when reachable
- $\rightarrow$ So the correct behavioral strategy is:

$$(0.5[a_1] + 0.5[b_1],\ y_1)$$

Key Insight

- You cannot infer local randomization at information states just by looking at weights in mixed strategies.
- Behavioral strategies specify **what the player would do** at each node **if reached**, while mixed strategies include **entire plans**, some of which may never be realized.
- $\rightarrow$ To convert a mixed strategy into a behavioral one, you must **condition on reachability**.

Compatibility of Pure Strategies and Information States

- Let $c_i \in C_i$ be a pure strategy and $s \in S_i$ an information state for player i.
- We say s and c_i are **compatible** if there's some strategy profile $c = (c_{-i}, c_i)$ such that:

$$\sum_{x \in Y_s} P(x \mid c) > 0$$

  where Y_s is the set of nodes in information state s, and $P(x \mid c)$ is the probability of node x under profile c.
- Let $C_i^*(s)$ be the set of all c_i compatible with s.

- *Intuition:* $C_i^*(s)$ is the set of player i's plans that don't *self-sabotage* their way out of s – the plans under which s is still "live"

Compatibility for Mixed Strategies

- Let $\tau_i \in \Delta(C_i)$ be a mixed strategy for player i.
- We say $s \in S_i$ is **compatible with** τ_i if:

$$\exists\, c_i \in C_i^*(s) \text{ such that } \tau_i(c_i) > 0$$

- Define $C_i^{**}(d_s, s)$ as:

$$C_i^{**}(d_s, s) = \{c_i \in C_i^*(s) \mid c_i(s) = d_s\}$$

- That is, $C_i^{**}(d_s, s)$ includes all pure strategies c_i compatible with s that specify move d_s at s.

Constructing $C_i^{**}(d_s, s)$ – Two Filters

- Building $C_i^{**}(d_s, s)$ applies **two filters** to the full plan set:
- **Filter 1 – Reachability:** $c_i \in C_i^*(s)$
  - keep only plans that can actually reach s
  - the compatibility filter from before – about the *path into* s
- **Filter 2 – Action at s:** $c_i(s) = d_s$
  - among those, keep only plans whose instruction *at* s is the move d_s
  - about the *choice made at* s itself
- So $C_i^{**}(d_s, s)$ is the set of plans that **both**:
  - get you to s, **and**
  - tell you to play d_s once you're there
- A plan must pass **both** filters to count. Filter 1 concerns the path leading *into* s; Filter 2 concerns the action taken *at* s.

Behavioral Representation of a Mixed Strategy

- A behavioral strategy $\sigma_i = (\sigma_{i,s})_{s \in S_i}$ is a **behavioral representation** of a mixed strategy $\tau_i \in \Delta(C_i)$ if:

  For all $s \in S_i$ and all $d_s \in D_s$,

$$\sigma_{i,s}(d_s)\left(\sum_{e_i \in C_i^*(s)} \tau_i(e_i)\right) = \sum_{c_i \in C_i^{**}(d_s,s)} \tau_i(c_i) \quad (1)$$

- This ensures that $\sigma_{i,s}(d_s)$ is the **conditional probability** that i would play d_s at s, given that a strategy compatible with s was used.

From Defining Equation to Conditional Probability

- Defining equation (1), with the two sums named:

$$\sigma_{i,s}(d_s)\underbrace{\left(\sum_{e_i\in C_i^*(s)}\tau_i(e_i)\right)}_{\text{denominator } D}=\underbrace{\sum_{c_i\in C_i^{**}(d_s,s)}\tau_i(c_i)}_{\text{numerator } N}$$

- Whenever $D>0$, divide both sides by D:

$$\boxed{\sigma_{i,s}(d_s)=\frac{\displaystyle\sum_{c_i\in C_i^{**}(d_s,s)}\tau_i(c_i)}{\displaystyle\sum_{e_i\in C_i^*(s)}\tau_i(e_i)}}$$

- This is exactly the conversion formula from the earlier slides.

The Conditional-Probability Reading

- Draw a single pure plan from the lottery τ_i, and define two events:
  - $A = \{\text{drawn plan} \in C_i^*(s)\}$ – the plan keeps s live; $P(A) = D$
  - $B = \{\text{drawn plan prescribes } d_s \text{ at } s\}$
- Since $C_i^{**}(d_s, s) = C_i^*(s) \cap \{c_i(s) = d_s\}$, the numerator is $N = P(A \cap B)$. Hence

$$\sigma_{i,s}(d_s) = \frac{P(A \cap B)}{P(A)} = P(d_s \text{ at } s \mid \text{plan compatible with } s)$$

- This is the textbook conditional probability $P(B \mid A)$ – the slide's wording in symbols.

Behavioral Representation of Strategy Profiles

- A strategy profile $\sigma = (\sigma_i)_{i \in N}$ is a behavioral representation of a mixed-strategy profile $\tau = (\tau_i)_{i \in N}$ if each σ_i is a behavioral representation of τ_i.
- Any $\tau_i \in \Delta(C_i)$ that is compatible with $s \in S_i$ admits **at least one** behavioral representation σ_i satisfying (1).

Mixed Representation of Behavioral Strategies

- Given $\sigma_i = (\sigma_{i,s})_{s \in S_i}$ in $\times_{s \in S_i} \Delta(D_s)$, define the mixed representation $\tau_i \in \Delta(C_i)$ as:

$$\tau_i(c_i) = \prod_{s \in S_i} \sigma_{i,s}(c_i(s)) \tag{2}$$

for all $c_i \in C_i$

Interpretation

- Equation (2) says:

  The mixed strategy τ_i assigns probability to c_i as the product of the behavioral probabilities of the actions prescribed by c_i at each state.

- This construction assumes **independence** across information sets.

- $\rightarrow$ A behavioral strategy randomizes **locally**; its mixed representation randomizes **over full plans**, based on local behavior.

Equivalence

- If τ_i is the mixed representation of σ_i, then σ_i is a behavioral representation of τ_i.
- A behavioral-strategy profile σ has a mixed representation τ such that:

$$\forall i \in N, \quad \tau_i(c_i) = \prod_{s \in S_i} \sigma_{i,s}(c_i(s))$$

- The behavioral and mixed representations are equivalent **as long as the game has perfect recall**.

Behavioral and Payoff Equivalence

- Two mixed strategies in $\Delta(C_i)$ are **behaviorally equivalent**
    - **if they share a common behavioral representation.**
- That is, they induce the same local action probabilities at each information state.

Payoff Equivalence

- Two mixed strategies τ_i and ρ_i in $\Delta(C_i)$ are **payoff equivalent** if for every player j and every profile $\tau_{-i} \in \times_{\ell \in N \setminus \{i\}} \Delta(C_\ell)$,

$$u_j(\tau_{-i}, \tau_i) = u_j(\tau_{-i}, \rho_i)$$

  where $u_j(\cdot)$ is j's utility function in the normal-form representation of Γ^e.
- This means that τ_i and ρ_i lead to the same expected utility for all players.

Theorem: Kuhn's Equivalence Result

Theorem 1 (Kuhn's Equivalence Result)

If Γ^e is a game with **perfect recall**, then any two mixed strategies in $\Delta(C_i)$ that are **behaviorally equivalent** are also **payoff equivalent**.

Equilibria in Behavioral Strategies

Motivation for Behavioral Equilibrium

- An equilibrium defined only via the **normal representation** may produce many equivalent mixed-strategy equilibria.
- An equilibrium defined only via the **multiagent representation** may allow players to assign inconsistent actions across information states.
- $\rightarrow$ Both approaches have shortcomings when used in isolation.

Equilibrium in Behavioral Strategies

- To resolve this, we define a **(Nash) equilibrium in behavioral strategies** as follows:

Equilibrium in behavioral strategies

- A behavioral strategy profile σ is an equilibrium of Γ^e if:
  - σ is an equilibrium in the multiagent representation, and
  - its **mixed representation** is an equilibrium of the normal representation.

- This ensures that player behavior is both **locally rational** (at each node) and **globally consistent** across their information sets.

Consistency Across Representations

- Let Γ^e be an extensive-form game with perfect recall, and let Γ be its normal representation.
- If τ is a Nash equilibrium in Γ, and $\hat{\tau}$ is behaviorally equivalent to τ, then:

$$\hat{\tau} \text{ is also a Nash equilibrium of } \Gamma$$

- So, σ is an equilibrium of Γ^e if and only if:
    - **Every mixed strategy profile with σ as a behavioral representation is a Nash equilibrium of Γ.**

Practical Implication

- To find an equilibrium of Γ^e, it is sufficient to find an equilibrium of the normal representation Γ.
- By **Kuhn's theorem**, every equilibrium of Γ corresponds to an equilibrium of the multiagent representation.
- $\rightarrow$ This justifies solving Γ^e via its normal form, as long as we interpret results behaviorally.

Theorems on Behavioral Equilibrium Existence

Theorem 2

If Γ^e is an extensive-form game with perfect recall and τ is an equilibrium of the normal representation of Γ^e, then any behavioral representation of τ is an equilibrium of the multiagent representation of Γ^e.

Theorem 3

For any extensive-form game Γ^e with perfect recall, a Nash equilibrium in behavioral strategies exists.

Sequential Rationality

Sequential Rationality

- A strategy σ_i is **sequentially rational** for player i at information state $s \in S_i$ if player i would actually want to follow σ_i at s in case s is reached.
- This concept tests **whether a player's plan remains optimal when actually implemented**, rather than simply being part of a pre-chosen complete strategy.

Dropping the Prior-Planning Assumption

- Traditional analysis assumes players choose full strategies at the game's beginning and follow them mechanically – even at unreached information sets.
- This assumption is often unrealistic.
  Instead, we check whether strategies are optimal **at each stage**, given **beliefs** about reaching that stage.

Probabilities Along Paths

- In an extensive-form game Γ^e, the game tree structure specifies **chance probabilities**.
- A behavioral strategy profile σ specifies a **move probability** $\sigma_{i,s}(m)$ for every branch with move label m at player i's node in information state s.
- Let $\bar{P}(y \mid \sigma, x)$ be the product of chance probabilities and behavioral move probabilities along the path from node x to terminal node y.
- If y does not follow x, then $\bar{P}(y \mid \sigma, x) = 0$

Expected Utility from a Node

- Let Ω be the set of all terminal nodes in Γ^e. Define:

$$U_i(\sigma \mid x) = \sum_{y \in \Omega} \bar{P}(y \mid \sigma, x) w_i(y)$$

- $w_i(y)$ is the payoff to player i at terminal node y.
- $U_i(\sigma \mid x)$ is the **expected utility for player** i if the game starts at node x and all players follow σ from that point onward.

Interpretation

- $U_i(\sigma \mid x)$ depends only on the **components of σ used after node x**.
- This framework allows us to test whether following σ at any point in the tree remains the best response – the foundation of **sequential rationality**.

Sequential Rationality at a Singleton Information State

- Suppose player i has an information state s such that $Y_s = \{x\}$ (i.e., it occurs at only one node).
- Then the behavioral strategy profile σ is **sequentially rational** for player i at s if and only if:

$$\sigma_{i,s} \in \arg \max_{\rho_s \in \Delta(D_s)} U_i(\sigma_{-i,s}, \rho_s \mid x)$$

  - $\sigma_{-i,s}$: rest of σ, excluding player i's behavior at s
  - ρ_s: an alternative behavior at state s
  - The profile $(\sigma_{-i,s}, \rho_s)$ is like σ, except that at s player i follows ρ_s.

Sequential Rationality with Beliefs

- At information states with multiple nodes, expected utility depends on **beliefs** over which node in Y_s was reached.
- For any information state s of player i, a **belief-probability distribution** $\pi_{i,s} \in \Delta(Y_s)$ assigns probabilities over the nodes in Y_s.
- $\pi_{i,s}(x)$ represents player i's belief that node x occurred, given he is at some node in Y_s.

Belief Vectors

- A **beliefs vector** $\pi = (\pi_{i,s})_{i \in N,\, s \in S_i}$ is a profile assigning a belief distribution $\pi_{i,s} \in \Delta(Y_s)$ for each information state s of each player.
- This allows expected utility to be defined conditionally at each s.

Definition of Sequential Rationality with Beliefs

- Given beliefs π, σ is sequentially rational for player i at $s \in S_i$ if:

$$\sigma_{i,s} \in \arg \max_{\rho_s \in \Delta(D_s)} \sum_{x \in Y_s} \pi_{i,s}(x) \cdot U_i(\sigma_{-i,s}, \rho_s \mid x) \tag{3}$$

  - Here, $\sigma_{-i,s}$ is the behavioral profile excluding i's behavior at s.
  - ρ_s is an alternative move distribution at s.

Determining Beliefs: Bayesian Updating

- The definition of sequential rationality raises the question:
  How should a rational player determine belief probabilities?
- Beliefs must reflect **Bayes' rule**, relating observed information during play to initial probabilities.

Prior Probability of Nodes

- Let $\bar{P}(y \mid \sigma)$ be the probability that node y is reached from the root, when all players follow behavioral strategy profile σ.

- That is,
$$\bar{P}(y \mid \sigma) = \bar{P}(y \mid \sigma, x^0)$$
where x^0 is the initial node of the game tree.

- $\bar{P}(y \mid \sigma)$ is the product of all relevant chance and move probabilities.

Weak Consistency of Beliefs

- Let $\pi_{i,s}(x)$ be player i's belief that node $x \in Y_s$ occurred, given s is reached.
- Then, by **Bayes' rule**, beliefs must satisfy:

$$\pi_{i,s}(x) \sum_{y \in Y_s} \bar{P}(y \mid \sigma) = \bar{P}(x \mid \sigma) \tag{4}$$

- A belief vector π is **weakly consistent** with σ if it satisfies (4) for all players, information states s, and nodes $x \in Y_s$.

When Beliefs Are Well-Defined

- Equation (4) characterizes beliefs only if:

$$\sum_{y \in Y_s} \bar{P}(y \mid \sigma) > 0 \tag{5}$$

- In that case, beliefs are fully determined:

$$\pi_{i,s}(x) = \frac{\bar{P}(x \mid \sigma)}{\sum_{y \in Y_s} \bar{P}(y \mid \sigma)} \quad \forall x \in Y_s \tag{6}$$

- Otherwise, s is said to be **off the path of** σ, and beliefs cannot be updated using Bayes' rule.

Sequential Rationality On the Path

Theorem 4

Suppose that σ is an equilibrium in behavioral strategies of an extensive-form game with perfect recall. Let $s \in S_i$ be an information state that occurs with positive probability under σ (so condition (4) is satisfied). Let π be a belief vector weakly consistent with σ. Then σ is sequentially rational for player i at s with beliefs π.

Representing Equilibrium and Beliefs in Extensive Form

- When diagramming an equilibrium in an extensive-form game, we must show both:
  - Move probabilities from the behavioral strategy profile σ
  - Belief probabilities from a beliefs vector π that is consistent with σ

Notation convention:

- **Move probabilities** (from σ) are shown in **parentheses**: e.g., (0.5)
- **Belief probabilities** (from π) are shown in **angle brackets**: e.g., $\langle 0.25 \rangle$

- At each **branch** after a player's decision node:
  $\rightarrow$ show the move probability (component of σ)
- At each **decision node**:
  $\rightarrow$ show the belief probability (component of π)

Full Consistency of Beliefs

- A belief vector π is **fully consistent** with a behavioral strategy profile σ if:

  There exist behavioral strategies arbitrarily close to σ that assign strictly positive probability to every move, and for which the beliefs vectors (satisfying Bayes's formula) converge to π.

- That is, π arises as the limit of beliefs computed from sequences of strictly positive strategy profiles near σ.

Formal Definition and Consequences

- Let Ψ be the set of all pairs (σ, π) such that π is fully consistent with σ.
- Then Ψ is the **smallest closed subset** of:

$$(\times_{s \in S^*} \Delta(D_s)) \times (\times_{s \in S^*} \Delta(Y_s))$$

that contains all pairs (σ^0, π^0) derived from **strictly positive** strategy profiles σ^0 and their associated beliefs π^0.

- Full consistency implies **weak consistency**:
  - If π is fully consistent with σ, then it also satisfies the Bayes condition (4.3) wherever defined.

Intuition Behind Full Consistency

- Sometimes, players assign zero probability to certain moves – this creates trouble for applying Bayes' rule at some nodes.
- So instead, we imagine a nearby strategy σ^0 where:
  - Every move has **positive** probability (no zeros!)
  - Players still behave **almost like** σ

- We compute beliefs π^0 from such σ^0 using Bayes' rule (since all paths are possible).
- Then we take the **limit** of these beliefs as $\sigma^0 \to \sigma$.
  $\to$ The limit π is said to be **fully consistent** with σ.

- Think of it this way:
  - If we *perturb* σ just a little so that all moves are possible,
  - then compute beliefs normally, and
  - those beliefs converge to π,
  - then we say π is fully consistent with σ.

Definition: Sequential Equilibrium

Definition: Sequential Equilibrium (Kreps and Wilson, 1982)

A *sequential equilibrium* (or a *full sequential equilibrium*) of Γ^e is any pair (σ, π) in $(\times_{s \in S^*} \Delta(D_s)) \times (\times_{s \in S^*} \Delta(Y_s))$ such that:

- The belief vector π is fully consistent with σ, and
- With beliefs π, the scenario σ is sequentially rational for every player at every information state.

A behavioral-strategy profile σ is a *sequential-equilibrium scenario* if there exists some π such that (σ, π) is a sequential equilibrium. That is, σ can be *extended* to a sequential equilibrium.

Theorems on Sequential Equilibrium

Theorem 5

If (σ, π) is a sequential equilibrium of an extensive-form game with perfect recall, then σ is an equilibrium in behavioral strategies.

Theorem 6

For any finite extensive-form game, the set of sequential equilibria is nonempty.

Behavioral vs Sequential Equilibrium

The Core Question

- Both concepts ask the same thing: **is each player choosing a best response?**
- They differ in two respects:
    - **Where** optimality is demanded – on the path only, or at every information state.
    - **What** supports that demand – expected payoff from the root, or payoff conditional on beliefs.
- The whole distinction lives **off the equilibrium path**.

Behavioral (Nash) Equilibrium – Ex Ante View

- A behavioral profile σ is an equilibrium if no player can raise their **expected payoff, computed from the root**, by deviating.
- Recall the root-level expected utility:

$$U_i(\sigma) = U_i(\sigma \mid x^0) = \sum_{y \in \Omega} \bar{P}(y \mid \sigma)\, w_i(y)$$

- Because payoffs are evaluated from the start, behavior only matters at information states reached with **positive probability**.

- $\rightarrow$ At states **off the path** of σ, a player may specify essentially anything: what they "would do" there never enters $U_i(\sigma)$.

The Loophole

- Off-path behavior is unconstrained under Nash.
- This is exactly what lets **non-credible threats** survive as equilibria.

> **Note**
>
> A player can "threaten" a move at a node that is never reached. Since the threat is never tested, Nash equilibrium does not require it to be optimal – even when carrying it out would hurt the threatener.

Sequential Equilibrium – Closing the Loophole

- Sequential equilibrium demands optimality at **every** information state, including unreached ones.
- To ask "is this move optimal here?" at an off-path node, the player needs a **belief** over which node in Y_s was reached.
- A sequential equilibrium is a **pair** (σ, π) satisfying two conditions:

Two Requirements

1. **Full consistency:** π is the limit of Bayes-updated beliefs from strictly positive perturbations of σ.
2. **Sequential rationality:** given π, the player would actually want to follow σ at every s.

The One-Line Difference

Behavioral vs Sequential

A **behavioral (Nash) equilibrium** requires rational play only **on** the equilibrium path. A **sequential equilibrium** requires rational play at **every** information set, supported by beliefs consistent with the strategies even **off** the path.

Side by Side

Feature	Behavioral (Nash)	Sequential
Object	profile σ	pair (σ, π)
Optimality demanded	on-path states only	every information state
Beliefs	not required	required, fully consistent
Off-path play	unconstrained	disciplined by beliefs

They Coincide On the Path

- The added content of sequential equilibrium is **entirely off-path**.

Theorem 4 (restated)

If σ is a behavioral equilibrium and s occurs with positive probability under σ, then with weakly consistent beliefs σ is **already** sequentially rational at s.

- $\rightarrow$ Sequential equilibrium does not change what happens where you can observe play; it disciplines what players *claim* they would do where you cannot.

A Strict Refinement

- Every sequential equilibrium is a behavioral (Nash) equilibrium (**Theorem 5**), but **not** conversely.

A Subtlety Worth Flagging

- A common student summary: *"Nash ignores beliefs, sequential adds them."*
- That is only **half right**.
- The deeper point is **which beliefs are admissible**:
  - Sequential equilibrium does **not** let a player rationalize an off-path action with any convenient belief.
  - Consistency forces beliefs to be the **limit of genuine Bayesian updating**.
- $\rightarrow$ This is what makes the refinement **bite** rather than being vacuous.

Computing Sequential Equilibria

Computing Sequential Equilibria

- Let's consider the following game:

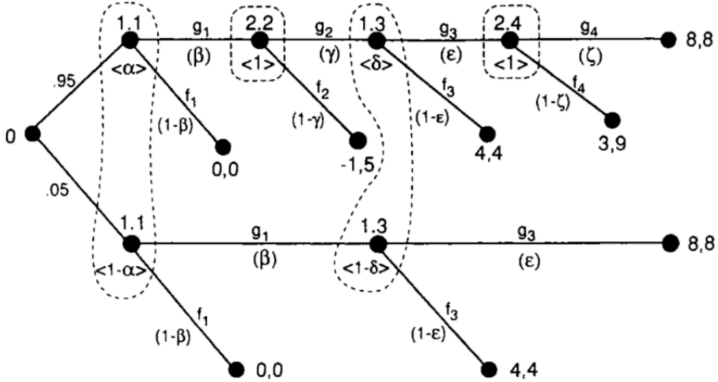

Figure 3: Game 3

Game Interpretation

- Nature chooses:
  - Upper branch (player 2 may be selfish): 0.95
  - Lower branch (player 2 always generous): 0.05
- Player 1 does **not observe** Nature's move.
- Players alternate between:
  - Generous (g_k): costs 1, gives other 5
  - Selfish (f_k): ends game
- Game ends after first selfish move or 2 generous moves by each player.

Variables and Beliefs

- α: Prior belief that player 2 can be selfish $\to \alpha = 0.95$
- δ: Belief 1 assigns that 2 can be selfish after one generous act

- β: Prob. 1 chooses g_1
- γ: Prob. 2 chooses g_2 if selfish
- ε: Prob. 1 chooses g_3
- ζ: Prob. 2 chooses g_4 if selfish

- **Goal:** Find a **sequential equilibrium**
  $\to$ Values of $\beta, \gamma, \varepsilon, \zeta, \delta$
  that make σ and π consistent and sequentially rational.

Easy Deductions

- Two variables are straightforward to determine:
  - $\alpha = 0.95$ – this is the probability of the upper branch in the initial chance move.
  - $\zeta = 0$ – player 2 has no incentive to be generous in the final move if she is capable of selfishness.

Strategy for Solving Sequential Equilibrium

- To find other components of the sequential equilibrium, use the concept of **support**:
  - The **support** of a strategy is the set of moves played with positive probability.

- At each information state:
  1. Guess the support of the sequential equilibrium.
  2. Use sequential rationality to verify or reject this guess.
  3. Repeat with alternative supports if necessary.
- Solving often proceeds **backward** from terminal nodes.

Game 3

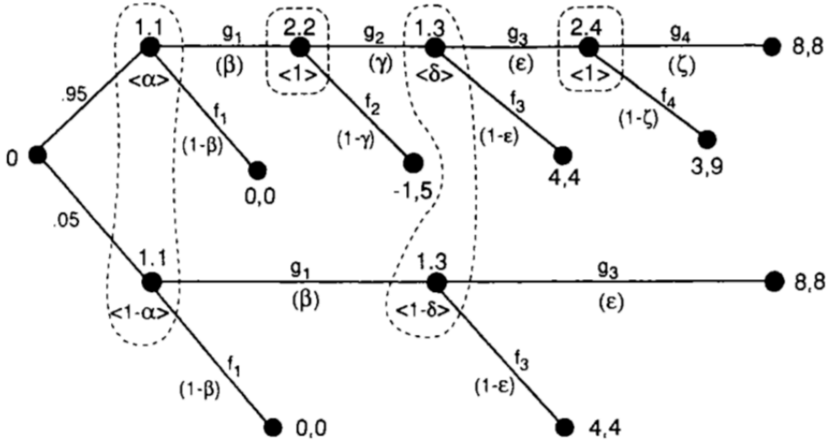

0
.95
.05
1.1
<α>
g_1
(β)
f_1
(1-β)
0,0
2.2
<1>
g_2
(γ)
f_2
(1-γ)
-1,5
1.3
<δ>
g_3
(ε)
f_3
(1-ε)
4,4
2.4
<1>
g_4
(ζ)
f_4
(1-ζ)
3,9
8,8
1.1
<1-α>
g_1
(β)
f_1
(1-β)
0,0
1.3
<1-δ>
g_3
(ε)
f_3
(1-ε)
4,4
8,8

Backward Induction from State 3

- At player 1's second information state (State 3), three possible supports:
  - $\{g_3\}$, $\{f_3\}$, or $\{g_3, f_3\}$
- Given $\zeta = 0$, compute expected payoffs:
  - From g_3: $3\delta + 8(1 - \delta)$
  - From f_3: $4\delta + 4(1 - \delta) = 4$
- From Bayes' formula:

$$\delta = \frac{0.95\beta\gamma}{0.95\beta\gamma + 0.05\beta} = \frac{19\gamma}{19\gamma + 1}$$

Ruling Out $\{g_3\}$ as Support

- Suppose $\varepsilon = 1$ (i.e., support is $\{g_3\}$)
- Sequential rationality at State 3 requires:

$$3\delta + 8(1-\delta) \geq 4 \quad \Rightarrow \quad \delta \leq 0.8$$

- But from Bayes' formula:

$$\delta = \frac{19\gamma}{19\gamma + 1} \Rightarrow \gamma \leq \frac{4}{19}$$

- However, if $\varepsilon = 1$, then:
  - Player 2's expected payoff from g_2 is 9
  - From f_2, it is 5
  - So sequential rationality at her info state requires $\gamma = 1$
- Contradiction: $\gamma = 1$ and $\gamma \leq 4/19$ can't both hold.

$\Rightarrow$ No sequential equilibrium with support $\{g_3\}$

Second Hypothesis: Support is $\{f_3\}$

- Try the hypothesis: Player 1's support at state 3 is $\{f_3\}$
- Then $\varepsilon = 0$
- Sequential rationality at state 3 implies:

$$3\delta + 8(1 - \delta) \leq 4$$

- This inequality implies:

$$\delta \leq 0.2$$

Consistency Condition

- From Bayes' formula:

$$\delta = \frac{19\gamma}{19\gamma + 1}$$

- So the inequality becomes:

$$\frac{19\gamma}{19\gamma + 1} \leq 0.2 \Rightarrow \gamma \geq \frac{4}{19}$$

Contradiction with Player 2's Incentives

- $\varepsilon = 0$ implies:
  - Player 2 expects:
    - g_2: payoff $= 4$
    - f_2: payoff $= 5$
- Sequential rationality at state 2 implies:
  - $\gamma = 0$
- But then: $\gamma = 0$ and $\gamma \geq \frac{4}{19}$ cannot both hold.

$\Rightarrow$ No sequential equilibrium exists where support at Player 1's state 3 is $\{f_3\}$.

Recall the Game 3

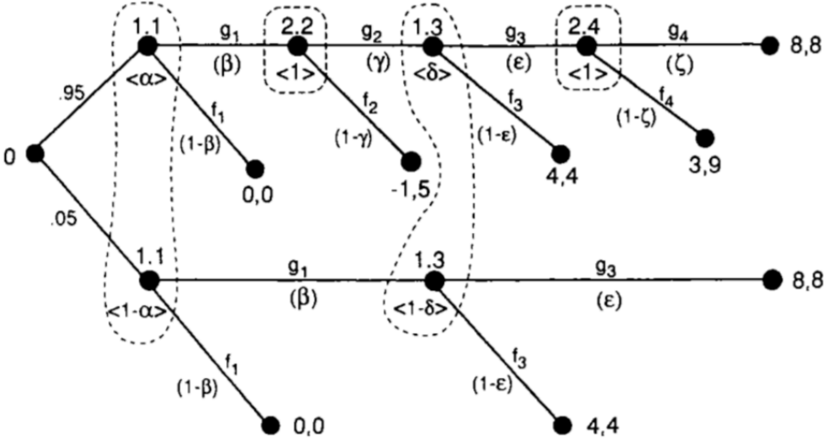

Final Hypothesis: Support is $\{g_3, f_3\}$

- Try the last possible support: Player 1's support at state 3 is $\{g_3, f_3\}$
- Then $0 < \varepsilon < 1$
- Sequential rationality at state 3 implies:

$$3\delta + 8(1 - \delta) = 4 \Rightarrow \delta = 0.8$$

- Using Bayes' formula:

$$\delta = \frac{19\gamma}{19\gamma + 1} \Rightarrow \gamma = \frac{4}{19}$$

Player 2's Payoffs and Choice

- Sequential rationality at state 2 requires:

$$5 = 9\varepsilon + 4(1 - \varepsilon) \Rightarrow \varepsilon = 0.2$$

- So Player 2 randomizes between g_2 and f_2

Player 1's Move at State 1

- Compare payoffs from f_1 and g_1
- Payoff from f_1: 0
- Payoff from g_1:

$$0.95 \cdot \frac{4}{19} \cdot 0.2 \cdot 3 + 0.95 \cdot \frac{4}{19} \cdot 0.8 \cdot 4 + 0.95 \cdot \frac{15}{19} \cdot (-1) + 0.05 \cdot 0.2 \cdot 8 + 0.05 \cdot 0.8 \cdot 4 = 0.25$$

- Since $0.25 > 0$, Player 1 prefers $g_1 \Rightarrow \beta = 1$

The Sequential Equilibrium – Summary

- Collecting the results of the support analysis, the scenario σ is:

Variable	Meaning	Value
β	Prob. 1 plays g_1 at state 1	1
γ	Prob. 2 plays g_2 if selfish	$\frac{4}{19}$
ε	Prob. 1 plays g_3 at state 3	0.2
ζ	Prob. 2 plays g_4 if selfish	0
δ	Belief 1 holds at state 3 (2 selfish)	0.8

- With the prior $\alpha = 0.95$, the pair (σ, π) is **fully consistent** and **sequentially rational** at every information state.
- This is the unique sequential equilibrium: Player 1 enters generously, Player 2 mixes when selfish, and cooperation partially unravels near the end.

Subgame-Perfect Equilibria

Subgame-Perfect Equilibrium

- The concept of **subgame-perfect equilibrium (SPE)** was introduced by Selten (1965, 1975, 1978).
- SPE is older and weaker than the concept of **sequential equilibrium**.
- However, SPE is very useful, especially in games with **infinite pure-strategy sets**.
- In some settings, sequential equilibrium is not yet formally defined.

Subgames

- A **subgame** of Γ^e is defined as:
  - Start from a subroot x
  - Delete all nodes and branches **not** following x
  - Make x the **root** of the subgame
- Let Γ^e_x be the subgame starting at subroot x.
- If x occurs during play, players know they are in the subgame beginning at x.
- Thus, the subgame structure becomes **common knowledge**.

Interpretation of Subgames

- A game theorist observing x can describe the subgame Γ_x^e.
- They can analyze player behavior in this subgame alone.
- Rational behavior **after** x must also be rational **within** the subgame.

- Selten defined an **SPE** of Γ^e as:

Definition: Subgame-Perfect Equilibrium (Selten)

An equilibrium in behavioral strategies such that, in **every subgame** of Γ^e, the restriction of the equilibrium strategy is itself an equilibrium.

Relationship to Sequential Equilibrium

- Let (σ, π) be a sequential equilibrium of Γ^e.
- Then σ is a **subgame-perfect equilibrium**.
- However, the reverse is not always true:

> **Note**
>
> The set of subgame-perfect equilibria can be **much larger** than the set of sequential equilibria.

When Subgame Perfection Has No Bite

- A subgame can start at x only if the singleton $\{x\}$ is an information set – i.e. reaching x is **common knowledge**.
- In games of **imperfect information**, this often fails: the only subgame is the whole game.

- Example: in our generosity game, Player 1 never observes Nature's move, so no node below the root starts a proper subgame.
- There, **subgame perfection reduces to Nash** and rules nothing out.

> **Note**
>
> This is exactly why we need sequential equilibrium: it imposes rationality and consistent beliefs at *every* information state, not only at the rare nodes that begin a subgame.

Example

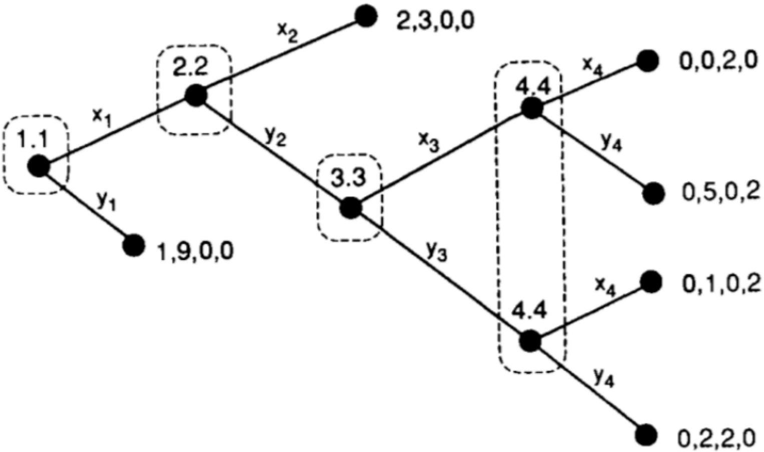

Figure 4: Game 4

Games with Perfect Information

- A game has **perfect information** if each information set of each player contains **exactly one decision node**.
- That is, for every player i and every information state $s \in S_i$, the set Y_s satisfies:

$$|Y_s| = 1$$

- Such games represent settings where:
  - Players move **one at a time**
  - All previous actions are **fully observed** before a move

Beliefs in Perfect Information

- In a game with perfect information:
    - If $x \in Y_s$, then $Y_s = \{x\}$
    - Hence, $\pi_s(x) = 1$
- The **only belief vector** consistent with perfect information is:

$$\pi = \textbf{vector of all ones}$$

Subgame-Perfect = Sequential Equilibrium

- Every node is a **subroot** in perfect information games
- Therefore:

Coincidence under Perfect Information

In games with perfect information, **subgame-perfect equilibria** and **sequential equilibria** coincide.

Zermelo's Theorem

Theorem 7 (Zermelo)

If Γ^e is an extensive-form game with perfect information, then there exists at least one sequential equilibrium of Γ^e in pure strategies (so every move probability is either 0 or 1).

Forward Induction

Forward Induction

- **Backward induction** supports sequential rationality and subgame-perfect equilibrium.
- But it assumes players only reason backward from the end.

- **Forward induction** instead reasons from **past moves**:
  - Players may draw inferences from earlier moves about **future intentions**.

Example

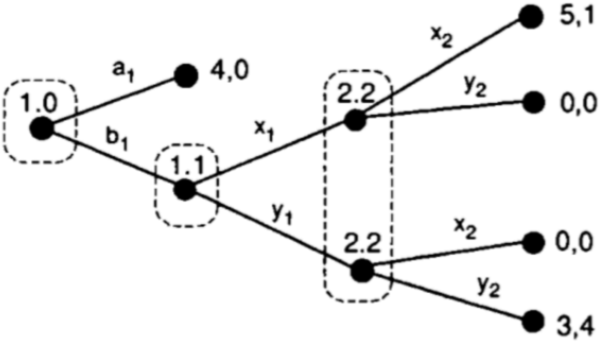

Figure 5: Game 5

Example (Cont.)

- The game has **three sequential equilibria**:
  - $([\alpha_1], [y_1], [y_2])$
  - $([\alpha_1], 0.8[x_1] + 0.2[y_1], \frac{3}{8}[x_2] + \frac{5}{8}[y_2])$
  - $([\beta_1], [x_1], [x_2])$
- Forward induction helps eliminate some of these:
  - Only $([\beta_1], [x_1], [x_2])$ gives player 1 a payoff **not worse than 4**, which they could get by playing α_1.

Reasoning Behind Forward Induction

- Suppose player 1 chooses β_1 instead of α_1.
- This causes the game to continue into the subgame at node 1.1.
- Player 2 should reason:
  - "Player 1 chose β_1 expecting that I would play x_2 in the subgame."
  - So, her **best response** is to play x_2.
- Thus, player 1 **expects** this outcome and chooses β_1 at the root.

Role of Representation

- The key: the person who chooses between x_1 and y_1 is **the same** person who chose β_1.
- This connection allows forward induction reasoning.

- To formalize this:
  - Use the **normal form** (not multi-agent form)
  - This preserves **cross-stage reasoning**

The Centipede Game

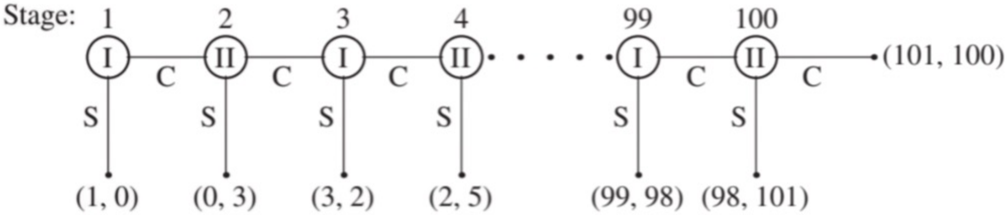

Figure 6: Game 6

Summary

The Refinement Hierarchy (1/2)

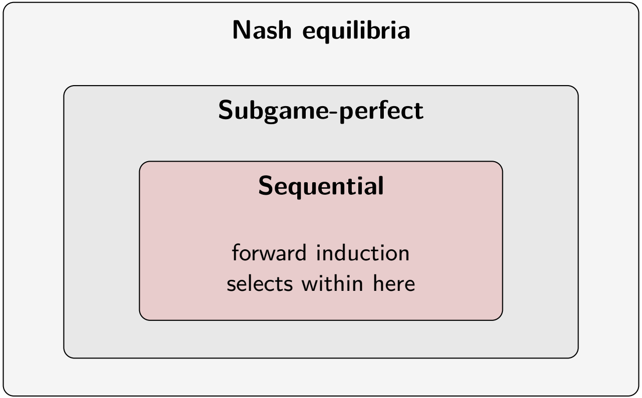

The Refinement Hierarchy (2/2)

- Each concept refines the one outside it: every sequential equilibrium is subgame-perfect, and every subgame-perfect equilibrium is Nash.
- Under **perfect recall**, behavioral and mixed strategies are equivalent (Kuhn), so we may solve via the normal form and read off behavior.
- Under **perfect information**, the three inner notions coincide and a pure-strategy equilibrium exists (Zermelo).
- **Forward induction** is not a larger set – it is a logic for selecting *among* sequential equilibria.

Required Reading

- **Myerson (1997)**
  - Chapter 4 – 4.1–4.9

Sequential Equilibrium: A Reminder

Sequential Equilibrium: Reminder

Sequential Equilibrium (Kreps and Wilson, 1982)

A *sequential equilibrium* (or a *full sequential equilibrium*) of Γ^e is any pair (σ, π) in $(\times_{s\in S^*}\Delta(D_s)) \times (\times_{s\in S^*}\Delta(Y_s))$ such that:

- The belief vector π is **fully** (strong consistency) consistent with σ, and
- With beliefs π, the scenario σ is sequentially rational for every player at every information state.

A behavioral-strategy profile σ is a *sequential-equilibrium scenario* if there exists some π such that (σ, π) is a sequential equilibrium. That is, σ can be *extended* to a sequential equilibrium.

Sequential Equilibrium: An Example

- Let's consider the following game:

Game 0

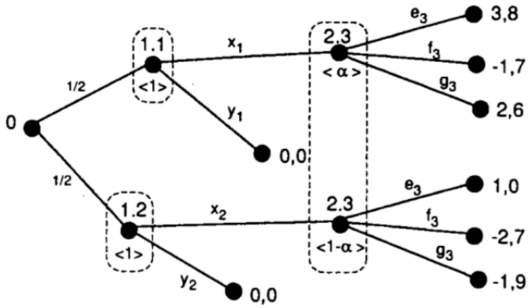

Game Setup and Beliefs

- Consider the **Game 0**.
- Let α be the belief probability that player 2 assigns to the top node in information state 3 (after x_1).
- At information state 3:
  - If player 2 chooses e_3: payoff is $8\alpha + 0(1-\alpha) = 8\alpha$
  - If player 2 chooses f_3: payoff is $7\alpha + 7(1-\alpha) = 7$
  - If player 2 chooses g_3: payoff is $6\alpha + 9(1-\alpha) = 9 - 3\alpha$

Optimal Moves at Information State 3

- e_3 is optimal when both:

$$8\alpha \geq 7 \quad \text{and} \quad 8\alpha \geq 9 - 3\alpha$$

  - This implies $\alpha \geq \frac{7}{8}$

- f_3 is optimal when both:

$$7 \geq 8\alpha \quad \text{and} \quad 7 \geq 9 - 3\alpha$$

  - So $\frac{2}{3} \leq \alpha \leq \frac{7}{8}$

- g_3 is optimal when:

$$9 - 3\alpha \geq 8\alpha \quad \text{and} \quad 9 - 3\alpha \geq 7$$

  - So $\alpha \leq \frac{2}{3}$

Candidate Supports for SE

- No value of α satisfies optimality of both e_3 and g_3.
- Possible supports for player 2 at information state 3:
  - $\{e_3\}$, $\{f_3\}$, $\{g_3\}$, $\{e_3, f_3\}$, $\{f_3, g_3\}$

Case 1: Support $\{e_3\}$

- Player 2 plays e_3, then for sure $\Rightarrow$ Player 1 plays x_1 and x_2
- $\Rightarrow \alpha = \frac{1}{2}$ (equal probability for both 2.3 nodes)
- But e_3 not optimal at $\alpha = \frac{1}{2}$ since g_3 gives higher payoff
- **No SE with support $\{e_3\}$**

Case 2: Support $\{f_3\}$

- Player 2 plays f_3, then for sure $\Rightarrow$ Player 1 plays y_1 and y_2
- α not pinned down, since 2.3 nodes have prior probability 0
- Sequential rationality requires $\frac{2}{3} \leq \alpha \leq \frac{7}{8}$
- **So** $([y_1], [y_2], [f_3])$ **is SE for** $\frac{2}{3} \leq \alpha \leq \frac{7}{8}$

Case 3: Support $\{g_3\}$

- Player 2 plays g_3, then for sure $\Rightarrow$ Player 1 plays x_1 and y_2
- Then $\alpha = 1$
- g_3 not optimal when $\alpha = 1$ (since e_3 better)
- **No SE with support** $\{g_3\}$

Case 4: Support $\{e_3, f_3\}$

- Randomization between e_3 and f_3 requires:

$$8\alpha = 7 = 9 - 3\alpha \Rightarrow \alpha = \frac{7}{8}$$

- **Option 1**:
  - Player 1 plays y_1, y_2 (so α not pinned down)
  - To ensure x_1 not better than y_1, e_3 must be played with probability $\beta \leq \frac{1}{4}$
  - **SE:** $([y_1], [y_2], \beta[e_3] + (1-\beta)[f_3])$ **with** $\alpha = \frac{7}{8}$ **and** $\beta \leq \frac{1}{4}$
- **Option 2**:
  - Player 1 plays x_1 and randomizes at node 1.2:

$$[x_1], \left(\frac{1}{7}\right)[x_2] + \left(\frac{6}{7}\right)[y_2]$$

  - Then 2.3 nodes get weights $\alpha = \frac{7}{8}$ (as required)
  - **SE:** $([x_1], (1/7)[x_2] + (6/7)[y_2], (2/3)[e_3] + (1/3)[f_3])$

Case 5: Support $\{f_3, g_3\}$

- For f_3 and g_3 to be optimal:

$$7 = 6\alpha + 9(1-\alpha) \Rightarrow \alpha = \frac{2}{3}$$

- Then player 1 must play y_1 at state 1 so that 2.3 nodes have prior probability 0
- Let player 2 play $(1-\gamma)[f_3] + \gamma[g_3]$ with $\gamma \leq \frac{1}{3}$
- **SE:** $([y_1], [y_2], (1-\gamma)[f_3] + \gamma[g_3])$ **with** $\alpha = \frac{2}{3}$ **and** $\gamma \leq \frac{1}{3}$