PROFESSOR:

# MICROECONOMICS 602 UNIVERSITY OF MICHIGAN

# *Contents*

# *Class 1: Introduction to Game Theory*

## *Defining Feature: Strategic Interaction*

The defining feature of game theory is *strategic interaction*. An environment is strategic whenever an agent's optimal decision depends explicitly on the decisions taken by others, and, symmetrically, when one's own decision affects the incentives faced by other agents.

In such settings, decision-making cannot be understood in isolation. Each player must form expectations about the behavior of others, because payoffs depend on the entire profile of actions, not solely on individual choices. As a result, rational behavior involves reasoning about how others reason.

**Remark** (Relation to earlier models in the core sequence)**.** It is useful to contrast this perspective with earlier models studied in the core:

- In **Micro 601**, uncertainty was present, but it was not strategic: uncertainty came from nature, not from the choices of other agents.

- In **Micro 603**, agents are price takers; individual actions do not influence market prices or the behavior of others.

- In **Micro 602**, by contrast, interaction itself becomes central: agents internalize that their actions affect others, and vice versa.

Game theory provides the formal language to study precisely this type of interaction.

## *Dimensions in Which Games Differ*

Games can be classified along several dimensions. At this stage of the course, two distinctions are particularly important, as they determine both how players reason and which solution concepts are appropriate.

**Two basic dimensions**

1) **Timing of moves.**

The first dimension concerns whether players move simultaneously or sequentially.

*Simultaneous-move games:*
Players choose their actions without observing the actions chosen by others. Each player must therefore form beliefs about what others are doing at the same time.

*Sequential-move games:*
Some players observe the actions of others before choosing their own. This allows later movers to condition their choices on earlier actions, introducing an explicit temporal structure into strategic reasoning.

2) **Information.**

The second dimension concerns what players know about the environment in which they interact.

*Complete-information games:*
All relevant aspects of the game—such as players, available actions, and payoff functions—are commonly known among the players.

*Incomplete-information games:*
Some aspects of the game are not known to at least one player. This typically involves uncertainty about preferences, payoffs, or types of other players.

**Focus of the course:** We begin by studying games with complete information, starting with simultaneous-move games.

## *Simultaneous-Move Games of Complete Information*

A simultaneous-move game of complete information can be modeled as a *strategic-form* (a.k.a. normal-form) game.

**Definition** (Strategic-form game)**.** A *strategic-form game* is a tuple

$$G = \big((S_i)_{i=1}^{I},\ (u_i)_{i=1}^{I}\big)$$

such that:

(i) $\{1, \ldots, I\}$ is a finite set of players.

(ii) For each player $i$, $S_i$ is a set of strategies for player $i$.

(iii) For each player $i$, $u_i : \prod_{j=1}^{I} S_j \to \mathbb{R}$ is a payoff function for player $i$.

**Remark** (Interpretation)**.** In simultaneous-move games, it is often useful to think of a *strategy* as simply an *action*. Payoffs $u_i$ can be interpreted as, e.g., a utility function (if $i$ is an individual) or a profit function (if $i$ is a firm).

**Remark.** Each $u_i$ is "integrable".

## *Notation for Strategy Profiles*

Let

$$S = \prod_{i=1}^{I} S_i$$

denote the set of *strategy profiles*. A typical element is

$$s = (s_1, \dots, s_I) \in S.$$

For any player $i$, define the product set of opponents' strategies

$$S_{-i} = \prod_{\substack{j=1 \\ j \neq i}}^{I} S_j,$$

and write a profile as $s = (s_i, s_{-i})$ where

$$s_{-i} = (s_1, \dots, s_{i-1}, s_{i+1}, \dots, s_I) \in S_{-i}.$$

**Example** (Binary strategies with three players)**.** Consider a game with three players, indexed by $I = \{1, 2, 3\}$. Each player has two possible actions. For concreteness, interpret

$$\ell_i \quad \text{as "Left"} \qquad \text{and} \qquad r_i \quad \text{as "Right"}$$

for player $i$. Thus, for each player $i$,

$$S_i = \{\ell_i, r_i\}.$$

The set of all possible *strategy profiles* is given by the Cartesian product

$$S = \prod_{i=1}^{3} S_i = \{\ell_1, r_1\} \times \{\ell_2, r_2\} \times \{\ell_3, r_3\}.$$

An element of $S$ specifies one action for each player. For example,

$$s = (\ell_1, r_2, r_3)$$

describes the situation in which player 1 chooses Left, while players 2 and 3 choose Right.

It is often useful to separate the action of a given player from the actions of the others. For player 2, the set of strategies of the *other players* is

$$S_{-2} = \{\ell_1, r_1\} \times \{\ell_3, r_3\}.$$

In the profile $s = (\ell_1, r_2, r_3)$, player 2's own action is $s_2 = r_2$, while the actions of the other players are

$$s_{-2} = (\ell_1, r_3).$$

Hence, the full strategy profile can be written as the pair $(s_2, s_{-2})$, which simply reconstructs $s$ as $s = (s_2, s_{-2})$.

## *An Illustrative Example of a Strategic-Form Game*

To fix ideas, consider a simple strategic-form game with two players, labeled 1 and 2. Player 1 chooses an action from the set

$$S_1 = \{U, D\},$$

while player 2 chooses an action from

$$S_2 = \{L, R\}.$$

The game is represented in matrix form: player 1 chooses rows and player 2 chooses columns. Each cell of the matrix lists a pair of numbers, where the first component is the payoff of player 1 and the second component is the payoff of player 2.

**Remark** (Reading the payoff matrix)**.** A payoff matrix is simply a compact way of representing the payoff functions

$$u_i : S_1 \times S_2 \to \mathbb{R}, \qquad i = 1, 2.$$

Fixing a row and a column selects an action for each player, and the corresponding cell reports the payoffs associated with that strategy profile.

We now consider two versions of this game, labeled *Example 1a* and *Example 1b*. The set of players and the available actions are the same in both cases. Moreover, player 2's payoff function is identical across the two examples. What changes is the payoff function of player 1.

*Example 1a*

| | $L$ | $R$ |
|---|---|---|
| $U$ | $(-1, 1)$ | $(0, 0)$ |
| $D$ | $(2, 0)$ | $(1, 3)$ |

*Example 1b*

| | $L$ | $R$ |
|---|---|---|
| $U$ | $(3, 1)$ | $(0, 0)$ |
| $D$ | $(2, 0)$ | $(1, 3)$ |

These two payoff matrices illustrate a central idea of game theory: the strategic nature of a problem is entirely encoded in the payoff functions. Even when players, actions, and timing are held fixed, changes in payoffs can fundamentally alter how players reason about their choices.

At this stage, the purpose of the example is not to analyze optimal behavior or to predict outcomes. Rather, it serves to familiarize us with the formal description of a game in strategic form. In the next class, we will use these same examples to study strategic reasoning and equilibrium behavior.

# *Class 2: Strategic Reasoning, Best Responses, and Nash Equilibrium*

## *Strategic Reasoning in Static Games*

We return to the two games introduced at the end of Class 1. The purpose now is no longer merely descriptive. Instead, we ask how rational players reason about their choices when payoffs depend on the actions of others.

The key distinction that emerges is between decision-theoretic problems—where optimal choices do not depend on beliefs about others—and genuinely game-theoretic problems, where expectations about opponents' behavior are central.

## *Example 1a: Asymmetric Strategic Structure*

Consider first Example 1a. In this game, player 1's incentives are independent of player 2's action. Regardless of whether player 2 chooses $L$ or $R$, player 1 strictly prefers action $D$ over $U$.

As a consequence, player 1's problem is essentially decision-theoretic: there is a uniquely optimal action that does not require forming beliefs about the opponent. Player 1 will choose $D$ no matter what player 2 does.

Player 2's problem, by contrast, is genuinely strategic. Player 2's optimal action depends on what she expects player 1 to do. If player 2 expects player 1 to choose $U$, then $L$ is optimal. If she expects player 1 to choose $D$, then $R$ is optimal.

Since player 1 has a dominant action, player 2 should expect player 1 to choose $D$ and therefore respond by choosing $R$. This line of reasoning—anticipating others' optimal behavior and responding to it—is an example of *iterative strategic reasoning*.

## *Example 1b: Mutual Strategic Dependence*

In Example 1b, player 2's incentives remain unchanged, but player 1's payoff function is modified so that her optimal action now depends on player 2's behavior.

If player 1 expects player 2 to choose $L$, then $U$ yields a higher payoff than $D$. If instead player 1 expects player 2 to choose $R$, then $D$ is optimal. Thus, player 1's problem is now also strategic.

In this case, iterative reasoning alone does not eliminate any action. Different combinations of expectations may be self-consistent, and outcomes such as $(U, L)$ are no longer ruled out by simple dominance arguments.

## *Expectations and Consistency*

These examples highlight a fundamental issue in strategic interaction. Players choose actions based on expectations about what others will do. However, those expectations must themselves be consistent with the actions that are actually chosen.

A mismatch of expectations may arise when, for instance, player 1 chooses $U$ because she expects player 2 to play $L$, while player 2 chooses $R$ because she expects player 1 to play $D$. In such a situation, both players are best-responding to beliefs that are not mutually consistent.

This observation motivates the equilibrium concept introduced next.

## *Nash Equilibrium*

**Definition** (Nash Equilibrium)**.** A strategy profile $s^* = (s_1^*, \ldots, s_I^*)$ is a *Nash equilibrium* if, for every player $i = 1, \ldots, I$,

$$u_i(s_i^*, s_{-i}^*) \;\geq\; u_i(s_i, s_{-i}^*) \quad \text{for all } s_i \in S_i.$$

At a Nash equilibrium, each player's strategy is a best response to the strategies chosen by the other players. No player has an incentive to deviate unilaterally.

**Remark** (Interpretation)**.** A Nash equilibrium corresponds to a situation in which players' expectations about others' behavior are correct. Each player correctly anticipates what the others are doing and responds optimally.

## *Best Responses*

The definition of Nash equilibrium can be reformulated using the notion of best responses.

**Definition** (Best Response)**.** Fix $s_{-i} \in S_{-i}$. An action $s_i \in S_i$ is a *best response* to $s_{-i}$ if

$$u_i(s_i, s_{-i}) \ \geq \ u_i(s_i', s_{-i}) \quad \text{for all } s_i' \in S_i.$$

**Definition** (Best Response Correspondence)**.** The best response correspondence of player $i$ is the mapping

$$BR_i : S_{-i} \rightrightarrows S_i$$

defined by

$$BR_i(s_{-i}) = \{\, s_i \in S_i : s_i \text{ is a best response to } s_{-i} \,\}.$$

A Nash equilibrium is precisely a strategy profile $s^*$ such that, for every player $i$, $s_i^* \in BR_i(s_{-i}^*)$.

**Remark.** A strategy profile $s^* = (s_1^*, \ldots, s_I^*)$ is a Nash equilibrium if and only if

$$s_i^* \in BR_i(s_{-i}^*) \quad \text{for every } i = 1, \ldots, I.$$

Thus, a Nash equilibrium is a fixed point of the best response correspondences.

We will explore more of NE as a fix-point later in the course.

## *Best Responses in the Examples*

In Example 1a, player 1's best response correspondence is constant:

$$BR_1(s_2) = \{D\} \quad \text{for all } s_2 \in \{L, R\}.$$

Player 2's best responses satisfy

$$BR_2(U) = \{L\}, \qquad BR_2(D) = \{R\}.$$

The unique Nash equilibrium is therefore $(D, R)$.

In Example 1b, player 1's best responses depend on player 2's action:

$$BR_1(L) = \{U\}, \qquad BR_1(R) = \{D\}.$$

Player 2's best responses remain unchanged. As a result, the game admits two Nash equilibria: $(U, L)$ and $(D, R)$.

These examples illustrate how equilibrium outcomes emerge from the interaction of players' best responses.

## *Example 2: Quantity Competition with Linear Demand*

We now study a strategic-form game with a continuum of strategies, which illustrates how the concept of Nash equilibrium extends beyond finite action sets.

*Environment*

There are two firms, indexed by $i = 1, 2$, producing a homogeneous good. Each firm chooses a quantity, $s_i \in \mathbb{R}_+$.

Production costs are zero. Market price is determined by a linear inverse demand function, subject to the constraint that price cannot be negative. Specifically, the price is given by

$$P(Q) = \max\{\, p - \phi Q,\, 0\,\},$$

where $Q = s_i + s_{-i} = s_1 + s_2$, with parameters $p > 0$ and $\phi > 0$.

Firm $i$'s payoff (profit) function is therefore

$$u_i(s_i, s_{-i}) = s_i \cdot \max\{\, p - \phi(s_i + s_{-i}),\, 0\,\}.$$

*Step 1: Deriving the Best Response*

Fix $s_{-i} \geq 0$. We begin by ignoring the non-negativity constraint on price and consider the unconstrained problem

$$\max_{s_i \geq 0}\ s_i\big(p - \phi(s_i + s_{-i})\big).$$

The first-order condition for an interior solution is

$$p - 2\phi s_i - \phi s_{-i} = 0,$$

which yields

$$s_i = \frac{p}{2\phi} - \frac{s_{-i}}{2}.$$

We must now verify when this candidate solution is compatible with a non-negative price. Substituting back,

$$p - \phi(s_i + s_{-i}) \geq 0 \quad \Longleftrightarrow \quad \frac{p}{\phi} > s_{-i}.$$

*Best Response Correspondence*

Combining these observations, the best response correspondence of firm $i$ is given by

$$BR_i(s_{-i}) = \begin{cases} \left\{ \dfrac{p}{2\phi} - \dfrac{s_{-i}}{2} \right\} & \text{if } s_{-i} < \dfrac{p}{\phi}, \\ \mathbb{R}_+ & \text{if } s_{-i} \geq \dfrac{p}{\phi}. \end{cases}$$

The second case reflects the fact that when the rival produces so much that price is driven to zero, firm $i$ is indifferent across all non-negative quantities.

*Step 2: Solving for Nash Equilibrium*

A pair $(s_1^*, s_2^*)$ is a Nash equilibrium if

$$s_1^* \in BR_1(s_2^*) \quad \text{and} \quad s_2^* \in BR_2(s_1^*).$$

We first look for an equilibrium with strictly positive prices, that is,

$$s_1^* < \frac{p}{\phi} \quad \text{and} \quad s_2^* < \frac{p}{\phi}.$$

In this case, both firms' best responses satisfy

$$s_1^* = \frac{p}{2\phi} - \frac{s_2^*}{2}, \qquad s_2^* = \frac{p}{2\phi} - \frac{s_1^*}{2}.$$

Solving this system yields

$$s_1^* = s_2^* = \frac{p}{3\phi}.$$

This pair constitutes a Nash equilibrium with positive prices.

**Remark.** This example illustrates how Nash equilibrium can be computed as the intersection of best response correspondences, even in games with continuous strategy spaces.

*Case Analysis for Nash Equilibria*

So far, we have identified a Nash equilibrium with strictly positive prices. We now examine whether additional equilibria may arise once the price constraint becomes binding.

*Case 2: One Firm Produces Too Much.* We ask whether there exists a Nash equilibrium $(s_1^*, s_2^*)$ such that

$$s_i^* \ \geq \ \frac{p}{\phi} \quad \text{and} \quad s_{-i}^* \ < \ \frac{p}{\phi}$$

for some firm $i$.

The answer is no. If $s_i^* \geq \frac{p}{\phi}$, then total output satisfies

$$s_i^* + s_{-i}^* \ \geq \ \frac{p}{\phi},$$

which implies that the market price is zero. Hence,

$$u_i(s_i^*, s_{-i}^*) = 0.$$

However, consider a deviation $\hat{s}_i > 0$ small enough so that

$$p - \phi(\hat{s}_i + s_{-i}^*) > 0.$$

Such a deviation exists because $s_{-i}^* < \frac{p}{\phi}$. For this deviation,

$$u_i(\hat{s}_i, s_{-i}^*) = \hat{s}_i\big(p - \phi(\hat{s}_i + s_{-i}^*)\big) > 0,$$

which strictly exceeds the payoff under $(s_i^*, s_{-i}^*)$. Therefore, $s_i^*$ cannot be a best response, and no Nash equilibrium of this form exists.

*Case 3: Both Firms Produce Too Much.* We now ask whether there exists a Nash equilibrium $(s_1^*, s_2^*)$ such that

$$s_1^* \geq \frac{p}{\phi} \quad \text{and} \quad s_2^* \geq \frac{p}{\phi}.$$

In this case, total output satisfies $s_1^* + s_2^* \geq \frac{p}{\phi}$, so the market price is zero. Each firm therefore earns zero profit regardless of its own output level.

Since profits are identically zero for all $s_i \geq 0$, every non-negative quantity is a best response. Consequently, for any pair

$$(s_1^*, s_2^*) \in \left[\frac{p}{\phi}, \infty\right) \times \left[\frac{p}{\phi}, \infty\right),$$

both firms are best-responding to each other.

**Remark** (Intuition)**.** When the rival produces at least $\frac{p}{\phi}$, total output is so large that the market price is zero regardless of one's own production level. Even if a firm deviates by producing less, it can never reduce total output enough to make the price positive. As a result, no unilateral deviation can yield strictly positive profits, and every quantity is a best response.

**Remark.** This example illustrates that games with payoff discontinuities or binding constraints may admit a continuum of Nash equilibria, many of which are economically uninteresting. In applications, attention is often restricted to equilibria with positive prices and positive profits.

# *Class 3: Mixed Strategies*

## *Game Theory versus Decision Theory*

A useful way to motivate the introduction of mixed strategies is to contrast game theory with standard decision theory.

- In decision theory, increasing your payoffs cannot make you worse off.
- In decision theory, giving you additional options cannot make you worse off.
- Both principles may fail in game theory, because changing payoffs or available actions can alter the incentives and behavior of other players.

The key distinction is that, in games, outcomes depend not only on one's own choice but also on how others react to that choice.

## *Example 3: Sensitivity to Payoffs*

### *Example 3a*

Consider the following two-player game:

| | $L$ | $R$ |
|---|---|---|
| $U$ | $(2,2)$ | $(1,3)$ |
| $D$ | $(3,1)$ | $(0,0)$ |

This game has two pure-strategy Nash equilibria:

$$(D,L) \quad \text{and} \quad (U,R).$$

Suppose that the realized outcome happens to be $(U,R)$.

### *Example 3b*

Now modify the game by slightly perturbing player 2's payoffs:

| | $L$ | $R$ |
|---|---|---|
| $U$ | $(2, 2 + x_{UL})$ | $(1, 3 + x_{UR})$ |
| $D$ | $(3, 1 + x_{DL})$ | $(0, x_{DR})$ |

where all perturbations satisfy $x_{UL}, x_{UR}, x_{DL}, x_{DR} \geq 0$ and

$$x_{UL} - x_{UR} > 1, \qquad 3 > 1 + x_{DL} \geq x_{DR}.$$

Under these assumptions:

- $L$ is a best response to $D$.
- $L$ is the unique best response to $U$.

Hence, the game now has a *unique* Nash equilibrium, namely

$$(D, L).$$

This example illustrates how small changes in payoffs can eliminate equilibria and fundamentally alter strategic behavior.

## *Example 4: No Pure-Strategy Nash Equilibrium*

Consider the following game:

| | $H$ | $T$ |
|---|---|---|
| $H$ | $(1, -1)$ | $(-1, 1)$ |
| $T$ | $(-1, 1)$ | $(1, -1)$ |

The best responses are:

$$BR_1(H) = \{H\}, \quad BR_1(T) = \{T\},$$

$$BR_2(H) = \{T\}, \quad BR_2(T) = \{H\}.$$

Player 1 wants to *match* player 2's action, while player 2 wants to *mismatch*. As a consequence, there is **no pure-strategy Nash equilibrium**.

- Player 2 does not want player 1 to know what she is choosing.
- Player 2 has an incentive to randomize.
- More generally, players may want to randomize in strategic interactions.[1]

[1] The role of randomization is not to increase expected payoffs directly, but to prevent opponents from exploiting predictable behavior. In games without pure-strategy equilibria, mixing may be the only way to remain optimal.

## *Mixed Strategies*

Let $S_i$ denote the set of pure strategies of player $i$.

- Let $\Delta(S_i)$ denote the set of probability distributions on $S_i$; these are player $i$'s *mixed strategies*.

- We write $\sigma_i \in \Delta(S_i)$ for a particular mixed strategy of player $i$.
- $\sigma_i(s_i)$ is the probability that pure strategy $s_i$ is played under $\sigma_i$.
- If $E_i \subseteq S_i$, then $\sigma_i(E_i)$ is the probability that a strategy in $E_i$ is played.

**Remark.** If $S_i$ is finite, then

$$\sigma_i(E_i) = \sum_{s_i \in E_i} \sigma_i(s_i).$$

For infinite $S_i$, this need not be true. For example, if $S_i = [0,1]$ and $\sigma_i$ is the uniform distribution, then $\sigma_i(s_i) = 0$ for all $s_i \in [0,1]$, while $\sigma_i([0,1]) = 1$.

### *Notation*

A profile of mixed strategies is denoted by

$$\sigma = (\sigma_1, \ldots, \sigma_I) \in \prod_{i=1}^{I} \Delta(S_i).$$

We write $\sigma_{-i}$ for the mixed-strategy profile of all players other than $i$.

## *Expected Payoffs and the Mixed Extension*

We extend the payoff functions from pure strategies to mixed strategies by defining

$$u_i(\sigma_1, \ldots, \sigma_I) = \int_{S_I} \cdots \int_{S_1} u_i(s_1, \ldots, s_I)\, d\sigma_1 \cdots d\sigma_I.$$

Technically, this defines a new payoff function

$$\tilde{u}_i : \prod_{i=1}^{I} \Delta(S_i) \to \mathbb{R},$$

but we will continue to write $u_i$ for simplicity.

**Remark.** The game with strategy sets $\Delta(S_i)$ and payoff functions $u_i$ is called the *mixed extension* of the original game.

**Remark.** If the game is finite, then expected payoffs can be written as

$$u_i(\sigma_1, \ldots, \sigma_I) = \sum_{s_I \in S_I} \cdots \sum_{s_1 \in S_1} u_i(s_1, \ldots, s_I)\, \sigma_1(s_1) \cdots \sigma_I(s_I).$$

# Class 4: Mixed-Strategy Nash Equilibrium

## Mixed strategies and the equilibrium concept

Throughout, we work in the *mixed extension* of a strategic-form game. For each player $i \in \{1, \dots, I\}$, let $S_i$ denote the set of pure strategies and let

$$\Delta(S_i)$$

denote the set of probability distributions (more generally, probability measures) over $S_i$. A mixed strategy of player $i$ is denoted by $\sigma_i \in \Delta(S_i)$, and a mixed strategy profile is

$$\sigma = (\sigma_1, \dots, \sigma_I) \in \prod_{i=1}^{I} \Delta(S_i).$$

We write $\sigma_{-i} = (\sigma_1, \dots, \sigma_{i-1}, \sigma_{i+1}, \dots, \sigma_I) \in \prod_{j \neq i} \Delta(S_j)$.

**Remark** (Expected payoffs under independent randomization)**.** A mixed strategy $\sigma_i \in \Delta(S_i)$ is a probability distribution over player $i$'s pure strategies. Allowing randomization is strategically useful because it can make a player *hard to predict*: when an opponent cannot anticipate with certainty which pure action will be played, it becomes harder to exploit deterministic behavior.

Given a mixed-strategy profile $\sigma = (\sigma_1, \dots, \sigma_I)$, the payoff $u_i(\sigma)$ is defined as the *expected value* of $u_i$ when each player draws a pure strategy according to $\sigma_i$, independently across players.

In particular, when the game is finite, one may write

$$u_i(\sigma_1, \dots, \sigma_I) = \sum_{s_1 \in S_1} \cdots \sum_{s_I \in S_I} u_i(s_1, \dots, s_I)\, \sigma_1(s_1) \cdots \sigma_I(s_I).$$

**Remark** (Expected payoffs under independent randomization)**.** Fix a finite normal-form game. For each player $i$, let

$$S_i = \{s_i^1, \dots, s_i^{m_i}\}, \qquad m_i := |S_i|.$$

A mixed strategy $\sigma_i \in \Delta(S_i)$ can be identified with a vector

$$\sigma_i = \big(\sigma_i(s_i^1), \dots, \sigma_i(s_i^{m_i})\big) \in \mathbb{R}^{m_i} \quad \text{such that} \quad \sigma_i(s_i^k) \geq 0, \ \sum_{k=1}^{m_i} \sigma_i(s_i^k) = 1.$$

Thus $\Delta(S_i)$ is the probability simplex in $\mathbb{R}^{m_i}$.
**Independent randomization.** Under a mixed-strategy profile $\sigma = (\sigma_1, \ldots, \sigma_I)$, each player $j$ draws a pure strategy $s_j \in S_j$ according to $\sigma_j$, and these draws are independent across players. Therefore, for any pure profile $s = (s_1, \ldots, s_I) \in S_1 \times \cdots \times S_I$, the probability that $s$ is realized equals

$$\mathbb{P}_\sigma(s) = \sigma_1(s_1) \cdots \sigma_I(s_I).$$

**Expected payoff.** The payoff $u_i(\sigma)$ is defined as the expected value of $u_i(s)$ under this probability distribution:

$$u_i(\sigma) = \mathbb{E}_\sigma\big[u_i(s)\big] = \sum_{s_1 \in S_1} \cdots \sum_{s_I \in S_I} u_i(s_1, \ldots, s_I)\, \mathbb{P}_\sigma(s_1, \ldots, s_I).$$

Substituting $\mathbb{P}_\sigma(s_1, \ldots, s_I) = \sigma_1(s_1) \cdots \sigma_I(s_I)$ yields

$$u_i(\sigma_1, \ldots, \sigma_I) = \sum_{s_1 \in S_1} \cdots \sum_{s_I \in S_I} u_i(s_1, \ldots, s_I)\, \sigma_1(s_1) \cdots \sigma_I(s_I),$$

which is precisely a weighted average of the pure payoffs, with weights given by the probability that each pure profile occurs.

**Definition** (Mixed-strategy Nash equilibrium (MSNE))**.** A mixed strategy profile $\sigma^* = (\sigma_1^*, \ldots, \sigma_I^*)$ is a *mixed-strategy Nash equilibrium* if, for each $i = 1, \ldots, I$,

$$u_i(\sigma_i^*, \sigma_{-i}^*) \;\geq\; u_i(\sigma_i, \sigma_{-i}^*) \qquad \text{for all } \sigma_i \in \Delta(S_i). \tag{$*$}$$

**Definition** (Best response in mixed strategies)**.** Fix $\sigma_{-i} \in \prod_{j \neq i} \Delta(S_j)$. A mixed strategy $\sigma_i \in \Delta(S_i)$ is a *best response* to $\sigma_{-i}$ if

$$u_i(\sigma_i, \sigma_{-i}) \;\geq\; u_i(\sigma_i', \sigma_{-i}) \qquad \text{for all } \sigma_i' \in \Delta(S_i).$$

**Definition** (Best-response correspondence)**.** The (mixed) best-response correspondence of player $i$ is the set-valued map

$$\widetilde{\mathrm{BR}}_i : \prod_{j \neq i} \Delta(S_j) \rightrightarrows \Delta(S_i),$$

$$\widetilde{\mathrm{BR}}_i(\sigma_{-i}) = \big\{\sigma_i \in \Delta(S_i) : \sigma_i \text{ is a best response to } \sigma_{-i}\big\}.$$

**Remark** (Fixed point formulation)**.** A profile $\sigma^*$ is an MSNE if and only if, for each $i$,

$$\sigma_i^* \in \widetilde{\mathrm{BR}}_i(\sigma_{-i}^*).$$

That is, an MSNE is a fixed point of the best-response correspondence.

**Remark** (Why checking $(*)$ looks hard)**.** Condition $(*)$ quantifies over *all* mixed deviations $\sigma_i \in \Delta(S_i)$. Even if $S_i$ is finite, the set $\Delta(S_i)$ is infinite (a simplex), so in principle there is a continuum of deviations to check. The next result shows that one can restrict attention to deviations to *pure* strategies.

### *Reducing deviations: it suffices to check pure strategies*

**Lemma 0.0.1** (Checking only pure deviations)**.** *A strategy profile $\sigma^*$ is an MSNE if and only if, for each $i = 1, \dots, I$,*

$$u_i(\sigma_i^*, \sigma_{-i}^*) \;\geq\; u_i(s_i, \sigma_{-i}^*) \qquad \text{for all } s_i \in S_i. \tag{$**$}$$

*Proof.* $(*) \Rightarrow (**)$: Fix $s_i \in S_i$. Consider the degenerate mixed strategy $\sigma_i$ that assigns probability 1 to $s_i$. Plugging this $\sigma_i$ into $(*)$ yields $(**)$.

$(**) \Rightarrow (*)$: Fix any $\sigma_i \in \Delta(S_i)$. By linearity of expected utility in the player's own mixed strategy,

$$u_i(\sigma_i, \sigma_{-i}^*) = \int_{S_i} u_i(s_i, \sigma_{-i}^*)\, d\sigma_i(s_i) \leq \sup_{s_i \in S_i} u_i(s_i, \sigma_{-i}^*) \leq u_i(\sigma_i^*, \sigma_{-i}^*),$$

where the last inequality uses $(**)$. This establishes $(*)$. □

**Lemma 0.0.2** (Checking only pure deviations)**.** *A strategy profile $\sigma^*$ is an MSNE if and only if, for each $i = 1, \dots, I$,*

$$u_i(\sigma_i^*, \sigma_{-i}^*) \;\geq\; u_i(s_i, \sigma_{-i}^*) \qquad \text{for all } s_i \in S_i. \tag{$**$}$$

*Proof sketch (intuition).* Fix a player $i$ and hold $\sigma_{-i}^*$ fixed. Consider the map

$$\phi(\sigma_i) := u_i(\sigma_i, \sigma_{-i}^*), \qquad \sigma_i \in \Delta(S_i).$$

In a finite game, $u_i(\sigma_i, \sigma_{-i}^*)$ is *linear* in $\sigma_i$: indeed, it can be written as

$$\phi(\sigma_i) = \sum_{s_i \in S_i} \sigma_i(s_i)\, u_i(s_i, \sigma_{-i}^*).$$

Hence $\phi$ is a linear objective over the simplex $\Delta(S_i)$. A linear function attains its maximum over a simplex at an *extreme point* (a vertex), and the extreme points of $\Delta(S_i)$ are exactly the degenerate distributions that put probability 1 on a single pure strategy $s_i \in S_i$.

Therefore, if no pure strategy $s_i$ yields a higher payoff than $\sigma_i^*$ against $\sigma_{-i}^*$, then no mixed deviation can yield a higher payoff either. The converse implication is immediate since pure strategies are particular mixed strategies. □

### *Finite games: support and indifference*

**Lemma 0.0.3** (Support characterization in finite games)**.** *Fix a finite game. A mixed strategy profile $\sigma^*$ is an MSNE if and only if, for each $i$ and each $s_i \in S_i$ with $\sigma_i^*(s_i) > 0$,*

$$u_i(s_i, \sigma_{-i}^*) \;\geq\; u_i(r_i, \sigma_{-i}^*) \qquad \text{for all } r_i \in S_i.$$

**Corollary 0.0.4** (Equal payoffs on the support)**.** *Fix a finite game. A mixed strategy profile $\sigma^*$ is an MSNE if and only if, for each $i$ and each $s_i \in S_i$ with $\sigma_i^*(s_i) > 0$, the following two conditions hold:*

1. *If* $\sigma_i^*(r_i) > 0$*, then* $u_i(s_i, \sigma_{-i}^*) = u_i(r_i, \sigma_{-i}^*)$.
2. *If* $\sigma_i^*(r_i) = 0$*, then* $u_i(s_i, \sigma_{-i}^*) \geq u_i(r_i, \sigma_{-i}^*)$.

**Remark** (Interpretation)**.** If a player assigns strictly positive probability to multiple pure strategies, then she must be indifferent among them: otherwise she would shift probability mass toward the strictly better one. Equivalently, *mixing forces indifference for the opponent*: if player $j$ mixes over several actions, then player $i$ must choose a mixture that makes those actions yield the same expected payoff for $j$.

*Example 4 revisited: Matching pennies*

**Example** (Matching pennies)**.** Consider the two-player game with $S_1 = S_2 = \{H, T\}$ and payoff matrix

| | $H$ | $T$ |
|---|---|---|
| $H$ | $(1, -1)$ | $(-1, 1)$ |
| $T$ | $(-1, 1)$ | $(1, -1)$ |

Player 1 prefers to *match* (gets 1 on the diagonal), while player 2 prefers to *mismatch* (gets 1 off the diagonal). The game is strictly competitive: player 2's payoff is always the negative of player 1's.

**Goal.** We look for a mixed-strategy Nash equilibrium (MSNE), i.e., probabilities with which each player randomizes between $H$ and $T$ so that no unilateral change in these probabilities can increase expected payoff.

**Step 1 (no pure equilibrium).** There is no pure-strategy Nash equilibrium: at each pure profile, one of the players can profitably deviate by switching $H \leftrightarrow T$.

**Step 2 (why both players must mix with full support).** Suppose, toward a contradiction, that in an MSNE some player assigns probability 1 to one action, say player 1 plays $H$ for sure. Then player 2 would have a strict best reply (here: play $T$ for sure, to mismatch), so player 2 would not mix. But then player 1 would strictly prefer to switch (from $H$ to $T$), contradicting equilibrium. The same reasoning applies to any deterministic choice by either player. Hence, in any MSNE, both players must assign positive probability to both actions:

$$\sigma_1^*(H), \sigma_1^*(T) > 0 \quad \text{and} \quad \sigma_2^*(H), \sigma_2^*(T) > 0.$$

In other words, the equilibrium must have *full support*.

**Step 3 (mixing requires indifference of the opponent).** A key implication of mixing is *indifference*: if player 2 uses both $H$ and $T$ with positive probability, then player 2 must be indifferent between $H$ and $T$ given player 1's mixing probabilities; otherwise player 2 would place probability 1 on the strictly better action.

Let $p = \sigma_1^*(H)$, so $\sigma_1^*(T) = 1 - p$. Compute player 2's expected payoff from each pure action:

$$u_2(H, \sigma_1^*) = (-1) \cdot p + (1) \cdot (1 - p) = 1 - 2p,$$

$$u_2(T, \sigma_1^*) = (1) \cdot p + (-1) \cdot (1 - p) = 2p - 1.$$

Indifference $u_2(H, \sigma_1^*) = u_2(T, \sigma_1^*)$ implies $1 - 2p = 2p - 1$, hence $p = \frac{1}{2}$.

Symmetrically, let $q = \sigma_2^*(H)$ so $\sigma_2^*(T) = 1 - q$. Since player 1 also mixes with full support, player 1 must be indifferent between $H$ and $T$ given $\sigma_2^*$:

$$u_1(H, \sigma_2^*) = u_1(T, \sigma_2^*).$$

A parallel calculation yields $q = \frac{1}{2}$.

**Step 4 (verification and uniqueness).** When $p = q = \frac{1}{2}$, each player makes the opponent exactly indifferent between $H$ and $T$. Thus each player is willing to randomize, and no unilateral deviation can improve payoffs. Moreover, the indifference equations force $p = q = \frac{1}{2}$, so the MSNE is unique.

**Punchline.** The unique MSNE is

$$\sigma_1^*(H) = \sigma_1^*(T) = \frac{1}{2}, \qquad \sigma_2^*(H) = \sigma_2^*(T) = \frac{1}{2}.$$

**Remark** (Why a "biased" candidate cannot be an equilibrium)**.** Suppose one guesses a candidate where player 2 plays $H$ with probability $\frac{3}{4}$ and $T$ with probability $\frac{1}{4}$. Given this $\sigma_2$, player 1's expected payoff from $H$ and $T$ is

$$u_1(H, \sigma_2) = \tfrac{3}{4}(1) + \tfrac{1}{4}(-1) = \tfrac{1}{2}, \qquad u_1(T, \sigma_2) = \tfrac{3}{4}(-1) + \tfrac{1}{4}(1) = -\tfrac{1}{2}.$$

So player 1 has a *strict* best reply: play $H$ with probability 1. But if player 1 plays $H$ for sure, then player 2 has a strict best reply (play $T$ for sure, to mismatch). Hence the proposed "biased" mixing by player 2 cannot be part of an equilibrium: strict incentives force at least one player to become deterministic, and the resulting pure profile cannot be stable (as established in Step 1).

# *Class 5: Mixed Strategies with Infinite Strategy Sets*

This class clarifies why the familiar "support/indifference" logic from finite games must be reformulated when strategy sets are infinite. The key issue is that, under many continuous mixed strategies (e.g. a uniform distribution on $[0,1]$), every singleton $\{s_i\}$ has probability zero, so statements of the form "$\sigma_i(s_i) > 0$" lose their bite.

## *Remember: two lemmas from the finite setting*

**Lemma 0.0.5** (Checking only pure deviations)**.** *A mixed-strategy profile $\sigma^* = (\sigma_1^*, \dots, \sigma_I^*)$ is an MSNE if and only if, for each $i = 1, \dots, I$,*

$$u_i(\sigma_i^*, \sigma_{-i}^*) \;\geq\; u_i(s_i, \sigma_{-i}^*) \qquad \textit{for all } s_i \in S_i.$$

**Lemma 0.0.6** (Finite games: support/indifference test)**.** *Suppose each $S_i$ is finite. Then $\sigma^* = (\sigma_1^*, \dots, \sigma_I^*)$ is an MSNE if and only if, for each $i$, whenever $\sigma_i^*(s_i) > 0$,*

$$u_i(s_i, \sigma_{-i}^*) \;\geq\; u_i(r_i, \sigma_{-i}^*) \qquad \textit{for all } r_i \in S_i.$$

*In particular, if $\sigma_i^*(s_i) > 0$ and $\sigma_i^*(r_i) > 0$, then*

$$u_i(s_i, \sigma_{-i}^*) = u_i(r_i, \sigma_{-i}^*),$$

*i.e. player $i$ is indifferent across all pure strategies in the support of $\sigma_i^*$.*

**Remark** (Why Lemma 0.0.6 is a *finite*-game statement)**.** Lemma 0.0.6 is useful because, in finite games, "$\sigma_i^*(s_i) > 0$" identifies the actions that are actually played with positive probability. This is exactly what breaks down in many infinite settings.

## *Why the finite "support test" can fail in infinite games*

**Remark** (The uniform distribution example)**.** Lemma 0.0.6 does not have the right content for the infinite setting. For instance, let $S_i = [0,1]$

and suppose $\sigma_i$ is the uniform distribution on $[0,1]$. Then for every $s_i \in [0,1]$,

$$\sigma_i(\{s_i\}) = 0.$$

Hence the premise "$\sigma_i(s_i) > 0$" (as used in the finite-game lemma) never holds, so the lemma's requirement becomes *trivially satisfied*. But it is not true that such a $\sigma_i$ must be a best response in general. The correct replacement must talk about the probability assigned to *sets* of strategies, not to individual points.

## *The correct analogue for infinite games*

**Remark** (Notation in the infinite setting)**.** When $S_i$ is infinite, a mixed strategy $\sigma_i$ is naturally interpreted as a probability measure over $S_i$. Thus expressions like $\sigma_i(A)$, for a set $A \subseteq S_i$, denote the probability that player $i$'s random draw lies in $A$. In particular, it is common that $\sigma_i(\{s_i\}) = 0$ for every single point $s_i$ even though $\sigma_i$ is non-degenerate.

**Lemma 0.0.7** (Infinite-game support condition)**.** *A mixed-strategy profile $\sigma^* = (\sigma_1^*, \ldots, \sigma_I^*)$ is an MSNE if and only if, for each $i = 1, \ldots, I$, the following two conditions hold:*

1.
$$u_i(\sigma_i^*, \sigma_{-i}^*) \;\geq\; u_i(s_i, \sigma_{-i}^*) \qquad \textit{for all } s_i \in S_i.$$

2.
$$\sigma_i^*\Big(\{\, s_i \in S_i :\; u_i(\sigma_i^*, \sigma_{-i}^*) > u_i(s_i, \sigma_{-i}^*) \,\}\Big) = 0.$$

**Remark** (Interpretation of condition 2)**.** Condition 2 says that, under $\sigma_i^*$, the set of pure strategies that are *strictly worse* (for player $i$ given $\sigma_{-i}^*$) than the equilibrium payoff level has probability zero. **Equivalently**: player $i$ assigns probability 1 to strategies that are not strictly worse than $\sigma_i^*$ against $\sigma_{-i}^*$.

**Remark** (Logical relation between 1 and 2)**.** Condition 1 implies condition 2, but not conversely. Indeed, 2 only rules out putting positive probability on *strictly worse* strategies, yet it still allows the possibility that there exists some $s_i$ with

$$u_i(s_i, \sigma_{-i}^*) > u_i(\sigma_i^*, \sigma_{-i}^*),$$

i.e. an even better deviation than $\sigma_i^*$.

## *Representing mixed strategies by CDFs*

**Remark** (CDF representation when $S_i \subseteq \mathbb{R}$)**.** Suppose $S_i$ is a subset of $\mathbb{R}$ (e.g. $S_i = \mathbb{R}$, $S_i = \mathbb{R}_+$, or $S_i = [\underline{x}, \overline{x}]$). Given a mixed strategy

$\sigma_i \in \Delta(S_i)$, there is a cumulative distribution function (CDF)

$$F_i : \mathbb{R} \to [0,1] \quad \text{such that} \quad F_i(x) = \sigma_i(\{\, s_i \in S_i : \ s_i \le x \,\}) \qquad \forall x \in \mathbb{R}. \tag{$*$}$$

Moreover, if $F_i : \mathbb{R} \to [0,1]$ is a CDF, then there is a unique mixed strategy $\sigma_i \in \Delta(S_i)$ such that $(*)$ holds for all $x \in \mathbb{R}$.

When we represent mixed strategies by CDFs, we will write $u_i(F_i, F_{-i})$ for player $i$'s expected payoff induced by the profile $(F_1, \ldots, F_I)$.

## *Example 5: A pricing game with capacity constraints*

**Example** (Two firms choose prices)**.** Two firms $i = 1, 2$ simultaneously choose prices

$$p_i \in \mathbb{R}_+.$$

There is a continuum of non-strategic consumers indexed by $[0, 1]$. Each consumer buys from the firm with the lower price, provided that the price is less than or equal to 1. Each firm has a stock (capacity) of goods $\bar{q} \in \left(\frac{1}{2}, 1\right)$ to sell, and all costs are sunk.

Let $q_i(p_i, p_{-i})$ denote firm $i$'s sales as a function of the price profile $(p_i, p_{-i})$:

$$q_i(p_i, p_{-i}) = \begin{cases} \bar{q} & \text{if } p_i < p_{-i} \text{ and } p_i \le 1, \\ 1 - \bar{q} & \text{if } p_{-i} < p_i \le 1, \\ \frac{1}{2} & \text{if } p_i = p_{-i} \le 1 \quad \text{(consumers split the market evenly)}, \\ 0 & \text{if } p_i > 1. \end{cases}$$

Hence firm $i$'s payoff (profit) is

$$u_i(p_i, p_{-i}) \ = \ p_i \, q_i(p_i, p_{-i}).$$

### *There is no pure-strategy Nash equilibrium*

**Proposition 0.0.8** (No PSNE in Example )**.** *Example  has no pure-strategy Nash equilibrium.*

*Proof.* Suppose, toward a contradiction, that $(p_1^*, p_2^*)$ is a Nash equilibrium.

**Step 1: cannot have $p_i^* > 1$ for any $i$.** If $p_i^* > 1$, then $q_i(p_i^*, p_{-i}^*) = 0$ and hence $u_i(p_i^*, p_{-i}^*) = 0$.

- If also $p_{-i}^* > 1$, then any deviation $p_i \in (0, 1]$ yields $u_i(p_i, p_{-i}^*) = p_i\bar{q} > 0$, contradicting optimality of $p_i^*$.
- If $p_{-i}^* \le 1$, then deviating to $p_i = p_{-i}^*$ yields $u_i(p_i, p_{-i}^*) = \frac{1}{2} p_i > 0$, again contradicting $u_i(p_i^*, p_{-i}^*) = 0$. Here we can also think in a small deviation towards a lower price than $p_{-i}^*$.

Therefore $p_i^* \leq 1$ for both firms.

**Step 2: cannot have** $p_1^* = p_2^* = 0$**.** If $p_1^* = p_2^* = 0$, then each firm earns 0. But if firm 1 deviates to any $p_1 \in (0,1]$ while $p_2^* = 0$, then firm 1 is the higher-priced firm (still within $[0,1]$) and sells $1 - \bar{q}$, so

$$u_1(p_1, p_2^*) = p_1(1 - \bar{q}) > 0,$$

a profitable deviation. Hence $(0,0)$ cannot be an equilibrium.

**Step 3: cannot have** $p_1^* = p_2^* = p^*$ **with** $0 < p^* \leq 1$**.** If $p_1^* = p_2^* = p^*$, then each firm sells $\frac{1}{2}$ and earns

$$u_1(p_1^*, p_2^*) = \tfrac{1}{2}p^*.$$

Consider firm 1 deviating to $p_1 = p^* - \varepsilon < p^*$ for some $\varepsilon > 0$. Then firm 1 becomes the lower-priced firm and sells $\bar{q}$, earning

$$u_1(p_1, p_2^*) = (p^* - \varepsilon)\bar{q}.$$

Because $\bar{q} > \frac{1}{2}$, we can choose $\varepsilon > 0$ small enough so that

$$(p^* - \varepsilon)\bar{q} > \tfrac{1}{2}p^*,$$

equivalently $\varepsilon < \dfrac{p^*(\bar{q} - \frac{1}{2})}{\bar{q}}$. Thus the deviation is profitable, contradicting equilibrium.

**Step 4: cannot have** $0 \leq p_i^* < p_{-i}^* \leq 1$**.** If $p_i^* < p_{-i}^*$, then firm $i$ is the lower-priced firm and sells $\bar{q}$, earning $u_i(p_i^*, p_{-i}^*) = p_i^*\bar{q}$. But then firm $i$ can raise its price slightly while remaining below $p_{-i}^*$: pick any $p_i$ with $p_i^* < p_i < p_{-i}^*$. Demand remains $\bar{q}$, so

$$u_i(p_i, p_{-i}^*) = p_i\bar{q} > p_i^*\bar{q} = u_i(p_i^*, p_{-i}^*),$$

a profitable deviation. Contradiction.

Steps 2–4 exhaust all possibilities with $p_i^* \leq 1$. Hence no pure-strategy Nash equilibrium exists. □

**Remark** (Goal going forward)**.** The goal is to construct a mixed-strategy Nash equilibrium (MSNE) for Example 5. Since we are only trying to *construct* an MSNE (and not characterize the full set of equilibria), we will aim to build one that satisfies "nice" properties. (If this approach fails, that is fine.)

**Remark** (Goal and "nice" properties for the construction)**.** We will try to construct a mixed-strategy Nash equilibrium $(F_1^*, F_2^*)$ with the following "nice" properties:

1. **Symmetry:** $F_1^* = F_2^* = F^*$.
   - *Rationale.* The game is symmetric (same action sets, demand rule, and capacity), so a symmetric equilibrium is a natural candidate.

2. **No mass above 1:** $\text{supp}(F^*) \subseteq [0, 1]$.

   - *Rationale.* If $p > 1$, demand is 0, hence profit is 0; such prices cannot be optimal if equilibrium yields positive profit.

3. **Interval support:** $\text{supp}(F^*) = [\underline{p}, 1]$ for some $\underline{p} \in (0, 1)$ (i.e. no gaps).

   - *Rationale.* A gap typically creates a profitable deviation into the gap: one can raise price without changing the event of winning/losing the market over that region.

4. **Continuity (no atoms):** $F^*$ is continuous on $(\underline{p}, 1)$ (often: on $[\underline{p}, 1]$).

   - *Rationale.* Atoms generate ties with positive probability, complicating payoffs; in many continuous-price games a tractable equilibrium can be constructed with no atoms.

5. **Indifference on the support:** the expected profit $\pi(p; F^*)$ is constant for all $p \in [\underline{p}, 1]$.

   - *Rationale.* If multiple prices are played with positive probability, they must all be best responses; otherwise the player would shift probability toward strictly better prices.

6. **Boundary conditions:** $F^*(\underline{p}) = 0$ and $F^*(1) = 1$.

   - *Rationale.* These normalize the support: $\underline{p}$ is the lowest price used, and all equilibrium mass lies at or below 1.

# *Class 6: Constructing a Symmetric Mixed-Strategy Equilibrium*

## *Candidate symmetric MSNE and support assumptions*

We now try to construct a mixed-strategy Nash equilibrium $(F_1^*, F_2^*)$ in the pricing game of the last class satisfying two key properties.

A) **Symmetry.** We look for a symmetric equilibrium:

$$F_1^* = F_2^* =: F^*.$$

B) **Interval support inside** $[0, 1]$**.** We assume that the equilibrium CDF $F^*$ has support contained in $[0, 1]$ and, more precisely, that there exist

$$\underline{p} < \overline{p} \qquad \text{with} \qquad [\underline{p}, \overline{p}] \subseteq [0, 1]$$

such that

$$F^*(p) = 0 \quad \text{for } p < \underline{p}, \qquad F^*(p) = 1 \quad \text{for } p > \overline{p},$$

and $F^*(\cdot)$ is strictly increasing on the interval $[\underline{p}, \overline{p}]$.

Intuitively, prices below $\underline{p}$ and above $\overline{p}$ are never played in equilibrium. All equilibrium mass is concentrated on the closed interval $[\underline{p}, \overline{p}]$, and within this interval the CDF has no flat segments.

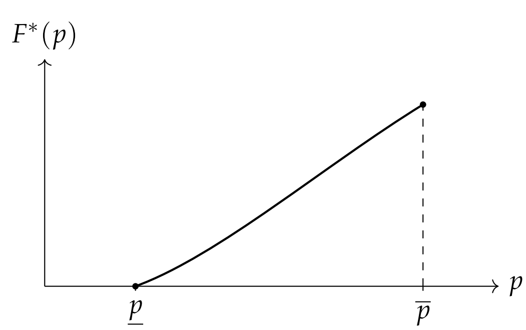


Candidate equilibrium CDF $F^*$ with support $[\underline{p}, \overline{p}] \subseteq [0, 1]$ is strictly increasing in the support.

### *No atoms and continuity of the CDF*

We now impose that $F^*$ has no atoms on its support.

**Remark** (No atoms)**.** The "no atoms" requirement is that every singleton $\{p\}$ with $p \in \left[\underline{p}, \overline{p}\right]$ has zero probability under $F^*$. Equivalently,

$$\sigma^*\big(\{p\}\big) = 0 \qquad \text{for all } p \in [\underline{p}, \overline{p}],$$

where $\sigma^*$ is the probability measure associated with $F^*$.

Graphically, an *atom* at some $\tilde{p} \in [\underline{p}, \overline{p}]$ corresponds to a jump of the CDF at $\tilde{p}$: the left limit and the value at $\tilde{p}$ are different.

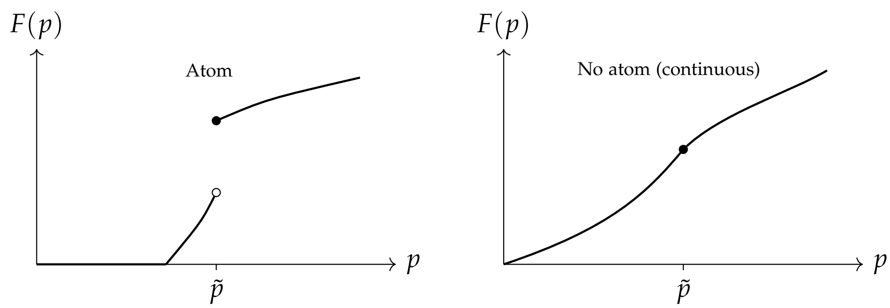


A jump (atom) at $\tilde{p}$ versus a continuous CDF with no atom.

Formally, suppose there were an atom at $\tilde{p} \in [\underline{p}, \overline{p}]$. Then there exists a sequence $(p_n)_{n\in\mathbb{N}}$ with

(i) $p_n < \tilde{p}$ for all $n$,

(ii) $\lim_{n\to\infty} p_n = \tilde{p}$,

(iii) $\lim_{n\to\infty} F^*(p_n) < F^*(\tilde{p})$.

Condition (iii) captures the jump of the CDF at $\tilde{p}$. Ruling out atoms exactly rules out such jumps and, combined with the standard right-continuity of a CDF, implies that $F^*$ is continuous on $[\underline{p}, \overline{p}]$.

### *Expected payoff given a candidate CDF*

Fix a candidate CDF $F^*_{-i}$ for the opponent (in a symmetric equilibrium we will have $F^*_{-i} = F^*$). For any price $p_i \in [0,1]$, firm $i$'s expected payoff can be written as

$$u_i(p_i, F^*_{-i}) = p_i\Big[\Pr(p_{-i} < p_i \mid F^*_{-i})\,(1-\bar{q}) + \Pr(p_{-i} > p_i \mid F^*_{-i})\,\bar{q} + \Pr(p_{-i} = p_i \mid F^*_{-i})\cdot\tfrac{1}{2}\Big].$$

The three terms in brackets correspond to the events in which firm $i$ is the high-priced seller (sells $1-\bar{q}$), the low-priced seller (sells $\bar{q}$), and the case of a tie (each sells $\frac{1}{2}$), respectively.

Under the "no atoms" assumption, we have

$$\Pr(p_{-i} = p_i \mid F^*_{-i}) = 0,$$

so that

$$\Pr(p_{-i} < p_i \mid F^*_{-i}) + \Pr(p_{-i} > p_i \mid F^*_{-i}) = 1.$$

Moreover, since $F^*_{-i}$ is the CDF of the opponent's price,

$$\Pr(p_{-i} < p_i \mid F^*_{-i}) = F^*_{-i}(p_i), \qquad \Pr(p_{-i} > p_i \mid F^*_{-i}) = 1 - F^*_{-i}(p_i).$$

Substituting these expressions into the expected payoff yields

$$u_i(p_i, F^*_{-i}) = p_i \Big[ F^*_{-i}(p_i)(1 - \bar{q}) + \big(1 - F^*_{-i}(p_i)\big)\bar{q} \Big].$$

In a symmetric candidate equilibrium, we set $F^*_{-i} = F^*$ and view $u_i(p_i, F^*)$ as a function of the own price $p_i$ and the common CDF $F^*$.

### *Using the finite-game trick: indifference on the support*

We now borrow the "finite-game trick" from the characterization of mixed equilibria.

**Remark** (Indifference on the support as an equilibrium heuristic)**.** We will impose the following requirements on a candidate symmetric MSNE:

- For every $p_i \in [\underline{p}, \overline{p}]$, firm $i$'s expected payoff $u_i(p_i, F^*)$ is at least as high as the expected payoff from any deviation $p_i \notin [\underline{p}, \overline{p}]$.
- For any $p_i, p_i' \in [\underline{p}, \overline{p}]$, the expected payoffs coincide:

$$u_i(p_i, F^*) = u_i(p_i', F^*).$$

The second requirement enforces *indifference on the support*: all prices that are played with positive probability must yield the same expected payoff.

Starting from these requirements, we will derive conditions pinning down $F^*$ and the endpoints $\underline{p}, \overline{p}$, and then verify ex post that the resulting $F^*$ indeed defines a mixed-strategy Nash equilibrium.

### *Deriving the equilibrium CDF and verifying optimality*

*Step 1: Pinning down the upper bound of the support.* We first determine the upper endpoint $\overline{p}$ of the support.

For $\overline{p}$ to belong to the support of $F^*$, pricing at $\overline{p}$ must yield at least as high an expected payoff as any deviation. In particular, we must have

$$u_i(\overline{p}, F^*_{-i}) \;\geq\; u_i(1, F^*_{-i}).$$

Using the payoff formula derived earlier,

$$u_i(p_i, F^*_{-i}) = p_i\Big[F^*_{-i}(p_i)(1-\bar{q}) + \big(1 - F^*_{-i}(p_i)\big)\bar{q}\Big],$$

and recalling that $F^*_{-i}(\overline{p}) = 1$, we obtain

$$u_i(\overline{p}, F^*_{-i}) = \overline{p}(1-\bar{q}), \qquad u_i(1, F^*_{-i}) = 1 \cdot (1-\bar{q}).$$

Thus the incentive constraint becomes

$$\overline{p}(1-\bar{q}) \;\geq\; (1-\bar{q}).$$

Since $1-\bar{q} > 0$, this implies $\overline{p} \geq 1$. Because prices above 1 yield zero demand, we conclude

$$\overline{p} = 1.$$

*Step 2: Indifference on the support.* Since $\overline{p} = 1$ belongs to the support and yields expected payoff $(1-\bar{q})$, every price in the support must deliver the same expected payoff.

Hence, for all $p_i \in [\underline{p}, 1]$,

$$u_i(p_i, F^*_{-i}) = u_i(1, F^*_{-i}) = (1-\bar{q}).$$

Substituting the payoff expression gives

$$p_i\Big[F^*(p_i)(1-\bar{q}) + \big(1 - F^*(p_i)\big)\bar{q}\Big] = (1-\bar{q}).$$

Rearranging,

$$p_i\Big[F^*(p_i)(1-2\bar{q}) + \bar{q}\Big] = (1-\bar{q}),$$

which yields

$$F^*(p_i) = \frac{p_i\bar{q} - (1-\bar{q})}{p_i(2\bar{q}-1)}.$$

*Step 3: Determining the lower bound $\underline{p}$.* By the no-atoms assumption, the CDF satisfies $F^*(\underline{p}) = 0$. Imposing this condition gives

$$\frac{\underline{p}\,\bar{q} - (1-\bar{q})}{\underline{p}(2\bar{q}-1)} = 0,$$

which implies

$$\underline{p} = \frac{1-\bar{q}}{\bar{q}}.$$

Thus the candidate equilibrium CDF is

$$F^*(p) = \begin{cases} 0 & \text{if } p < \dfrac{1-\bar{q}}{\bar{q}}, \\[2ex] \dfrac{p\bar{q} - (1-\bar{q})}{p(2\bar{q}-1)} & \text{if } p \in \left[\dfrac{1-\bar{q}}{\bar{q}}, 1\right], \\[2ex] 1 & \text{if } p > 1. \end{cases}$$

*Step 4: Properties of the candidate CDF.* The function $F^*$ is:

- continuous on $\left[\frac{1-\bar{q}}{\bar{q}}, 1\right]$,
- strictly increasing on its support,
- satisfies $F^*(\underline{p}) = 0$ and $F^*(1) = 1$,
- assigns zero probability to every singleton (no atoms).

Hence $F^*$ is a valid CDF satisfying all maintained assumptions.

*Step 5: Verifying that $(F^*, F^*)$ is an MSNE.* By Lemma 0.0.5, it suffices to check that no pure-price deviation can yield a higher payoff.

We already know that:

$$u_i(p_i, F^*) = (1 - \bar{q}) \quad \text{for all } p_i \in [\underline{p}, 1].$$

*Deviations above the support.* If $p_i > 1$, then demand is zero and

$$u_i(p_i, F^*) = 0 < (1 - \bar{q}) = u_i(F^*, F^*).$$

*Deviations below the support.* If $p_i < \underline{p}$, then $F^*(p_i) = 0$, so firm $i$ is always the low-priced seller. Thus,

$$u_i(p_i, F^*) = p_i\bar{q}.$$

Since $p_i < \underline{p} = \frac{1-\bar{q}}{\bar{q}}$, we have

$$p_i\bar{q} < \underline{p}\,\bar{q} = 1 - \bar{q} = u_i(F^*, F^*).$$

Therefore, no deviation below the support is profitable.

Combining all cases, we conclude that

$$u_i(F^*, F^*) \;\geq\; u_i(p_i, F^*) \qquad \text{for all } p_i \in \mathbb{R}_+,$$

with equality for $p_i \in [\underline{p}, 1]$. Hence $(F^*, F^*)$ is a *mixed-strategy Nash equilibrium*.

# *Class 7: The Pricing Game Equilibrium and Equilibrium Existence*

## *Equilibrium CDF revisited*

We continue with the pricing game of Classes 5 and 6. Recall that each firm $i = 1, 2$ chooses a price

$$S_i = \mathbb{R}_+,$$

and that capacities satisfy

$$\bar{q} \in \left(\tfrac{1}{2}, 1\right).$$

**Proposition 0.0.9** (Symmetric mixed-strategy equilibrium)**.** *In the pricing game of Example , there exists a symmetric mixed-strategy Nash equilibrium* $(F_1^*, F_2^*)$ *with*

$$F_1^* = F_2^* =: F^*,$$

*where the common CDF* $F^*$ *is given by*

$$F^*(p) = \begin{cases} 0 & \text{if } p < \underline{p} := \dfrac{1-\bar{q}}{\bar{q}}, \\ \dfrac{\bar{q}}{2\bar{q}-1} - \dfrac{1-\bar{q}}{p(2\bar{q}-1)} & \text{if } p \in [\underline{p}, 1], \\ 1 & \text{if } p > 1. \end{cases}$$

*Equivalently, on* $[\underline{p}, 1]$ *one may write*

$$F^*(p) = \frac{p\bar{q} - (1-\bar{q})}{p(2\bar{q}-1)}.$$

**Remark.** The support of $F^*$ is $[\underline{p}, 1]$. On this interval, $F^*$ is continuous, strictly increasing, and has no atoms. Outside $[\underline{p}, 1]$, the CDF is flat at 0 (for $p < \underline{p}$) and at 1 (for $p > 1$).

## *A remark on starting from a PDF*

**Remark** (Starting instead from a PDF)**.** One could try to start the analysis by postulating directly an equilibrium $(F_1^*, F_2^*)$ such that $F_1^* = F_2^* =$

$F^*$ and $F^*$ is a CDF that *admits* a probability density function (PDF) $f^*$ with the following properties:

- there exist $\underline{p} < \overline{p}$ with $[\underline{p}, \overline{p}] \subseteq [0,1]$,
- $f^*(p) > 0$ for all $p \in (\underline{p}, \overline{p})$,
- $f^*(p) = 0$ for $p \notin (\underline{p}, \overline{p})$.

In other words, one would posit from the outset that prices are continuously distributed on an interval, with a strictly positive density in the interior and zero density outside.

The key technical notion here is that of *absolute continuity*.

**Definition** (Absolutely continuous CDF)**.** A CDF $F : \mathbb{R} \to [0,1]$ is *absolutely continuous* if there exists an integrable function $f : \mathbb{R} \to [0, \infty)$ such that, for all $x \in \mathbb{R}$,

$$F(x) = F(-\infty) + \int_{(-\infty,x]} f(t)\,dt.$$

In this case, $f$ is called a (version of the) PDF associated with $F$.

**Remark** (Absolute continuity versus continuity)**.** If $F$ is absolutely continuous, then $F$ is automatically continuous. However, the converse is not true: there are CDFs that are continuous but not absolutely continuous (e.g. the Cantor distribution). Such CDFs have no PDF in the usual sense.

Thus:

$$\text{absolute continuity} \implies \text{continuity}, \qquad \text{but not conversely.}$$

**Remark** (What does a PDF tell us, and what does it not?)**.** This connects directly to a natural question:

*If a CDF $F$ is continuous and admits a PDF $f$, does the PDF tell us that the probability of each specific point $p$ is zero? Or does it only tell us the probability of intervals (or more general sets)?*

The correct interpretation is as follows.

1. If $F$ is absolutely continuous with density $f$, then for any interval $[a, b]$,
$$\mathbb{P}(X \in [a,b]) = F(b) - F(a-) = \int_a^b f(t)\,dt.$$
That is, the PDF describes how probability is *distributed over intervals*, not the probability of a single point.

2. For any fixed point $p$, the probability of the singleton $\{p\}$ is
$$\mathbb{P}(X = p) = F(p) - \lim_{x \uparrow p} F(x).$$

**What are $P$, $p$, and $f(p)$?**

Here $P$ denotes a *random variable* (the opponent's price drawn from a mixed strategy), while $p$ denotes a *real number* (a specific price level).

The CDF $F(p) = \mathbb{P}(P \leq p)$ describes the distribution of the random variable $P$. When $F$ is absolutely continuous, it admits a density $f$ such that

$$f(p) = \frac{dF(p)}{dp},$$

i.e. $f$ is the derivative of the CDF wherever this derivative exists.

**Local interpretation.** For small $\varepsilon > 0$,

$$f(p) \approx \frac{\mathbb{P}(p \leq P \leq p + \varepsilon)}{\varepsilon}.$$

Thus, $f(p)$ measures how much probability mass *per unit of length* is concentrated in a small neighborhood around $p$.

- If $f(p)$ is large, a lot of probability is "packed" near $p$.
- If $f(p) = 0$, there is essentially no probability mass around $p$.

Crucially, $f(p)$ is *not* the probability of playing price $p$. Indeed,

$$\mathbb{P}(P = p) = 0$$

whenever $F$ is continuous, regardless of the value of $f(p)$. What matters is that

$$f(p) > 0 \implies \mathbb{P}(|P - p| < \varepsilon) > 0 \quad \forall \varepsilon > 0,$$

i.e. every neighborhood of $p$ carries positive probability mass.

If $F$ is continuous at $p$ (in particular, if $F$ is absolutely continuous), this difference is zero:

$$\mathbb{P}(X = p) = 0.$$

This conclusion does *not* come from the numerical value of $f(p)$, but from the continuity of $F$ (no jump at $p$) and the fact that the integral of $f$ over the single point $\{p\}$ is zero.

3. The value $f(p)$ may be positive, zero, or even undefined at individual points without changing any probabilities. Indeed, if $f$ and $g$ differ only on a set of Lebesgue measure zero (for example, at finitely many points), they generate the same CDF:

$$\int_a^b f(t)\,dt = \int_a^b g(t)\,dt \quad \text{for all } a < b.$$

Hence the PDF is only defined *up to changes on sets of measure zero*.

In summary:

- Continuity of $F$ (no jumps) implies that every singleton $\{p\}$ has probability zero.
- Absolute continuity is a stronger requirement: it gives a PDF $f$ such that probabilities of intervals come from integrating $f$.
- The fact that singletons have probability zero does *not* mean that $f(p)$ must be zero; it means that the integral of $f$ over a set of length zero is zero.

For our pricing game, assuming from the outset that $F^*$ admits a PDF $f^*$ with support on an interval is a way of "baking in" both the absence of atoms and the idea that mixed strategies are genuinely *continuous* on that interval. In the previous classes we instead derived these properties from the equilibrium conditions, rather than imposing absolute continuity directly.

## *Interpreting the equilibrium density*

In the previous discussion we noted that one could equivalently describe the symmetric equilibrium by specifying a probability density function (PDF) $f^*$ supported on an interval $[\underline{p}, \overline{p}] \subseteq [0, 1]$ and then setting

$$F^*(p) = \int_{-\infty}^{p} f^*(t)\,dt.$$

**Remark** (Support of the equilibrium density)**.** The requirement that

$$f^*(p) > 0 \quad \Longleftrightarrow \quad p \in [\underline{p}, \overline{p}]$$

simply says that the *support* of the density $f^*$ is the interval $[\underline{p}, \overline{p}]$:

- for $p \in (\underline{p}, \overline{p})$, every neighborhood of $p$ carries positive probability mass;
- for $p \notin [\underline{p}, \overline{p}]$, there is no probability mass in any neighborhood of $p$.

Because $F^*$ is continuous,

$$\Pr(p_i = p \mid F_i^*) = 0 \quad \text{for all } p \in \mathbb{R},$$

while for any $p$ we have

$$\Pr(p_i \leq p \mid F_i^*) = F_i^*(p).$$

Thus $f^*$ governs how probability is spread over *intervals*, whereas the CDF $F^*$ gives the cumulative probability up to each price level.

## *Existence of pure and mixed-strategy equilibrium*

Up to now our focus has been largely *constructive*: given a specific game, we have characterized Nash equilibria—sometimes in pure strategies, sometimes in mixed strategies—and in a few examples we have even computed them in closed form.

From the point of view of applications, however, this is rarely the end goal. There are two related but conceptually distinct aims:

- **Prediction in a given model.** In principle, one might want exact predictions: the precise price distribution in a pricing game, the exact quantities in Cournot competition, and so on. This requires solving explicitly for an equilibrium (or for the set of equilibria).
- **Comparative statics.** More often, we care about how equilibrium outcomes move when underlying primitives change: e.g. how an increase in costs shifts equilibrium prices, or how a change in preferences affects equilibrium allocations. Here the *functional form* of equilibrium strategies is often less important than the *direction* or qualitative nature of the response.

**Remark** (Equilibrium conditions without full solution)**.** If we are primarily interested in comparative statics, we can sometimes work directly with the *equilibrium conditions* (best-response inequalities, first-order conditions, indifference relations, market-clearing constraints) to derive testable restrictions, without ever computing an explicit closed form for the equilibrium strategies or payoffs.

This "implicit" approach uses equilibrium as a set of *restrictions* on behavior, rather than as an explicit formula.

There is, however, a crucial conceptual step in both approaches:

- When we write down and manipulate equilibrium conditions, we are implicitly assuming that *some* equilibrium exists whose strategies and payoffs satisfy those conditions.
- If no equilibrium exists, then the comparative-statics exercise is ill-posed: there is no well-defined object whose response to parameter changes we can study.

This raises the fundamental question:

*Does an equilibrium exist?*

In what follows, we will move from constructing equilibria in specific examples to general *existence theorems* for pure and mixed-strategy Nash equilibrium. These results justify treating equilibrium conditions as meaningful restrictions in applied work: they guarantee that under appropriate assumptions, there is at least one profile of strategies and beliefs satisfying those conditions.

## *Best-response correspondences and fixed points*

We now rephrase the definition of Nash equilibrium in a way that will be convenient for existence results: as a *fixed point* of a suitable best-response correspondence.

### *Pure strategies*

For each player $i$, recall the pure best-response correspondence

$$\mathrm{BR}_i : S_{-i} \rightrightarrows S_i, \qquad \mathrm{BR}_i(s_{-i}) = \{s_i \in S_i : u_i(s_i, s_{-i}) \geq u_i(r_i, s_{-i}) \,\forall r_i \in S_i\}.$$

It is convenient to aggregate all players' best responses into a single map defined on the set of strategy profiles,

$$S := \prod_{i=1}^{I} S_i.$$

**Definition** (Aggregate pure best-response correspondence)**.** Define $\mathrm{BR} : S \rightrightarrows S$ by

$$\mathrm{BR}(s_1, \ldots, s_I) = \prod_{i=1}^{I} \mathrm{BR}_i(s_{-i}),$$

i.e. the set of all profiles in which each player $i$ chooses a pure best response to $s_{-i}$.

**Proposition 0.0.10** (Pure-strategy NE as fixed points)**.** *A profile* $s^* = (s_1^*, \ldots, s_I^*) \in S$ *is a pure-strategy Nash equilibrium (PSNE) if and only if*

$$s^* \in \mathrm{BR}(s^*).$$

*Proof.* By definition, $s^*$ is a PSNE iff $s_i^* \in \text{BR}_i(s_{-i}^*)$ for all $i$. This is equivalent to $(s_1^*, \ldots, s_I^*) \in \prod_i \text{BR}_i(s_{-i}^*) = \text{BR}(s^*)$. □

**Example** (Fixed point intuition with single-valued best responses)**.** Consider the Cournot duopoly from Class 2 with linear demand $P(Q) = a - bQ$ and zero costs. Each firm $i$ chooses a quantity $q_i \geq 0$. Fix $q_{-i}$ and ignore the non-negativity constraint on price. Firm $i$'s best response is

$$\text{BR}_i(q_{-i}) = \frac{a}{2b} - \frac{1}{2}q_{-i},$$

so in this example $\text{BR}_i$ is a *function* (single-valued) from $\mathbb{R}_+$ to $\mathbb{R}_+$.

Graphically, each player's best-response function is a downward sloping line in the $(q_1, q_2)$ plane. A pair $(q_1^*, q_2^*)$ is a Cournot equilibrium if and only if each firm is best-responding to the other, i.e.

$$q_1^* = \text{BR}_1(q_2^*), \qquad q_2^* = \text{BR}_2(q_1^*).$$

In the diagram, this means $(q_1^*, q_2^*)$ lies *on both lines simultaneously*: it is the intersection point of the two best-response functions. Equivalently, the profile $q^* = (q_1^*, q_2^*)$ is a fixed point of the aggregate best-response mapping $q \mapsto (\text{BR}_1(q_2), \text{BR}_2(q_1))$.

### *Mixed strategies*

In the mixed extension, the strategy space for player $i$ is the simplex $\Delta(S_i)$ of probability distributions over $S_i$, and a mixed profile is $\sigma = (\sigma_1, \ldots, \sigma_I) \in \prod_i \Delta(S_i)$.

For each $i$, recall the mixed best-response correspondence

$$\widetilde{\text{BR}}_i : \prod_{j \neq i} \Delta(S_j) \rightrightarrows \Delta(S_i), \qquad \widetilde{\text{BR}}_i(\sigma_{-i}) = \left\{\sigma_i \in \Delta(S_i) : u_i(\sigma_i, \sigma_{-i}) \geq u_i(\sigma_i', \sigma_{-i}) \; \forall \sigma_i' \in \Delta(S_i)\right\}.$$

**Definition** (Aggregate mixed best-response correspondence)**.** Define

$$\widetilde{\text{BR}} : \prod_{i=1}^{I} \Delta(S_i) \rightrightarrows \prod_{i=1}^{I} \Delta(S_i)$$

by

$$\widetilde{\text{BR}}(\sigma_1, \ldots, \sigma_I) = \prod_{i=1}^{I} \widetilde{\text{BR}}_i(\sigma_{-i}).$$

**Proposition 0.0.11** (MSNE as fixed points)**.** *A profile $\sigma^* = (\sigma_1^*, \ldots, \sigma_I^*)$ is a mixed-strategy Nash equilibrium (MSNE) if and only if*

$$\sigma^* \in \widetilde{\text{BR}}(\sigma^*).$$

*Proof.* Exactly parallel to the pure-strategy case: by definition, $\sigma^*$ is an MSNE iff $\sigma_i^* \in \widetilde{\text{BR}}_i(\sigma_{-i}^*)$ for all $i$, which is equivalent to $\sigma^* \in \prod_i \widetilde{\text{BR}}_i(\sigma_{-i}^*) = \widetilde{\text{BR}}(\sigma^*)$. □

**Example** (Best responses in matching pennies as a fixed point)**.** Define all correspondences as functions first. Return to the matching-pennies game from Class 3. Let $p$ be the probability with which player 1 plays $H$, and $q$ the probability with which player 2 plays $H$. Thus the mixed-strategy space can be identified with the unit square $[0, 1] \times [0, 1]$.

Given $q$, player 1's expected payoffs from $H$ and $T$ are

$$u_1(H; q) = q \cdot 1 + (1 - q) \cdot (-1) = 2q - 1, \qquad u_1(T; q) = -u_1(H; q) = 1 - 2q.$$

Hence player 1's best response is:

$$\mathrm{BR}_1(q) = \begin{cases} \{1\} & \text{if } q > 1/2 \quad (H \text{ is strictly better}), \\ [0,1] & \text{if } q = 1/2 \quad (H \text{ and } T \text{ are tied}), \\ \{0\} & \text{if } q < 1/2 \quad (T \text{ is strictly better}), \end{cases}$$

where we now understand $\mathrm{BR}_1(q)$ as a subset of $[0, 1]$ (the admissible $p$'s).

A symmetric calculation gives player 2's best response $\mathrm{BR}_2(p)$. In the $(p, q)$ plane, each best response is a correspondence with a *kink* at $1/2$. The unique pair $(p^*, q^*)$ such that

$$p^* \in \mathrm{BR}_1(q^*), \qquad q^* \in \mathrm{BR}_2(p^*)$$

is $(1/2, 1/2)$, i.e. the unique MSNE we already computed. Thus the equilibrium is precisely the fixed point of the aggregate best-response correspondence on the compact, convex set $[0, 1] \times [0, 1]$.

### *From best responses to Kakutani*

The fixed-point characterization suggests a general existence strategy:

> View $\widetilde{\mathrm{BR}}$ as a set-valued map from a compact, convex subset of $\mathbb{R}^n$ into itself and apply a *fixed-point theorem* to guarantee that $\widetilde{\mathrm{BR}}$ has a fixed point.

**Recall** (Kakutani's fixed-point theorem)**.** Let $X \subset \mathbb{R}^n$ be non-empty, compact, and convex, and let $F : X \rightrightarrows X$ be non-empty-valued, convex-valued, and have a closed graph. Then $F$ has a fixed point, i.e. there exists $x^* \in X$ such that $x^* \in F(x^*)$.

In finite games:

- each $S_i$ is finite, so $\Delta(S_i)$ is a simplex in $\mathbb{R}^{m_i}$: non-empty, compact, and convex;
- payoffs $u_i$ are continuous in mixed strategies, and linear in each player's own mixed strategy, so each $\widetilde{\mathrm{BR}}_i$ is non-empty-valued, convex-valued, and has a closed graph;

- thus the product correspondence $\widetilde{\mathrm{BR}}$ inherits these properties on

$$X = \prod_{i=1}^{I} \Delta(S_i) \subset \mathbb{R}^{\sum_i m_i}.$$

Kakutani's theorem then ensures that $\widetilde{\mathrm{BR}}$ has a fixed point $\sigma^* \in X$, which, by the previous proposition, is precisely a mixed-strategy Nash equilibrium.

This is the backbone of the standard *existence theorem* for MSNE in finite games: even when no pure-strategy equilibrium exists (as in matching pennies), the geometry of the mixed-strategy simplices and the convexity of best responses guarantee that there *must* exist a fixed point of the best-response correspondence, i.e. a mixed equilibrium.

### *Mixed strategies as a simplex*

Fix a player $i$ with a finite pure strategy set

$$S_i = \{s_i^1, \ldots, s_i^{m_i}\}.$$

A mixed strategy $\sigma_i \in \Delta(S_i)$ can be identified with the probability vector

$$\sigma_i = \big(p(s_i^1), \ldots, p(s_i^{m_i})\big) \in \mathbb{R}^{m_i}$$

such that $p(s_i^k) \geq 0$ for all $k$ and $\sum_{k=1}^{m_i} p(s_i^k) = 1$. Hence $\Delta(S_i)$ is the standard simplex in $\mathbb{R}^{m_i}$: a non-empty, compact, convex subset of $\mathbb{R}^{m_i}$.

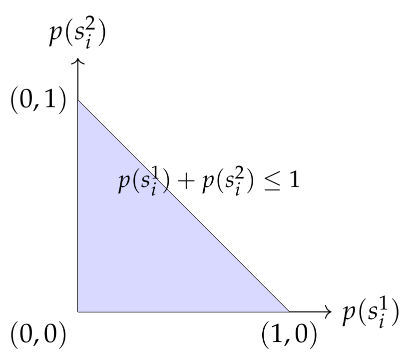


The simplex $\Delta(S_i)$ when $|S_i| = 3$. Here $p(s_i^3) = 1 - p(s_i^1) - p(s_i^2)$.

More generally, if $|S_i| = m_i$, then $\Delta(S_i)$ can be identified with a simplex in $\mathbb{R}^{m_i - 1}$, and

$$\prod_{i=1}^{I} \Delta(S_i)$$

is a non-empty, compact, convex subset of some Euclidean space $\mathbb{R}^N$.

**Step 0.** For a finite game, each mixed strategy set $\Delta(S_i)$ is a non-empty, compact, convex subset of $\mathbb{R}^{m_i}$, and the product $\prod_i \Delta(S_i)$ is non-empty, compact, and convex in $\mathbb{R}^N$. This makes the mixed-strategy best-response correspondence a natural candidate for applying Kakutani's fixed-point theorem.

*Step 1: non-emptiness of mixed best responses*

Fix $i$ and $\sigma_{-i} \in \prod_{j \neq i} \Delta(S_j)$. Recall the mixed best-response correspondence of player $i$,

$$\widetilde{BR}_i(\sigma_{-i}) = \Big\{\sigma_i \in \Delta(S_i) : u_i(\sigma_i, \sigma_{-i}) \geq u_i(\sigma_i', \sigma_{-i}) \; \forall \sigma_i' \in \Delta(S_i)\Big\}.$$

**Lemma 0.0.12** (Step 1: $\widetilde{BR}_i$ is non-empty valued)**.** *For every finite game and for every $\sigma_{-i}$, the set $\widetilde{BR}_i(\sigma_{-i})$ is non-empty.*

This point is very important to solve exercises. In many of them we only need to look at finiteness of the game and then apply this theorem to ensure there is a MSNE.

*Proof.* Since $S_i$ is finite, the set of pure-strategy payoffs

$$\big\{u_i(s_i, \sigma_{-i}) : s_i \in S_i\big\} \subset \mathbb{R}$$

is a finite subset of $\mathbb{R}$ and therefore has a maximum. Hence there exists at least one pure strategy $s_i^* \in S_i$ such that

$$u_i(s_i^*, \sigma_{-i}) \;\geq\; u_i(r_i, \sigma_{-i}) \quad \text{for all } r_i \in S_i.$$

Thus $s_i^*$ is a (pure) best response to $\sigma_{-i}$.

Now consider the degenerate mixed strategy $\sigma_i^* \in \Delta(S_i)$ that assigns probability 1 to $s_i^*$. By Lemma 0.0.5, checking deviations in mixed strategies is equivalent to checking deviations to pure strategies. Hence $\sigma_i^*$ is also a best response to $\sigma_{-i}$ in the mixed game, i.e. $\sigma_i^* \in \widetilde{BR}_i(\sigma_{-i})$. Therefore $\widetilde{BR}_i(\sigma_{-i}) \neq \varnothing$. □

**Remark** (Key: $S_i$ is finite.)**.** The key point in the proof is the existence of $s_i^* \in S_i$ that maximizes $u_i(\,\cdot\,, \sigma_{-i})$ over $S_i$. This does *not* require any heavy analysis: it follows from the elementary fact that a real-valued function on a finite set always attains a maximum. If one prefers a more general analytic statement, it is a special case of the extreme value theorem: a continuous function on a compact set attains its maximum and minimum, and any finite set is compact in the Euclidean topology.

# Class 8: Best–response correspondence and existence of MSNE

## Continuation

We keep working with a finite normal–form game

$$G = \big(N,\, (S_i)_{i\in N},\, (u_i)_{i\in N}\big),$$

and with the mixed–strategy version of the best–response correspondence. Recall that for each $i$,

$$\widetilde{BR}_i : \prod_{j\neq i} \Delta(S_j) \;\rightrightarrows\; \Delta(S_i)$$

maps a profile of opponents' mixed strategies $\sigma_{-i}$ into the set of $i$'s mixed best replies

$$\widetilde{BR}_i(\sigma_{-i}) = \arg\max_{\sigma_i\in\Delta(S_i)} u_i(\sigma_i,\sigma_{-i}).$$

*Step 2 (Convex–valued).*

For each fixed $\sigma_{-i}$, the set $\widetilde{BR}_i(\sigma_{-i})$ is convex.[2] Hence $\widetilde{BR}_i$ is convex–valued.

[2] This is proved in Problem Set 3. The intuition is that $u_i(\sigma_i,\sigma_{-i})$ is affine in $\sigma_i$, so whenever $\sigma_i^1$ and $\sigma_i^2$ are best replies, any convex combination $\lambda\sigma_i^1 + (1-\lambda)\sigma_i^2$ with $\lambda\in[0,1]$ yields the same expected payoff and is therefore also a best reply.

*Step 3 (Closed graph).*

We now show that the graph of $\widetilde{BR}_i$ is closed. Recall the definition:

$$\text{graph}\big(\widetilde{BR}_i\big) = \Big\{\, (\sigma_{-i},\sigma_i) \in \Big(\textstyle\prod_{j\neq i}\Delta(S_j)\Big)\times\Delta(S_i) : \sigma_i\in\widetilde{BR}_i(\sigma_{-i}) \,\Big\}.$$

Consider a sequence

$$\big(\sigma_{-i}^n,\sigma_i^n\big)_{n\in\mathbb{N}} \subset \text{graph}\big(\widetilde{BR}_i\big)$$

such that

$$\big(\sigma_{-i}^n,\sigma_i^n\big) \;\longrightarrow\; \big(\sigma_{-i}^*,\sigma_i^*\big) \quad \text{in} \quad \Big(\textstyle\prod_{j\neq i}\Delta(S_j)\Big)\times\Delta(S_i).$$

We want to prove that

$$\big(\sigma^*_{-i}, \sigma^*_i\big) \in \text{graph}\big(\widetilde{BR}_i\big), \quad \text{i.e.} \quad \sigma^*_i \in \widetilde{BR}_i(\sigma^*_{-i}).$$

By definition of $\widetilde{BR}_i$, for every $n$ and every $\sigma_i \in \Delta(S_i)$ we have

$$u_i(\sigma^n_i, \sigma^n_{-i}) \;\geq\; u_i(\sigma_i, \sigma^n_{-i}). \tag{1}$$

*Intuition.* If all the sequence is part of the graph then it is a best response at some extent. So its payoff should be greater or equal than the payoff of every other strategy (that is not part of the graph).

Fix an arbitrary $\sigma_i \in \Delta(S_i)$ and define the function

$$f : \Delta(S_i) \times \prod_{j \neq i} \Delta(S_j) \to \mathbb{R}, \qquad f(\sigma_i, \sigma_{-i}) := u_i(\sigma_i, \sigma_{-i}) - u_i(\sigma_i, \sigma_{-i}).$$

Then (1) is equivalent to

$$f(\sigma^n_i, \sigma^n_{-i}) \;\geq\; 0 \quad \text{for all } n \in \mathbb{N}.$$

Because the game is finite, $u_i$ is continuous on $\prod_{j \in N} \Delta(S_j)$; hence $f$ is continuous on $\Delta(S_i) \times \prod_{j \neq i} \Delta(S_j)$. The sequence $(\sigma^n_i, \sigma^n_{-i})$ converges to $(\sigma^*_i, \sigma^*_{-i})$, so by continuity of $f$ we obtain

$$f(\sigma^n_i, \sigma^n_{-i}) \;\longrightarrow\; f(\sigma^*_i, \sigma^*_{-i}).$$

The left–hand side is a sequence of non–negative real numbers (by (1)), therefore its limit is also non–negative:

$$f(\sigma^*_i, \sigma^*_{-i}) \;\geq\; 0.$$

Rewriting the definition of $f$,

$$u_i(\sigma^*_i, \sigma^*_{-i}) \;-\; u_i(\sigma_i, \sigma^*_{-i}) \;\geq\; 0 \quad \text{for every } \sigma_i \in \Delta(S_i).$$

Thus $\sigma^*_i$ yields at least as high an expected payoff as any $\sigma_i \in \Delta(S_i)$ against $\sigma^*_{-i}$, which means $\sigma^*_i \in \widetilde{BR}_i(\sigma^*_{-i})$. Hence

$$\big(\sigma^*_{-i}, \sigma^*_i\big) \in \text{graph}\big(\widetilde{BR}_i\big),$$

so $\text{graph}(\widetilde{BR}_i)$ is closed.

**Theorem 0.0.13** (Finite games)**.** *In any finite game there exists a mixed–strategy Nash equilibrium.*

## *Existence of equilibrium beyond finite games*

Up to now we have focused on *finite* games, i.e. games in which each player $i$ has a finite pure–strategy set $S_i$. What about infinite games?

**Theorem 0.0.14** (Glicksberg)**.** *Fix a game* $G = (S_i, u_i)_{i=1}^{I}$ *such that*

1. *For each player $i$, the strategy space $S_i$ is a nonempty compact metric space.*
2. *For each player $i$, the payoff function $u_i : \prod_{j=1}^{I} S_j \to \mathbb{R}$ is continuous.*

*Then there exists a mixed–strategy Nash equilibrium.*

**Remark.** Theorem 0.0.14 generalizes the finite–game existence result. In a finite game we can identify $S_i$ with a finite subset of $\mathbb{R}^k$, hence with a compact metric space, and the usual expected–payoff functions are continuous. Therefore every finite game satisfies the hypotheses and thus has at least one mixed–strategy Nash equilibrium.

### *Pure–strategy Nash equilibria in infinite games*

Suppose now that we are interested in the existence of **pure–strategy** Nash equilibria. We will assume that each player $i$ has a (possibly infinite) pure–strategy set $S_i$ satisfying

$$S_i \subset \mathbb{R}^k \quad \text{for some fixed } k$$

and that $S_i$ is a nonempty, compact, convex subset of $\mathbb{R}^k$. Hence each $S_i$ is uncountable.

**Definition** (Best–response correspondence)**.** Given a game $G = (S_i, u_i)_{i=1}^{I}$, the *best–response correspondence* of player $i$ is

$$BR_i(s_{-i}) \;:=\; \arg\max_{s_i \in S_i} u_i(s_i, s_{-i}) \subseteq S_i \qquad \text{for each } s_{-i} \in \prod_{j \neq i} S_j.$$

The global best–response correspondence is

$$BR(s_1, \ldots, s_I) \;:=\; \prod_{i=1}^{I} BR_i(s_{-i}) \subseteq \prod_{i=1}^{I} S_i.$$

**Definition** (Convex–valued correspondence)**.** Let $X \subset \mathbb{R}^k$ and $F : X \rightrightarrows X$ be a correspondence. We say that $F$ is *convex–valued* if for every $x \in X$ the set $F(x)$ is a convex subset of $X$.

**Definition** (Quasi–concavity)**.** Let $X \subset \mathbb{R}^k$ be convex and $f : X \to \mathbb{R}$. The function $f$ is *quasi–concave* if for every real number $\alpha$ the upper contour set

$$\{x \in X : f(x) \geq \alpha\}$$

is convex. In one dimension this means that all upper contour sets are intervals (possibly empty).

**Remark.** If $u_i(\cdot, s_{-i})$ is continuous and quasi–concave on $S_i$ for every fixed $s_{-i}$, then for each $s_{-i}$ the argmax set $BR_i(s_{-i}) = \arg\max_{s_i \in S_i} u_i(s_i, s_{-i})$ is nonempty, compact and convex. Graphically, $u_i(\cdot, s_{-i})$ has a *single peak or plateau*, so the set of best responses is either a point or a closed interval; see Figure 4.

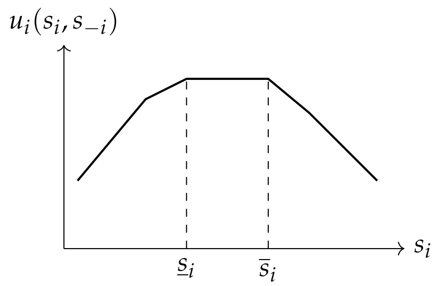


Quasi–concave payoff and convex best–response set $BR_i(s_{-i}) = [\underline{s}_i, \overline{s}_i]$.

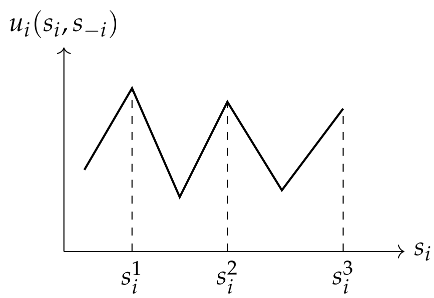


Non–quasi–concave payoff: $BR_i(s_{-i}) = \{s_i^1, s_i^2, s_i^3\}$ is not convex.

With this notation we can state an existence result for pure–strategy equilibrium.

**Theorem 0.0.15** (Pure–strategy Nash equilibrium in infinite games)**.** *Fix a game $G = (S_i, u_i)_{i=1}^I$ such that*

1. *For each $i$, $S_i$ is a nonempty, compact, convex subset of $\mathbb{R}^k$ for some fixed $k$.*
2. *For each $i$, $u_i : \prod_{j=1}^I S_j \to \mathbb{R}$ is continuous.*
3. *For each $i$ and every fixed $s_{-i}$, the function $u_i(\cdot, s_{-i}) : S_i \to \mathbb{R}$ is quasi–concave.*

*Then there exists a pure–strategy Nash equilibrium.*

*Idea of the proof.* Under the above assumptions, the best–response correspondence of each player $BR_i$ is

- nonempty–valued (Weierstrass theorem: continuous function on a compact set attains a maximum),
- convex–valued (quasi–concavity of $u_i(\cdot, s_{-i})$),
- and has a closed graph (by continuity of $u_i$).

Hence the product correspondence $BR : \prod_{j=1}^I S_j \rightrightarrows \prod_{j=1}^I S_j$ satisfies the conditions of Kakutani's fixed–point theorem, and therefore it has a fixed point $s^* \in BR(s^*)$. By definition of $BR$, $s^*$ is a pure–strategy Nash equilibrium. □

### *Topological assumptions (for reference)*

**Definition** (Metric space and compactness)**.** A *metric space* is a pair $(X, d)$ where $X$ is a set and $d : X \times X \to \mathbb{R}_+$ satisfies the usual axioms of a distance. A subset $K \subset X$ is *compact* if every open cover of $K$ admits a finite subcover. In $\mathbb{R}^k$ this is equivalent to $K$ being closed and bounded.

**Definition** (Compact convex subset of $\mathbb{R}^k$)**.** A set $C \subset \mathbb{R}^k$ is a *compact convex subset* of $\mathbb{R}^k$ if:

1. $C$ is convex: for any $x, y \in C$ and $\lambda \in [0, 1]$, $\lambda x + (1 - \lambda)y \in C$;
2. $C$ is compact as a subset of $\mathbb{R}^k$ (equivalently, closed and bounded).

## *Rationalizability*

Up to now, our approach has followed a common structure:

1. We start with a *description of the game*.
2. We then obtain predictions by imposing a *solution concept*.
3. The canonical solution concept studied so far has been Nash equilibrium.

A Nash equilibrium is a solution concept that presumes that players have *correct beliefs* about how other players play. In particular, each player best–responds to beliefs that are exactly consistent with the strategies actually chosen by others.

In many environments, however, this assumption may be too strong. Players may be uncertain about others' strategies, even if they are perfectly rational. This motivates the study of alternative solution concepts that allow for strategic uncertainty.

### *An illustrative example*

Consider the following two games.

Game (1a)

| | L | R |
|---|---|---|
| U | $(-1, 1)$ | $(0, 0)$ |
| D | $(0, 0)$ | $(1, 3)$ |

Game (1b)

| | L | R |
|---|---|---|
| U | $(3, 1)$ | $(0, 0)$ |
| D | $(2, 0)$ | $(1, 3)$ |

Previously, we argued that in Game (1a), player 1's optimal action depends on their belief about player 2's behavior. If player 1 believes that player 2 is rational, then player 2 will play $R$, and therefore player 1 will optimally choose $D$.

In contrast, in Game (1b), player 1 is rationally justified in choosing either $U$ or $D$, depending on their belief about player 2. There is no action of player 2 that is strictly dominant, and therefore player 1 cannot deduce player 2's behavior solely from rationality considerations.

**Key distinction.** Nash equilibrium assumes that player 1 has *correct beliefs* about player 2's strategy. Rationalizability relaxes this requirement and only imposes that beliefs are *consistent with rationality*.

This observation highlights a limitation of Nash equilibrium as a predictive tool: it imposes strong epistemic requirements on players' beliefs. In particular, Nash equilibrium requires players to correctly anticipate how often others play each strategy.

### *From Nash equilibrium to rationalizability*

Rationalizability is a solution concept designed to weaken these epistemic assumptions. Rather than requiring correct beliefs about others' strategies, it only requires that players' strategies are optimal for *some* beliefs about others, as long as those beliefs assign probability one to rational behavior.

Formally, rationalizability allows players to hold beliefs that are not correct, as long as they are consistent with the assumption that all players are rational.

- Nash equilibrium imposes both rationality and correct beliefs.
- Rationalizability imposes rationality, but allows for strategic uncertainty.

As a result, rationalizability typically yields a larger set of predicted strategies than Nash equilibrium. Its predictions are therefore weaker, but also more robust to misspecification of beliefs.

In what follows, we will formalize rationalizability and study how it can be characterized through iterated elimination of strategies that are never optimal for any belief consistent with rationality.

# *Class 9: Beliefs, strategic uncertainty, and rationalizability*

## *From Nash equilibrium to strategic uncertainty*

Up to now, our approach to games has followed a familiar structure:

1. Start with a description of the game.
2. Impose a solution concept to obtain predictions.
3. The canonical solution concept has been Nash equilibrium.

A Nash equilibrium is a solution concept that combines two strong assumptions:

- players are *rational*, and
- players hold *correct beliefs* about how others play.

**Key shift.** Nash equilibrium takes strategies as primitives and derives beliefs. Rationalizability takes beliefs as primitives and studies optimal responses to them.

That is, in equilibrium each player best–responds to beliefs that coincide exactly with the strategies actually chosen by the other players.

In many environments, this epistemic requirement may be too strong. Players may be perfectly rational, yet uncertain about others' strategies. This motivates the study of solution concepts that allow for *strategic uncertainty*.

## *Beliefs as primitive objects*

To model strategic uncertainty, we introduce beliefs explicitly.

**Definition** (Belief)**.** In a finite normal–form game, a *belief* of player $i$ is a probability distribution over the pure–strategy profiles of the other players:

$$\mu_i \in \Delta(S_{-i}).$$

A belief $\mu_i$ represents player $i$'s subjective assessment of how the other players might behave. Importantly, beliefs are *subjective* objects and need not coincide with actual play.

**Key distinction.** A mixed strategy $\sigma_i \in \Delta(S_i)$ describes how player $i$ *randomizes over their own actions*. A belief $\mu_i \in \Delta(S_{-i})$ describes how player $i$ *assigns probabilities to others' actions*. Mixed strategies are *choices*; beliefs are *epistemic*.

## *Beliefs versus mixed–strategy profiles*

Conceptually, a belief $\mu_i$ is different from a mixed–strategy profile $\sigma_{-i} = (\sigma_j)_{j \neq i}$.

- A mixed–strategy profile describes how players actually randomize.
- A belief describes what player $i$ thinks others will do.

Formally, every mixed–strategy profile induces a belief, but the converse is not true.

### *Beliefs induced by mixed strategies*

Given a mixed–strategy profile $\sigma_{-i}$, player $i$ can construct a belief $\mu_i^\sigma$ over $S_{-i}$ by assuming independence:

$$\mu_i^\sigma(s_{-i}) \;=\; \prod_{j \neq i} \sigma_j(s_j).$$

Thus, mixed strategies impose a strong restriction on beliefs: they must be *product measures*. This independence assumption plays no role in rationalizability.

## *An example: beliefs not induced by mixed strategies*

The distinction becomes clear in the following example.

Consider a three–player game where players 2 and 3 each have two actions:

$$S_2 = S_3 = \{L, R\}.$$

Hence

$$S_{-1} = \{(L, L), (L, R), (R, L), (R, R)\}.$$

Define a belief of player 1 by

$$\mu_1(L, L) = \tfrac{1}{2}, \qquad \mu_1(R, R) = \tfrac{1}{2},$$

and

$$\mu_1(L, R) = \mu_1(R, L) = 0.$$

*Interpretation.* Player 1 believes that players 2 and 3 always coordinate: either both play $L$ or both play $R$, with equal probability.

### *Why this belief cannot come from mixed strategies*

Suppose that $\mu_1$ were induced by some mixed–strategy profile $(\sigma_2, \sigma_3)$. Then independence would require

$$\mu_1(L, R) = \sigma_2(L)\sigma_3(R), \qquad \mu_1(R, L) = \sigma_2(R)\sigma_3(L).$$

Since both probabilities are zero, at least one term in each product must be zero. But this makes it impossible to assign positive probability to both $(L, L)$ and $(R, R)$ simultaneously.

Therefore, no mixed–strategy profile can generate $\mu_1$. The belief exhibits *correlation* between players' actions.

**Crucial insight.** Mixed strategies impose independence. Beliefs in rationalizability may involve correlation.

## *Best responses vs. justifiable strategies*

We now make precise the distinction between being a *best response* to a given belief and being a *justifiable* strategy. The second notion will be the building block of rationalizability.

### *Best response under a belief and justifiable strategies*

Fix a player $i$ and let $S_{-i} = \prod_{j \neq i} S_j$ denote the set of pure strategy profiles of the other players. A *belief* of player $i$ about the opponents' pure strategies is an element

$$\mu_i \in \Delta(S_{-i}).$$

Often we are interested in beliefs that put probability one on some subset $A_{-i} \subseteq S_{-i}$ which collects the strategies of the opponents that player *$i$ considers possible* (for instance, those that have not been ruled out by previous rounds of reasoning about rationality).

**Definition** (Best response under a belief given $Q_i$ and $Q_{-i}$)**.** Fix nonempty sets $Q_i \subseteq S_i$ and $Q_{-i} \subseteq S_{-i}$, and a belief $\mu_i \in \Delta(S_{-i})$ such that $\mu_i(Q_{-i}) = 1$. A strategy $s_i \in Q_i$ is a *best response under $\mu_i$ given $Q_i$ and $Q_{-i}$* if

$$\int_{S_{-i}} u_i(s_i, s_{-i})\, d\mu_i(s_{-i}) \;\geq\; \int_{S_{-i}} u_i(r_i, s_{-i})\, d\mu_i(s_{-i}) \quad \text{for all } r_i \in Q_i.$$

**Definition** (Justifiable strategy)**.** Fix nonempty sets $Q_i \subseteq S_i$ and $Q_{-i} \subseteq S_{-i}$. A strategy $s_i \in Q_i$ is *justifiable given $Q_i \times Q_{-i}$* if there exists a belief $\mu_i \in \Delta(S_{-i})$ such that:

1. $\mu_i(Q_{-i}) = 1$, and
2. $s_i$ is a best response under $\mu_i$ given $Q_i$ and $Q_{-i}$.

We say that $s_i$ is *justifiable* if it is justifiable given $S_i \times S_{-i}$.

Intuitively, "best response under $\mu_i$" is purely an optimization statement: *given* a particular belief (which could be arbitrarily unreasonable), $s_i$ maximizes expected utility. Justifiability adds an epistemic restriction: the belief used to justify $s_i$ must be concentrated on opponents' strategies that we are prepared to treat as rational candidates themselves (those in $A_{-i}$). Rationalizability will then require mutual consistency: the strategies in $A_{-i}$ must, in turn, be justifiable for the other players.

**Best response vs. justifiable.** A best response is defined *relative to a fixed belief* and a *fixed set of admissible deviations*. Justifiability is an *existence statement*: a strategy is justifiable if there exists some belief concentrated on opponents' strategies deemed rational and under which the strategy is optimal relative to the admissible deviations. Rationalizability iterates this restriction across players and belief levels.

*An example: best response but not justifiable*

The difference becomes clear in a simple $2 \times 2$ game. Consider the following payoff matrix, where the first entry in each cell is the payoff of player 1 and the second entry is the payoff of player 2:

| | L | R |
|---|---|---|
| U | (2,0) | (0,1) |
| D | (1,0) | (1,1) |

For player 2, action $R$ *strictly dominates* $L$, since $u_2(\cdot, R) = 1 > 0 = u_2(\cdot, L)$ in both rows. If we impose even this minimal notion of rationality on player 2, then player 1 should regard $L$ as impossible. Formally, a "rationality–based" set of possible strategies for player 2 is

$$A_{-1} = \{R\} \subseteq \{L, R\}.$$

Now examine player 1's incentives. Let $p = \mu_1(L)$ be the probability that player 1 assigns to action $L$ under some belief $\mu_1 \in \Delta(\{L, R\})$.

- Expected payoff from $U$: $EU_1(U) = 2p + 0 \cdot (1-p) = 2p$.
- Expected payoff from $D$: $EU_1(D) = 1 \cdot p + 1 \cdot (1-p) = 1$.

If $p > 1/2$, then $EU_1(U) > 1 = EU_1(D)$ and $U$ is a best response to $\mu_1$. Thus there exist beliefs–for example $\mu_1(L) = 0.8$ and $\mu_1(R) = 0.2$–under which $U$ is a best response.

However, such beliefs are *not* compatible with our rationality restriction on player 2. If we maintain that 2 never plays the dominated action $L$, the only admissible beliefs for player 1 are those with support on $A_{-1}$, that is, beliefs satisfying $\mu_1(R) = 1$ and $\mu_1(L) = 0$.

Under any belief with $\mu_1(R) = 1$ we have

$$EU_1(U) = 0 \qquad \text{and} \qquad EU_1(D) = 1,$$

so $D$ strictly dominates $U$ in expectation. Therefore:

- $U$ is a *best response* to some beliefs (those that put sufficiently high probability on the irrational action $L$ of player 2),
- but $U$ is *not justifiable* once we restrict attention to beliefs supported on rational strategies of the opponent; it can never be optimal under any $\mu_1$ with $\mu_1(A_{-1}) = 1$.

This example shows that "best response" and "justifiable" are genuinely different notions. Best responses are defined relative to a given belief, whereas justifiable strategies require the existence of a belief with support only on strategies of others that are themselves rational candidates. In the next step, rationalizability will iterate this restriction across all players to construct the set of strategies that survive common belief in rationality.

### *Inductive definition of rationalizability*

We now formalize rationalizability through an inductive construction. Fix a finite normal–form game

$$G = \big(N, (S_i)_{i\in N}, (u_i)_{i\in N}\big).$$

For each player $i$, recall that a belief is a probability distribution $\mu_i \in \Delta(S_{-i})$ and that a pure strategy $s_i \in S_i$ is *justifiable given* $Q_i \times Q_{-i} \subseteq S_i \times S_{-i}$ if there exists a belief $\mu_i \in \Delta(Q_{-i})$ such that

1. $s_i$ is a best response to $\mu_i$ given $Q_i$, and
2. $\mu_i(Q_{-i}) = 1$ (i.e. the belief assigns probability one to $Q_{-i}$).

In particular, $s_i$ is justifiable given $S_i \times Q_{-i}$ if it is a best response to some belief that puts probability one on $Q_{-i}$.

*Inductive construction.* For each $i$ we define a decreasing sequence of subsets of $S_i$.

- **Step 0.** Set
$$R_i^0 := S_i \qquad \text{for all } i \in N.$$
Initially, every strategy is admissible.
- **Step** $m+1$**.** Given $(R_j^m)_{j\in N}$, define
$$R_{-i}^m := \prod_{j\neq i} R_j^m \qquad \text{and} \qquad R_i^{m+1} := \big\{\, s_i \in S_i : s_i \text{ is justifiable given } S_i \times R_{-i}^m \big\}.$$

Thus $R_i^{m+1}$ collects those strategies of player $i$ that are best responses to some belief that puts probability one on the current surviving set $R_{-i}^m$ of opponents' strategies.

$R_i^1$ corresponds to rationality, $R_i^2$ to rationality and first–order belief in others' rationality, and in general $R_i^m$ to rationality and $(m-1)$th–order belief in rationality.

Finally, the set of *rationalizable strategies* for player $i$ is defined as the intersection of all surviving sets:

$$R_i^\infty := \bigcap_{m=0}^{\infty} R_i^m.$$

A strategy $s_i \in S_i$ is *rationalizable* if $s_i \in R_i^\infty$.

*Monotonicity: later rounds cannot add strategies.* The first basic property is that the inductive procedure is *monotone*: once a strategy is eliminated, it never reappears.

**Lemma 0.0.16.** *For every player $i$ and every $m \geq 0$,*

$$R_i^{m+1} \subseteq R_i^m.$$

*Proof.* Fix $i$ and $m \geq 0$, and let $s_i \in R_i^{m+1}$. By definition of $R_i^{m+1}$ there exists a belief $\mu_i \in \Delta(S_{-i})$ such that

1. $s_i$ is a best response to $\mu_i$ given $S_i$, and
2. $\mu_i(R_{-i}^m) = 1$.

By the inductive construction, $R_j^m \subseteq R_j^{m-1}$ for all $j$ and therefore

$$R_{-i}^m = \prod_{j \neq i} R_j^m \subseteq \prod_{j \neq i} R_j^{m-1} = R_{-i}^{m-1}.$$

Since $\mu_i$ assigns probability one to $R_{-i}^m$, it also assigns probability one to $R_{-i}^{m-1}$:

$$\mu_i(R_{-i}^{m-1}) \; = \; 1.$$

Hence the same belief $\mu_i$ witnesses that $s_i$ is justifiable given $S_i \times R_{-i}^{m-1}$, so by definition $s_i \in R_i^m$. Thus $R_i^{m+1} \subseteq R_i^m$. □

Lemma 0.0.16 formalizes the idea that the procedure only *eliminates* strategies: higher levels of rationality impose stronger restrictions on admissible beliefs and therefore can only shrink the set of surviving strategies.

*Checking only surviving strategies.* The next remark explains why, in finite games, one can restrict attention to strategies that have not yet been eliminated when moving to the next round.

**Remark** (Working inside the surviving set)**.** Because of Lemma 0.0.16, for every $m$ we have $R_i^m \subseteq S_i$. In a finite game, this allows us to equivalently define

$$R_i^{m+1} = \left\{ s_i \in R_i^m : s_i \text{ is justifiable given } R_i^m \times R_{-i}^m \right\}.$$

Intuitively, once a strategy of player $i$ has been eliminated at some earlier round $k \leq m$, it is never a best response to any belief that puts probability one on $R_{-i}^m$, so we do not need to reconsider it. Operationally, when computing $R_i^{m+1}$ we only check strategies that have survived up to round $m$.

This observation connects the epistemic definition of rationalizability with a computational procedure based on iterated elimination: each round discards strategies that are no longer justifiable given the surviving strategies of all players.

**Lemma 0.0.17** (Finite games)**.** *In a finite game, for every player $i$ and every $m \geq 0$, a strategy $s_i \in S_i$ is justifiable given $S_i \times R_{-i}^m$ if and only if $s_i$ is justifiable given $R_i^m \times R_{-i}^m$.*

*Idea of the proof.* Because the game is finite, expected utilities are linear in beliefs and the set $R_i^m$ contains exactly those strategies that have survived the previous $m$ rounds of the procedure. If a strategy $s_i \notin R_i^m$ were ever strictly worse than some surviving strategy $s_i' \in R_i^m$ for all beliefs with support in $R_{-i}^m$, then $s_i$ cannot be a best response to any such belief and therefore cannot be justifiable given $S_i \times R_{-i}^m$. Conversely, any strategy $s_i \in R_i^m$ that is justifiable given $S_i \times R_{-i}^m$ remains justifiable when we restrict the own–strategy space to $R_i^m$, because we are only removing strategies that are never optimal under beliefs with support in $R_{-i}^m$. Thus the two formulations are equivalent. □

In practical terms, Lemma 0.0.17 means that in a finite game the inductive construction of the sets $(R_i^m)_{m \geq 0}$ can be implemented as a genuine *iterated elimination procedure*: at each step we only need to examine the strategies that have not yet been eliminated.

# Class 10: Iterative rationalizability in finite games

In this class we specialize the general notion of justifiability to construct, step by step, the sets of $m$–rationalizable strategies in a *finite* game, and we work through a non–trivial example.

## Iterative construction in finite games

Recall the definition of justifiability: for nonempty $Q_i \subseteq S_i$ and $Q_{-i} \subseteq S_{-i}$, a strategy $s_i \in S_i$ is *justifiable given* $Q_i \times Q_{-i}$ if there exists a belief $\mu_i \in \Delta(S_{-i})$ such that

1. $s_i$ is a best response under $\mu_i$ given $Q_i$, i.e.
$$\int_{S_{-i}} u_i(s_i, s_{-i})\, d\mu_i(s_{-i}) \;\geq\; \int_{S_{-i}} u_i(r_i, s_{-i})\, d\mu_i(s_{-i}) \quad \text{for all } r_i \in Q_i,$$
2. $\mu_i(Q_{-i}) = 1$ (the belief is concentrated on $Q_{-i}$).

We now use this notion to define $m$–rationalizability.

**Definition** ($m$–rationalizable strategies)**.** For each player $i$, let
$$R_i^0 \;=\; S_i.$$
Having defined the sets $(R_j^m)_{j \in N}$ for some $m \geq 0$, define
$$R_i^{m+1} \;=\; \{s_i \in S_i : s_i \text{ is justifiable given } S_i \times R_{-i}^m\}, \qquad R_{-i}^m \;=\; \prod_{j \neq i} R_j^m.$$
The set $R_i^{m+1}$ collects the strategies of player $i$ that can be justified by a belief $\mu_i$ putting probability one on opponents' $m$–rationalizable strategies $R_{-i}^m$.

In a *finite* game, it is convenient to exploit the fact that once a strategy of player $i$ has been eliminated in some round $m$, we never need to use it again as a potential deviation.

**Remark** (Finite games)**.** In a finite game we can rewrite the previous definition as
$$R_i^{m+1} \;=\; \{s_i \in R_i^m : s_i \text{ is justifiable given } R_i^m \times R_{-i}^m\}.$$

Intuitively, when checking whether $s_i$ survives to round $m+1$, it suffices to compare it with deviations that themselves survived round $m$: deviations already ruled out as non–rationalizable never need to be considered again.

In finite games, the iterative process takes place inside the shrinking sets $R_i^m$. At each step we only check deviations in $R_i^m$, not in the whole $S_i$.

As before, the set of *rationalizable* strategies for player $i$ is

$$R_i^\infty = \bigcap_{m \geq 1} R_i^m.$$

## *Example 3: computing rationalizable strategies*

Consider the following two–player game. Player 1 has four pure strategies $S_1 = \{x_1, y_1, z_1, q_1\}$ and player 2 has $S_2 = \{x_2, y_2, z_2, q_2\}$. Payoffs are:

| | $x_2$ | $y_2$ | $z_2$ | $q_2$ |
|---|---|---|---|---|
| $x_1$ | $(3,0)$ | $(0,3)$ | $(3,1)$ | $(0,1)$ |
| $y_1$ | $(0,3)$ | $(3,0)$ | $(0,-1)$ | $(3,0)$ |
| $z_1$ | $(1,0)$ | $(-1,3)$ | $(-1,-1)$ | $(4,1)$ |
| $q_1$ | $(1,3)$ | $(0,0)$ | $(1,2)$ | $(0,0)$ |

We now compute $R_1^m$ and $R_2^m$.

### *First round*

We start from $R_1^0 = S_1$ and $R_2^0 = S_2$.

*Player 1.* We claim that $q_1$ is *never* a best response to any belief $\mu_1$ over $S_2$, hence $q_1 \notin R_1^1$.

Assume for contradiction that there exists a belief $\mu_1 \in \Delta(S_2)$ such that $q_1$ is a best response. Then in particular its expected payoff must be at least as large as that of $x_1$:

$$\sum_{s_2 \in S_2} u_1(q_1, s_2)\, \mu_1(s_2) \;\geq\; \sum_{s_2 \in S_2} u_1(x_1, s_2)\, \mu_1(s_2).$$

Using the entries in the table, this inequality becomes

$$\mu_1(x_2) + \mu_1(z_2) \;\geq\; 3\mu_1(x_2) + 3\mu_1(z_2),$$

which is only possible if

$$\mu_1(x_2) + \mu_1(z_2) = 0.$$

Thus any belief that could make $q_1$ optimal must put probability one on $\{y_2, q_2\}$.

But conditional on $\mu_1(x_2) = \mu_1(z_2) = 0$, the expected payoff of $y_1$ is

$$\sum_{s_2 \in S_2} u_1(y_1, s_2)\, \mu_1(s_2) = 3\mu_1(y_2) + 3\mu_1(q_2),$$

whereas the expected payoff of $q_1$ is

$$\sum_{s_2 \in S_2} u_1(q_1, s_2)\, \mu_1(s_2) = 0 \cdot \mu_1(y_2) + 0 \cdot \mu_1(q_2) = 0.$$

Since $\mu_1(y_2) + \mu_1(q_2) = 1$, we have $3\mu_1(y_2) + 3\mu_1(q_2) > 0$, so $y_1$ strictly dominates $q_1$ under this belief, contradicting that $q_1$ is a best response.

Hence there is no belief that makes $q_1$ a best response; $q_1$ is not justifiable and is eliminated in the first round. We obtain

$$R_1^1 = \{x_1, y_1, z_1\}.$$

*Player 2.* A symmetric argument shows that $q_2$ is never a best response for player 2: whatever belief $\mu_2 \in \Delta(S_1)$ we consider, some other strategy yields a higher expected payoff. Therefore

$$R_2^1 = \{x_2, y_2, z_2\}.$$

### *Second round*

We now work with $R_1^1$ and $R_2^1$ and ask which strategies are justifiable given $R_1^1 \times R_2^1$.

Observe that $z_1$ is *strictly dominated* by a mixed strategy over $\{x_1, y_1\}$ when player 2 is restricted to $\{x_2, y_2, z_2\}$. Consequently, there is no belief $\mu_1$ with support in $R_2^1$ under which $z_1$ is a best response. Hence $z_1$ is not justifiable given $R_1^1 \times R_2^1$ and is eliminated in the second round.

By symmetry, $z_2$ is eliminated for player 2. Thus

$$R_1^2 = \{x_1, y_1\}, \qquad R_2^2 = \{x_2, y_2\}.$$

A strictly dominated pure strategy can never be a best response under any belief. Therefore it is never justifiable, and will be eliminated at some finite step of the rationalizability procedure.

### *Third round and limit set*

In the reduced $2 \times 2$ subgame with strategies $\{x_1, y_1\} \times \{x_2, y_2\}$, all four strategies are best responses to some belief concentrated on the opponents' remaining strategies. Hence no further eliminations occur:

$$R_1^m = R_1^2 = \{x_1, y_1\}, \qquad R_2^m = R_2^2 = \{x_2, y_2\} \quad \text{for all } m \geq 2.$$

Therefore the sets of rationalizable strategies are

$$R_1^\infty = \{x_1, y_1\}, \qquad R_2^\infty = \{x_2, y_2\},$$

and the set of rationalizable strategy profiles is

$$R_1^\infty \times R_2^\infty \;=\; \{x_1, y_1\} \times \{x_2, y_2\}.$$

## *Rationalizability, existence and dominance by mixed strategies*

Recall the iterative construction

$$R_i^0 \;=\; S_i, \qquad R_i^{m+1} \;=\; \left\{\, s_i \in S_i : s_i \text{ is justifiable given } S_i \times R_{-i}^m \,\right\},$$

and $R_i^\infty := \bigcap_{m \geq 1} R_i^m$, the set of rationalizable strategies of player $i$.

### *Stability in finite games and existence in infinite games*

**Proposition 0.0.18** (Finite games). *For a finite game there exists an integer $M \geq 1$ such that*

$$R_i^m \;=\; R_i^\infty \neq \varnothing \qquad \text{for all } m \geq M \text{ and all } i.$$

*Idea.* For each player $i$, the sequence $R_i^0 \supseteq R_i^1 \supseteq R_i^2 \supseteq \cdots$ is a decreasing sequence of subsets of the finite set $S_i$. Hence membership of each pure strategy can change only finitely many times, so the sequence must stabilize after finitely many steps. Since we know from the definition that $R_i^\infty \neq \varnothing$ in finite games, the stabilized set must coincide with $R_i^\infty$. □

We now move to games with possibly infinite strategy sets.

**Proposition 0.0.19** (Infinite games). *Fix a game in which, for each player $i = 1, \ldots, I$,*

- *the strategy space $S_i$ is a nonempty compact metric space;*
- *the payoff function $u_i : \prod_{j=1}^I S_j \to \mathbb{R}$ is continuous.*

*Then the set*

$$R^\infty \;:=\; \prod_{i=1}^I R_i^\infty$$

*is nonempty, where $R_i^\infty = \bigcap_{m \geq 1} R_i^m$ for each $i$.*

*Idea.* Under the assumptions above one can show that, for each $i$ and each $m$,

- $R_i^m$ is nonempty,
- $R_i^m$ is closed in $S_i$,
- the sequence is decreasing: $R_i^{m+1} \subseteq R_i^m$.

Thus $\{R_i^m\}_{m \geq 1}$ is a nested family of nonempty, closed subsets of the compact metric space $S_i$. By the Cantor intersection property, the intersection $R_i^\infty = \bigcap_{m \geq 1} R_i^m$ is nonempty. Taking products across players yields $R^\infty \neq \varnothing$. □

**Remark.** No entendí. The closedness of the sets $R_i^m$ is essential. For instance, if we took $S_i = [0, 1]$ and the nested sequence $A^m = (0, 1/m)$, then each $A^m$ is nonempty and $A^{m+1} \subset A^m$, but $\bigcap_{m \geq 1} A^m = \varnothing$. This cannot happen with a nested family of nonempty *closed* subsets of a compact metric space.

*Example 4: pure vs. mixed dominance*

We now revisit rationalizability and dominance in two closely related $3 \times 2$ games. In both examples we focus on player 1 and treat player 2's payoffs as irrelevant for the argument, so we only display player 1's payoffs.

*Example 4a*

| | L | R |
|---|---|---|
| U | 4 | 0 |
| M | 0 | 4 |
| D | 3 | 3 |

For each pure strategy of player 2, player 1's pure best responses are:

$$BR_1(L) = \{U\}, \qquad BR_1(R) = \{M\}.$$

Hence $D$ is *never* a pure best response:

$$D \notin BR_1(L) \cup BR_1(R).$$

If we iteratively delete pure strategies that are never best responses to any pure strategy of the opponent, then in this example we obtain

$$R_1^1 = \{U, M\}.$$

However, $D$ is a best response to some *belief* over player 2's strategies. If player 1 believes that player 2 plays $L$ and $R$ each with probability $1/2$, then the expected payoffs are

$$u_1(U, \mu_2) = 4 \cdot \tfrac{1}{2} + 0 \cdot \tfrac{1}{2} = 2, \quad u_1(M, \mu_2) = 0 \cdot \tfrac{1}{2} + 4 \cdot \tfrac{1}{2} = 2, \quad u_1(D, \mu_2) = 3 \cdot \tfrac{1}{2} + 3 \cdot \tfrac{1}{2} = 3.$$

Thus $D$ is strictly better than $U$ and $M$ under this belief, and so $D$ is a best response to the belief $\mu_2(L) = \mu_2(R) = \frac{1}{2}$. In particular, $D$ is *justifiable* when we allow beliefs over opponents' strategies, even though it is never a pure best response.

This illustrates the tension: pure best-response elimination removes $D$, but the belief-based notion of justifiability (and hence rationalizability) keeps it.

*Example 4b* Consider now the slightly modified game:

| | L | R |
|---|---|---|
| U | 4 | 0 |
| M | 0 | 4 |
| D | 2 | 1 |

Again, $D$ is never a pure best response. We now ask whether $D$ can be a best response to some belief $\mu_2$ over $\{L, R\}$. Let $\mu_2(L) = q$ and $\mu_2(R) = 1 - q$. Then the expected payoff of $D$ is

$$u_1(D, \mu_2) = 2q + 1 \cdot (1 - q) = 1 + q.$$

The expected payoff of $M$ is

$$u_1(M, \mu_2) = 4(1 - q).$$

For $D$ to be (weakly) better than $M$, we would need

$$1 + q \ \geq \ 4(1 - q) \quad \Longleftrightarrow \quad 1 + q \geq 4 - 4q \quad \Longleftrightarrow \quad 1 \geq 3q \quad \Longleftrightarrow \quad q \leq \tfrac{1}{3},$$

or equivalently $\mu_2(R) \geq \frac{2}{3}$. But then,

$$u_1(M, \mu_2) = 4(1 - q) \ \geq \ 4 \cdot \tfrac{2}{3} = \tfrac{8}{3} \ > \ 2 \ \geq \ u_1(D, \mu_2),$$

so for such beliefs $M$ strictly outperforms $D$.

The previous inequalities suggest looking for a *mixed* strategy of player 1 that strictly dominates $D$. Consider a mixture between $U$ and $M$,

$$\sigma_1(p) = pU + (1 - p)M, \qquad p \in [0, 1].$$

Against $L$ this mixture yields payoff $4p$; against $R$ it yields $4(1 - p)$. In Example 4b, $D$ yields payoffs 2 against $L$ and 1 against $R$. We can choose $p$ so that

$$4p > 2 \quad \text{and} \quad 4(1 - p) > 1.$$

For instance, $p = 0.6$ works:

$$4p = 2.4 > 2, \qquad 4(1 - p) = 1.6 > 1.$$

Hence the mixed strategy $\sigma_1(0.6)$ *strictly dominates* $D$ in Example 4b: it does strictly better than $D$ against both $L$ and $R$.

By contrast, in Example 4a the same calculation would require

$$4p > 3 \quad \text{and} \quad 4(1 - p) > 3,$$

which is impossible to satisfy simultaneously. Thus, in Example 4a there is no mixture that always does better than $D$: for every value of $p$ there is some column (either $L$ or $R$) for which the mixture fails to beat $D$.

**Remark.** Example 4b shows a pure strategy $D$ that is never a pure best response, and in fact is strictly dominated by a mixed strategy. Example 4a shows that the same pattern of payoffs need not imply mixed dominance: $D$ is never a pure best response there either, but no mixed strategy strictly dominates it. This will be crucial when we connect justifiability with *mixed* dominance.

*Dominance by mixed strategies and justifiability*

We now formalize dominance by mixed strategies on a given restriction of the game.

**Definition** (Dominated and undominated given a restriction)**.** Fix nonempty sets $Q_1, \ldots, Q_I$ with $Q_j \subseteq S_j$ for each player $j$. We say that a strategy $s_i \in S_i$ is *dominated given* $\prod_{j=1}^{I} Q_j$ if

$$s_i \in Q_i \quad \text{and there exists} \quad \sigma_i \in \Delta(Q_i) \quad \text{such that} \quad u_i(\sigma_i, s_{-i}) > u_i(s_i, s_{-i}) \quad \text{for all } s_{-i} \in \prod_{j \neq i} Q_j.$$

That is, some mixed strategy with support in $Q_i$ yields a strictly higher payoff than $s_i$ against every profile of opponents' strategies in the restriction.

We say that $s_i$ is *dominated* (without qualification) if it is dominated given the full product $\prod_{j=1}^{I} S_j$.

We say that $s_i$ is *undominated given* $\prod_{j=1}^{I} Q_j$ if it is not dominated given $\prod_{j=1}^{I} Q_j$.

**Proposition 0.0.20.** *Fix a finite game. A strategy $s_i$ is undominated if and only if it is justifiable.*

*Idea.* ("Only if") If $s_i$ is dominated, then there exists a mixed strategy $\sigma_i$ that strictly improves upon $s_i$ against every opponents' profile. Any belief $\mu_i$ with support on the opponents' strategies cannot make $s_i$ a best response, because integrating the strict inequality $u_i(\sigma_i, s_{-i}) > u_i(s_i, s_{-i})$ over $\mu_i$ preserves strict inequality. Hence $s_i$ is not a best response under any belief and so cannot be justifiable.

("If") Conversely, if $s_i$ is not justifiable, then there is no belief with support on opponents' strategies under which $s_i$ is a best response. In a finite game one can show that this failure of best-response optimality can be aggregated into a single mixed strategy $\sigma_i$ that uniformly does better than $s_i$ against every opponents' profile, so $s_i$ is dominated. □

**Remark.** Proposition 0.0.20 gives a clean equivalence in finite games: rationalizability (via justifiable strategies and iterated belief-based reasoning) coincides with the standard game-theoretic notion of eliminating strategies that are dominated by some mixed strategy. In particular, a strategy survives iterated mixed-strategy dominance if and only if it is a candidate for rational play under some coherent system of beliefs about others' rationality.

# *Class 11: From rationalizability to sequential games*

In this class we complete the analysis of rationalizability in finite games by establishing its equivalence with iterated elimination of strategies dominated by mixed strategies. This result provides a belief–free characterization of rationalizable behavior.

We then turn to games with sequential moves. Through a simple entry–deterrence example, we show that the strategic form may conceal the timing and observability of actions, leading Nash equilibrium to rely on non–credible threats. This motivates the introduction of extensive–form games as the proper framework for modeling dynamic interaction.

## *Undominated strategies and justifiability*

Recall Proposition 0.0.20 from last class:

**Proposition 0.0.21** (Finite games)**.** *Fix a finite game. A strategy $s_i \in S_i$ is undominated if and only if it is justifiable.*

In particular, if the game is finite, then $s_i$ is undominated if and only if it is undominated *given the full product* $\prod_{j=1}^{I} S_j$ of strategy sets:

$$s_i \in S_i \text{ is undominated} \iff$$
$$\nexists \sigma_i \in \Delta(S_i) \text{ such that } u_i(\sigma_i, s_{-i}) > u_i(s_i, s_{-i}) \text{ for all } s_{-i} \in \prod_{j \neq i} S_j.$$

In any finite game, for each player $i$, the set of undominated strategies is nonempty.

Indeed, if every strategy of player $i$ were strictly dominated by some mixed strategy, then by finiteness one could aggregate these dominations into a single mixed strategy that strictly dominates all pure strategies in $S_i$, which is impossible. Hence at least one strategy must be undominated.

**Remark.** The equivalence in Proposition 0.0.21 relies heavily on finiteness of the game. In games with infinite strategy spaces, one can construct examples in which a strategy is undominated but fails to be justifiable, or vice versa. This is what Amanda wanted to show in exercise 5 of PS4.

### *Why the equivalence fails in infinite games*

Proposition 0.0.21 relies in an essential way on the finiteness of the strategy sets. In infinite games, the equivalence between undominated strategies and justifiability may fail.

To see why, recall the logic underlying the proof. If a strategy $s_i$ is not justifiable, then for every belief over opponents' strategies it fails to be a best response. In a finite game, this failure can be aggregated into a single mixed strategy $\sigma_i$ that strictly dominates $s_i$ against all opponents' strategy profiles. This aggregation step relies on two key properties:

- the payoff comparisons involve only finitely many opponent strategies,
- the set of mixed strategies $\Delta(S_i)$ is a finite-dimensional, compact, convex set.

In infinite games, both properties may fail. When the opponents' strategy sets are infinite, the condition that $s_i$ is never a best response may hold pointwise—belief by belief—without implying the existence of a *single* mixed strategy that uniformly dominates $s_i$ against all opponents' strategies.

Formally, the separation arguments that allow one to pass from

"no belief makes $s_i$ optimal"

to

"there exists a mixed strategy that strictly dominates $s_i$"

may break down in infinite-dimensional settings. In particular, the relevant sets of payoff vectors need not be compact, and hyperplane separation need not apply.

As a result, in infinite games it is possible for a strategy to be *undominated* while still failing to be justifiable, or vice versa. This is why the clean equivalence between rationalizability and iterated elimination of mixed-strategy dominance is a special feature of finite games.

## *Iterated elimination of dominated strategies*

Motivated by Proposition 0.0.21, we now define an iterative elimination procedure based on dominance by mixed strategies.

For each player $i$ let

$$U_i^0 = S_i.$$

For each $m \geq 0$, define recursively

$$U_i^{m+1} = \{s_i \in U_i^m : s_i \text{ is undominated given } \textstyle\prod_{j=1}^I U_j^m\}.$$

That is, $U_i^{m+1}$ consists of the strategies of player $i$ that are *not strictly dominated by any mixed strategy with support in $U_i^m$* when opponents are restricted to $U_{-i}^m$.

We then define the limit set

$$U_i^\infty = \bigcap_{m \geq 0} U_i^m.$$

It is convenient to collect these across players:

$$U^m = \prod_{i=1}^{I} U_i^m, \qquad U^\infty = \prod_{i=1}^{I} U_i^\infty.$$

We refer to $U^m$ as the set of *m–undominated strategy profiles* and to $U^\infty$ as the set of *iteratively undominated profiles*.

The iterative deletion process takes place inside the shrinking sets $U_i^m$, exactly as in the construction of $R_i^m$ for rationalizable strategies.

## *Equivalence with rationalizability in finite games*

We now relate the sets $(U_i^m)_m$ to the rationalizability sets $(R_i^m)_m$ defined in the previous class.

**Corollary 0.0.22.** *In a finite game, for every player i and every $m \geq 0$,*

$$R_i^m = U_i^m.$$

*Consequently,*

$$R_i^\infty = U_i^\infty \qquad \textit{for all } i.$$

*Idea.* The proof is by induction on $m$.
*Base step.* By construction, $R_i^0 = S_i = U_i^0$ for all $i$.
*Inductive step.* Fix $m \geq 0$ and suppose $R_i^m = U_i^m$ for all $i$. A strategy $s_i \in S_i$ belongs to $R_i^{m+1}$ if and only if it is justifiable given $R_i^m \times R_{-i}^m$. Using the inductive hypothesis, this is equivalent to being justifiable given $U_i^m \times U_{-i}^m$. By Proposition 0.0.21, in a finite game this is equivalent to being undominated given $\prod_{j=1}^{I} U_j^m$, i.e. to belonging to $U_i^{m+1}$. Hence $R_i^{m+1} = U_i^{m+1}$ for all $i$.

Thus, by induction, $R_i^m = U_i^m$ for all $m$ and all $i$, and therefore also $R_i^\infty = U_i^\infty$ for each $i$. □

**Remark.** Corollary 0.0.22 shows that in finite games the belief-based notion of rationalizability coincides with the purely strategic procedure of iterated elimination of strategies dominated by mixed strategies. This provides an operational way to compute rationalizable strategies without explicitly keeping track of beliefs.

## *Example 5 and weak dominance*

We now discuss a simple $2 \times 2$ example that motivates a weaker (not epistemic) notion of dominance.

Consider the following two-player game. Player 1 chooses $U$ or $D$, player 2 chooses $L$ or $R$:

| | $L$ | $R$ |
|---|---|---|
| $U$ | $(2,3)$ | $(1,0)$ |
| $D$ | $(2,0)$ | $(2,3)$ |

For player 1 we have $S_1 = \{U, D\}$ and for player 2 we have $S_2 = \{L, R\}$. Neither $U$ nor $D$ is strictly dominated:

- against $L$, both $U$ and $D$ yield payoff 2;

- against $R$, $D$ yields payoff 2 while $U$ yields payoff 1.

Thus no pure or mixed strategy of player 1 strictly dominates $D$, and likewise no strategy strictly dominates $U$. In the iterative procedure based on strict dominance we keep both strategies:

$$R_1^1 = U_1^1 = \{U, D\}, \qquad R_2^1 = U_2^1 = \{L, R\}.$$

From the point of view of player 1, however, always playing $U$ may feel "silly" if she expects $R$ with substantial probability: she is then indifferent between $U$ and $D$ against $L$, but strictly prefers $D$ to $U$ against $R$.

This suggests a *weak version of dominance*: a strategy that is never strictly better and sometimes strictly worse than another strategy might be discarded by a *cautious* player.

Formally, one can define an iterative procedure based on *weak* dominance. Let

$$W_i^0 = S_i,$$

and, for each $m \geq 0$,

$$W_i^{m+1} = \left\{s_i \in W_i^m : s_i \text{ is not weakly dominated given } \textstyle\prod_{j=1}^{I} W_j^m\right\},$$

and $W_i^\infty = \bigcap_{m \geq 0} W_i^m$.

We refer to $\{W_i^\infty\}_i$ as the result of *iterated weak dominance*.

**Remark.** Iterated weak dominance reflects a different idea than rationalizability. A weakly dominated strategy may still be a best response to some belief, so it can be justifiable. Moreover, iterated weak dominance may eliminate strategies that would survive rationalizability, and the outcome of the procedure can depend on the order in which strategies are removed. For this reason, rationalizability is formally tied to *strict* dominance by mixed strategies, not to weak dominance.

## *Sequential games: motivation*

We now turn to *sequential games*. The key new feature relative to the previous analysis is that players may move at different points in time, and some players may observe earlier actions before choosing their own.

### *Example 1: entry deterrence*

Consider the following example. An *entrant* ($E$) decides whether to *enter* the market or stay *out*. If the entrant stays out, the *incumbent* monopolist ($I$) earns monopoly profits. If the entrant enters, the incumbent can either *fight* the entrant or *accommodate* (entry is costly to fight for both players).

The payoffs are as follows:

- If the entrant stays out:

$$(0, \pi), \qquad \pi > 2.$$

- If the entrant enters and the incumbent accommodates:

$$\left(\frac{\pi}{2}, \frac{\pi}{2}\right).$$

- If the entrant enters and the incumbent fights:

$$(-1, 1).$$

### *Normal form representation*

The strategic form of this game can be written as

| | $F$ | $A$ |
|---|---|---|
| enter | $(-1, 1)$ | $\left(\frac{\pi}{2}, \frac{\pi}{2}\right)$ |
| out | $(0, \pi)$ | $(0, \pi)$ |

It is easy to verify that this game has two pure–strategy Nash equilibria:

$$(\text{enter}, A) \qquad \text{and} \qquad (\text{out}, F).$$

### *What is missing? Timing and observability*

The problem with the strategic–form analysis is not the representation per se, but the solution concept applied to it.

The normal form hides the fact that the incumbent chooses between $F$ and $A$ *after observing* whether entry has occurred. In particular, the profile $(\text{out}, F)$ relies on a threat that is never tested on the equilibrium path.

This creates a tension:

- In the strategic form, $F$ is weakly dominated for the incumbent.
- Yet Nash equilibrium allows outcomes supported by such non–credible threats.

This is why some people informally say that "the strategic form is not enough." The real issue, however, is not the strategic form itself, but the use of Nash equilibrium as the solution concept in environments with sequential moves.

### *Why Nash equilibrium is too weak here*

In the strategic form, the entrant reasons as follows:

> If I enter, the incumbent might fight, so staying out is safer.

But this ignores the incumbent's incentives *after* entry. Once entry has occurred, accommodation strictly dominates fighting for the incumbent:

$$\frac{\pi}{2} > 1.$$

Thus the threat to fight is not credible. Nash equilibrium does not rule out such behavior because it does not impose any restriction on actions that are taken off the equilibrium path.

### *How should we represent the game?*

To make timing explicit, we represent the game in *extensive form*. The game starts with the entrant choosing between `out` and `enter`. If the entrant chooses `out`, the game ends. If the entrant chooses `enter`, the incumbent observes this choice and then chooses between `F` and `A`. The game then ends.

This representation makes clear:

- which player moves at each node,
- which actions are available at each node,
- and which payoffs correspond to each terminal history.

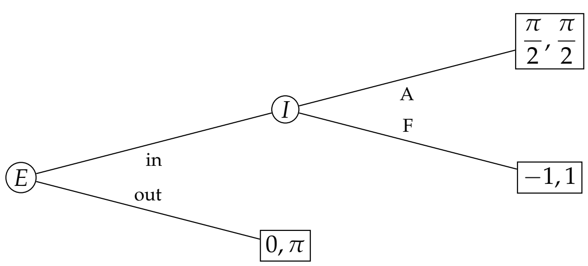


### *Formalization: Osborne–Rubinstein approach*

We adopt the formalization of extensive–form games used by Osborne and Rubinstein.

A key modeling device is the use of a *dummy action*, denoted by $*$, which allows us to represent the game as one in which *all players move at every non–terminal node*. Players who do not have a genuine choice at a node are assigned the dummy move.

Under this convention:

- At the initial node, the incumbent chooses $*$ while the entrant chooses `out` or `enter`.

- At the second node, the entrant chooses $*$ while the incumbent chooses `F` or `A`.

This allows us to describe the extensive form in a way that is compatible with standard definitions of strategies as complete contingent plans.

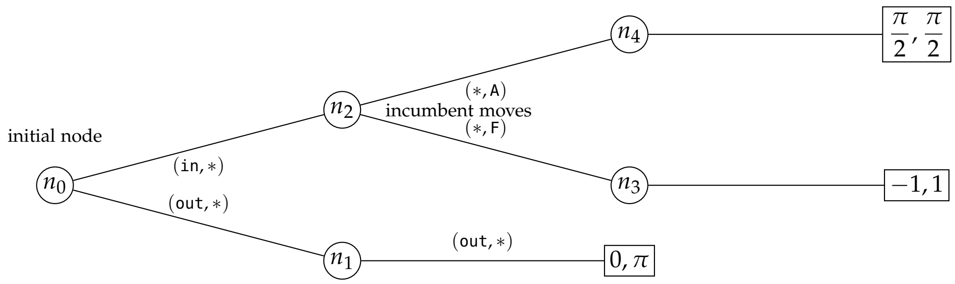


## *Extensive form games*

**Definition** (Extensive form game)**.** We model a sequential game as an *extensive form game*

$$\Gamma = \big(I,\ (A_i)_{i\in I},\ N,\ (A_i(n))_{i\in I,\, n\in N},\ Z,\ (u_i)_{i\in I}\big).$$

We describe each component in turn.

*EF–1: Players.* The set of players is

$$I = \{1, \ldots, I\}.$$

We typically use $i$ to denote a generic player.

*EF–2: Actions.* For each player $i$, $A_i$ is the set of actions that player $i$ may choose *somewhere in the game*. Not all actions in $A_i$ are necessarily available at every node.

*EF–3: Nodes.* Let $N$ denote the set of nodes. A useful way to think about nodes is as *finite sequences of action profiles*. Formally, a node records the history of play up to that point.

The initial node is the empty sequence, denoted by

$$\varnothing.$$

The set $N$ is a (typically strict) subset of the set of all finite sequences of action profiles. In particular:

- $\varnothing \in N$,
- if $n \in N$ and $n'$ extends $n$, then $n'$ represents a later point in the game.

This naturally induces a precedence relation on nodes: if $n$ is a prefix of $n'$, then $n$ occurs before $n'$.

*Tree structure.* The prefix property implies that nodes form a *tree*. In particular, if two action sequences differ at some action, they cannot later reconverge to the same node. That is, the game has no cycles and no merging paths.

*Available actions at a node.* Each non–terminal node $n$ is associated with a nonempty set of actions for each player,

$$A_i(n) \subseteq A_i, \qquad A_i(n) \neq \varnothing.$$

(Allowing $A_i(n)$ to be a singleton accommodates dummy actions.)

The set of action profiles available at node $n$ is

$$A(n) = \prod_{i \in I} A_i(n).$$

*Terminal nodes.* Let $Z \subset N$ denote the set of terminal nodes. At terminal nodes the game ends and payoffs are assigned.

### *Example: entry game revisited*

Consider again the entry game with two players: an entrant $E$ and an incumbent $I$.

*Players.*

$$I = \{E, I\}.$$

*Actions.*

$$A_E = \{\texttt{in}, \texttt{out}, *\}, \qquad A_I = \{\texttt{F}, \texttt{A}, *\},$$

where $*$ denotes a dummy action.

*Nodes.*

- Initial node: $n_0 = \varnothing$.
- After entry decision:

$$n_1 = (\texttt{in}, *), \qquad n_2 = (\texttt{out}, *).$$

- After incumbent response:

$$n_3 = \big((\texttt{in}, *), (*, \texttt{F})\big), \qquad n_4 = \big((\texttt{in}, *), (*, \texttt{A})\big).$$

Here $n_0$ is a prefix of $n_1$ and $n_2$, and $n_1$ is a prefix of $n_3$ and $n_4$, illustrating the tree structure.

*Available actions.* At the initial node,

$$A_E(n_0) = \{\texttt{in}, \texttt{out}\}, \qquad A_I(n_0) = \{*\}.$$

At node $n_1$,

$$A_E(n_1) = \{*\}, \qquad A_I(n_1) = \{\texttt{F}, \texttt{A}\}.$$

*Terminal nodes.*

$$Z = \{n_2, n_3, n_4\}.$$

This example illustrates how nodes, action sets, dummy moves, and the prefix structure jointly encode the timing and observability of actions in a sequential game.

# *Class 12: Information sets*

## *Information sets*

We now complete the definition of an extensive form game by introducing *information sets*.

*EF–4: Information sets for player i.* For each player $i$, let $H_i$ denote the collection of information sets of player $i$.

Each $H_i$ is a partition of the set of non–terminal nodes:

$$H_i \text{ is a partition of } N \setminus Z.$$

**Be careful**: not every node in $N \setminus Z$ belongs to $H_i$. More precisely, $H_i$ is a partition of the set of decision nodes at which player $i$ moves. If $P(n)$ denotes the player who moves at node $n$, then

$$H_i \text{ is a partition of } \{n \in N \setminus Z : P(n) = i\}.$$

Thus, each player partitions only the nodes where he has a decision.

*Idea.* If two nodes $n$ and $n'$ belong to the same information set $h_i \in H_i$, then when player $i$ moves at node $n$, he does not know whether he is at $n$ or at $n'$.

Conversely, if $n$ and $n'$ are not in the same information set of $i$, then when $i$ is at $n$, he knows he is not at $n'$.

### *Why must information sets form a partition?*

Suppose information sets did not form a partition.

Imagine three nodes $n, n', n''$ such that:

- at $n''$, player $i$ cannot distinguish between $n'$ and $n''$,
- at $n'$, player $i$ knows he is not at $n$,
- but at $n$, player $i$ cannot rule out being at $n'$.

This creates an inconsistency: reasoning about which node one is at would depend on the node itself. Requiring $H_i$ to be a partition prevents this kind of epistemic contradiction.

### *Consistency of available actions*

If $n$ and $n'$ belong to the same information set $h_i$, then the available actions must coincide:

$$A_i(n) = A_i(n').$$

Otherwise, the difference in available actions would reveal to player $i$ which node he is at.

For example, suppose

$$A_i(n) = \{L, R\} \qquad \text{and} \qquad A_i(n') = \{L, C, R\}.$$

Then at node $n$, player $i$ could infer that he is not at $n'$, because action $C$ is unavailable. This contradicts the idea that $n$ and $n'$ belong to the same information set.

Therefore, actions must be constant within each information set.

We denote by

$$A_i(h_i)$$

the set of actions available to player $i$ at any node in information set $h_i$.

### *Perfect information*

An extensive form game has *perfect information* if every information set is a singleton.

That is, for every player $i$ and every $h_i \in H_i$,

$$|h_i| = 1.$$

In this case, whenever a player moves, he knows exactly at which node he is.[3]

[3] **Perfect vs. Complete Information.** Perfect information refers to the information structure of the game tree: at every decision node, the moving player knows exactly where he is. Formally, each information set is a singleton, i.e., for every player $i$ and every $h_i \in H_i$, we have $|h_i| = 1$.

Complete information, instead, refers to payoffs: all players know the structure of the game and the payoff functions of all players.

These are logically distinct concepts. A game may have perfect information but incomplete information about payoffs (e.g., a sequential signaling game), or complete information but imperfect information about past moves.

*Example.* A simple two-period Stackelberg game has perfect information: Player 1 chooses $a \in A$, Player 2 observes $a$ and then chooses $b \in B$, and payoffs are realized. Since Player 2 observes Player 1's action before moving, every information set contains exactly one node, so $|h_i| = 1$ for all $i$.

### *Example: information sets and the Osborne–Rubinstein formalization*

The following example illustrates how information sets encode what a player knows at the moment of choosing an action.

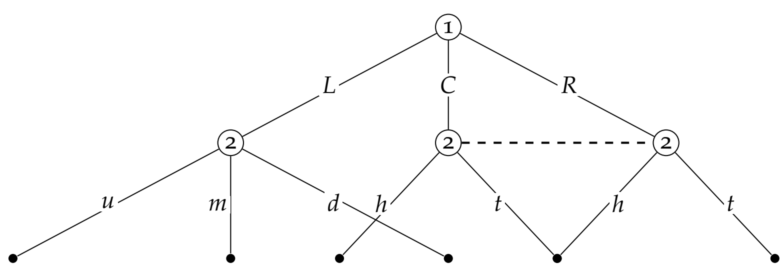


In the picture, player 1 moves first and chooses among $L, C, R$. If $L$ is chosen, player 2 moves at the left node and chooses among $u, m, d$. If $C$ or $R$ is chosen, player 2 chooses between $h$ and $t$, but the dashed line indicates that these two nodes are in the *same information set* of player 2: when moving there, player 2 cannot distinguish whether player 1 played $C$ or $R$.

We can represent the same game using the Osborne–Rubinstein convention in which every player chooses an action at every nonterminal node, and players who are inactive take a dummy action $*$.

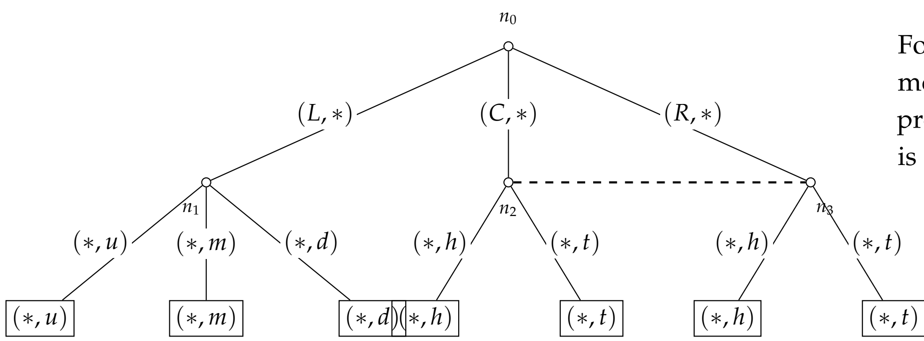


Formalization with dummy moves: actions are recorded as profiles $(a_1, a_2)$. The dashed line is player 2's information set.

**Remark.** Player 1 has perfect information in this example in the sense that each of his information sets is a singleton. By contrast, player 2 has imperfect information: at nodes $n_2$ and $n_3$ he does not know whether player 1 previously chose $C$ or $R$. In particular, if player 1 remembers his own past action, then after choosing $C$ or $R$ he knows which node is reached and therefore anticipates which outcomes correspond to player 2's subsequent move.

Formally, the set of (nonterminal) decision nodes is

$$N \setminus Z = \{n_0, n_1, n_2, n_3\}.$$

Player 1's information sets are

$$H_1 = \{\{n_0\}\},$$

since he is the only active player at the root node $n_0$.

Player 2 moves after $L$, $C$, or $R$, so his information sets are

$$H_2 = \{\{n_1\}, \{n_2, n_3\}\},$$

where the non-singleton information set $\{n_2, n_3\}$ is represented by the dashed line in Figure 6. The game therefore does *not* have perfect information (because player 2 has an information set with $|h_2| > 1$), but player 1 does have perfect information in the sense that all of his information sets are singletons.

*Payoffs and standard restrictions on information sets*

*EF–5: Payoffs.* Payoffs are given by functions

$$u_i : Z \to \mathbb{R}, \qquad i \in I,$$

which assign a real payoff to each player at each terminal node. This completes the description of the game.

In addition, we typically impose two restrictions on information sets.

*No absentmindedness* The restriction of *no absentmindedness* rules out situations in which a player may be unable to tell whether he has already moved (or is about to move) at a previous node in the same play of the game.

Intuitively, it rules out information sets that contain nodes that lie on the same path of play (so that reaching one implies that the other has occurred).

Let $H_i$ denote the collection of information sets of player $i$. The game satisfies *no absentmindedness* if for every player $i$ and every information set $h_i \in H_i$, the following holds:

$$\text{There do not exist } n, n' \in h_i \text{ such that } n \prec n'.$$

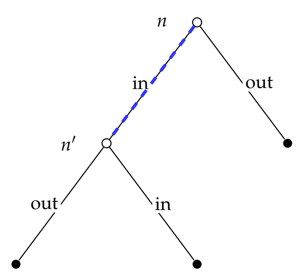


Absentmindedness (forbidden): an information set contains two nodes $n$ and $n'$ such that $n \prec n'$.

Here $n \prec n'$ means that node $n'$ is a successor of $n$ (i.e. $n$ lies on the unique path from the root to $n'$).

Equivalently, no information set contains two nodes that lie on the same path of play. Thus, if a player is at some node in an information set $h_i$, he can never be uncertain about whether he has already moved earlier along that same history.

*Perfect recall* The restriction of *perfect recall* rules out forgetting one's own past actions and past information. A common informal violation is:

> "I remember that I moved before, but I do not remember what action I chose."

Let $H_i$ denote the collection of information sets of player $i$. The game satisfies *perfect recall* if for every player $i$, the following holds:

If two nodes $n, n'$ belong to the same information set $h_i \in H_i$, then along the path from the root to $n$ and to $n'$:

1. player $i$ has taken the same number of past moves, and
2. the sequence of actions previously chosen by player $i$ is identical at $n$ and $n'$.

Equivalently, whenever a player reaches an information set, he remembers all of his own past actions and all of his past information sets.

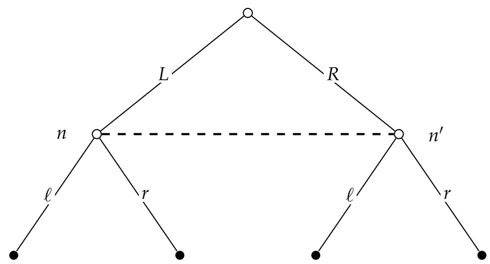


Perfect recall violation (forbidden): nodes $n$ and $n'$ belong to the same information set even though the player previously chose different actions ($L$ vs. $R$).

In this figure, the same player moves first by choosing $L$ or $R$, and then moves again at a later stage. However, the information set connects the two second-stage nodes reached after $L$ and $R$. Thus, at the second decision point, the player knows that he has moved before, but cannot distinguish whether he previously chose $L$ or $R$. This violates perfect recall because the sequence of the player's own past actions is not identical across the nodes in the information set.

## *Appendix to Class 12: Some Games from PS5 (Exercise 3)*

Let

$$\Gamma = (A_1, \dots, A_I, N, H_1, \dots, H_I, U_1, \dots, U_I)$$

be an extensive-form game.

**Definition** (Observable actions)**.** We say that $\Gamma$ has *observable actions* if, for each player $i = 1, \dots, I$, each information set $h_i \in H_i$ is a singleton.

Observe that if a game has observable actions, then

$$H_1 = \dots = H_I.$$

**Definition** (Perfect information)**.** We say that $\Gamma$ has *perfect information* if it has observable actions and at most one player has a non-trivial decision at each information set. That is, if $h \in H_1 = \dots = H_I$ and $|A_i(h)| \geq 2$, then $|A_j(h)| = 1$ for all $j \neq i$.

### *Two games without observable actions*

We construct two extensive-form games that satisfy:

- no absentmindedness,
- perfect recall,
- but do *not* have observable actions.

*Game 1: Timing Known, but Actions Not Observable* There are two players, 1 and 2.

At the initial node $n_0$, Player 1 chooses

$$A_1(n_0) = \{L, R\}.$$

If $L$ is chosen, the game proceeds to node $n_L$. If $R$ is chosen, the game proceeds to node $n_R$.

At both $n_L$ and $n_R$, Player 2 chooses

$$A_2(n_L) = A_2(n_R) = \{x, y\}.$$

Terminal payoffs are arbitrary.

The information structure is:

$$h_2 = \{n_L, n_R\}, \qquad H_2 = \{h_2\}.$$

Player 1 has only singleton information sets:

$$H_1 = \{\{n_0\}\}.$$

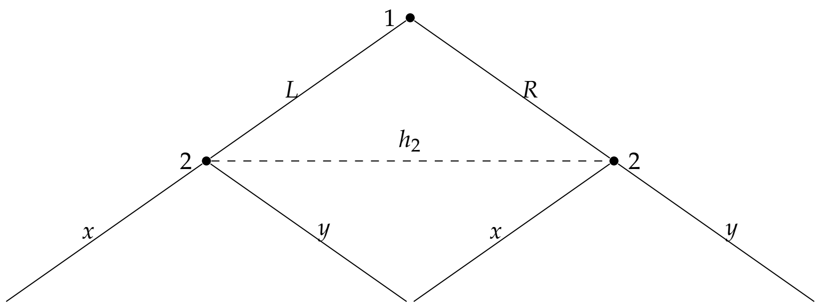


- The game does *not* have observable actions since $h_2$ is not a singleton.

- No absentmindedness: each player moves at most once.

- Perfect recall: each player remembers his own past actions.

- Timing is known: Player 1 knows he moves first; Player 2 knows he moves second.

*Game 2: A Player Does Not Know Whether He Is Second or Third* There are three players: $1, 2, 3$.

Player 1 moves first and chooses

$$A_1(n_0) = \{L, R\}.$$

If $L$ is chosen, Player 2 moves next at node $n_L$. If $R$ is chosen, Player 3 moves at node $n_R$.

After Player 3's move, Player 2 moves.

Thus, Player 2 moves:

- second along branch $L$,

- third along branch $R$.

We define the information set of Player 2 as

$$h_2 = \{n_L, n_R^m, n_R^n\},$$

where $n_R^m$ and $n_R^n$ are the nodes following Player 3's actions.

Thus, Player 2 does not know whether he is second or third.

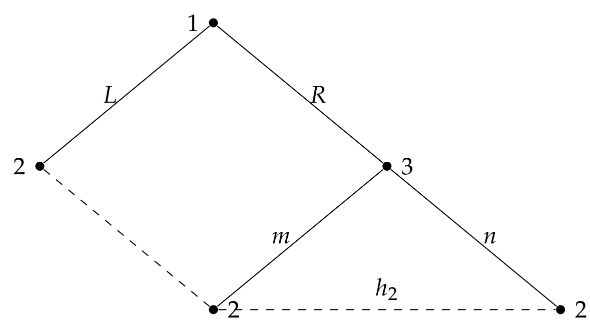


- The game does *not* have observable actions since $h_2$ is not a singleton.
- No absentmindedness: no player encounters the same information set twice.
- Perfect recall: Player 2 has not moved before reaching $h_2$, hence forgets nothing.
- Player 2 does not know whether he is second or third.

*Two games with observable actions*

We now construct two extensive-form games that satisfy:

- no absentmindedness,
- perfect recall,
- and *observable actions*.

The two examples differ in whether they satisfy *perfect information* (Definition 3).

*Game 3: Observable actions but* not *perfect information* There are two players, 1 and 2. There is a single decision node $h$ at which both players choose simultaneously.[4]

At the unique node $h$, the available actions are

$$A_1(h) = \{U, D\}, \qquad A_2(h) = \{L, R\}.$$

The terminal outcome is determined by the pair $(a_1, a_2) \in A_1(h) \times A_2(h)$, and payoffs are arbitrary.

Since there is only one decision node, each player's information set is a singleton:

$$H_1 = \{\{h\}\}, \qquad H_2 = \{\{h\}\}.$$

Hence, the game has observable actions.

[4] Formally, this can be represented as a single information set $h$ that belongs to $H_1$ and $H_2$, with both players having non-trivial action sets at $h$.

- *Observable actions:* every information set is a singleton.
- *No absentmindedness:* each player moves only once.
- *Perfect recall:* each player remembers his own past actions (trivial here).
- *Not perfect information:* Definition 3 fails, because at the same information set $h$ we have $|A_1(h)| \geq 2$ and $|A_2(h)| \geq 2$, i.e., more than one player has a non-trivial decision at $h$.

A convenient extensive-form representation is the following "simultaneous-move" tree:

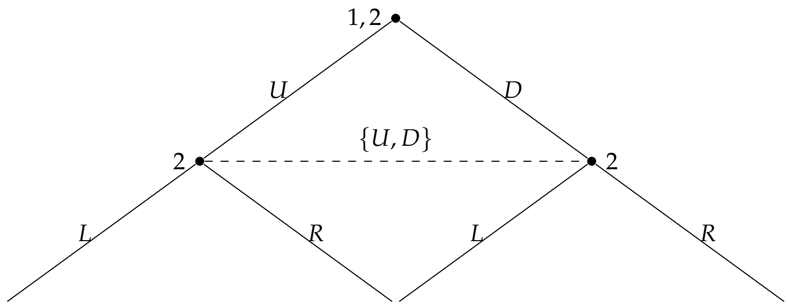


In this representation, Player 2 does not observe Player 1's action when choosing $L$ or $R$, which captures simultaneity. Notice, however, that Player 2's information set is *not* a singleton in this tree; the underlying simultaneous-move interpretation is that the information sets are singletons in the "one-node" model and the dashed line is merely a graphical device to represent simultaneity in a sequential tree. If you prefer a strict literal reading of Definition 2 in the drawn tree, use the "one-node" description above and omit the dashed line.

*Game 4: Observable actions and perfect information* There are two players, 1 and 2.

At the initial node $n_0$, Player 1 chooses

$$A_1(n_0) = \{U, D\}.$$

If Player 1 chooses $U$, the game moves to node $n_U$; if Player 1 chooses $D$, it moves to node $n_D$. At both $n_U$ and $n_D$, Player 2 chooses

$$A_2(n_U) = A_2(n_D) = \{L, R\}.$$

Terminal payoffs are arbitrary.

All information sets are singletons:

$$H_1 = \{\{n_0\}\}, \qquad H_2 = \{\{n_U\}, \{n_D\}\}.$$

Hence, the game has observable actions.

- *Observable actions:* every information set is a singleton.
- *No absentmindedness:* no player encounters the same information set twice.
- *Perfect recall:* the game has perfect information (hence perfect recall), and in addition each player moves at most once.
- *Perfect information:* at each node, only one player has a non-trivial action set.

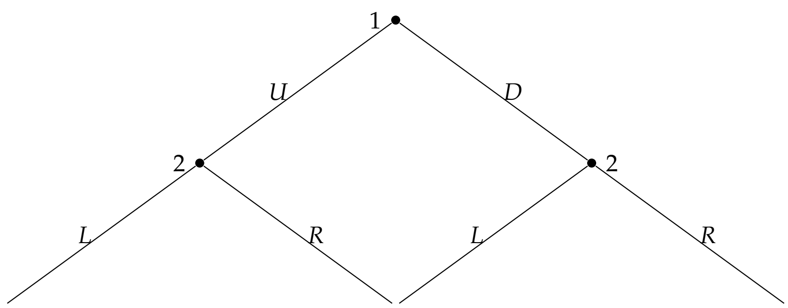


### *A Game with $H_1 = H_2$ but without Observable Actions*

Consider an extensive-form game with two players 1 and 2. Let the set of non-terminal nodes be $N = \{n_0, n_L, n_R\}$. At the initial node $n_0$, Player 1 chooses an action $a \in \{L, R\}$. If $L$ is chosen, the game proceeds to node $n_L$; if $R$ is chosen, it proceeds to node $n_R$. At both $n_L$ and $n_R$, Player 2 chooses an action $b \in \{x, y\}$. Each pair of actions $(a, b)$ leads to a terminal node with arbitrary payoffs.

The information structure is defined as follows. Let

$$h = \{n_L, n_R\}.$$

Define

$$H_1 = \{h\}, \qquad H_2 = \{h\}.$$

Hence,

$$H_1 = H_2.$$

This game does *not* have observable actions. Indeed, the information set $h$ is not a singleton, so Player 2 cannot distinguish between $n_L$ and $n_R$ when making a decision. Therefore, the condition for observable actions fails.

The game satisfies *no absentmindedness*. No player encounters the same information set more than once along any history of play. Player 1 moves only at $n_0$, and Player 2 moves at most once (either at $n_L$ or at $n_R$). Hence, no player can reach the same information set twice within a single path of the game.

The game also satisfies *perfect recall*. Each player remembers his own past actions and the information available at the time of choice. Since each player moves at most once, perfect recall holds trivially.

In conclusion, this example shows that it is possible to have $H_1 = \cdots = H_I$ while the game does not have observable actions, and that this does not imply absentmindedness.

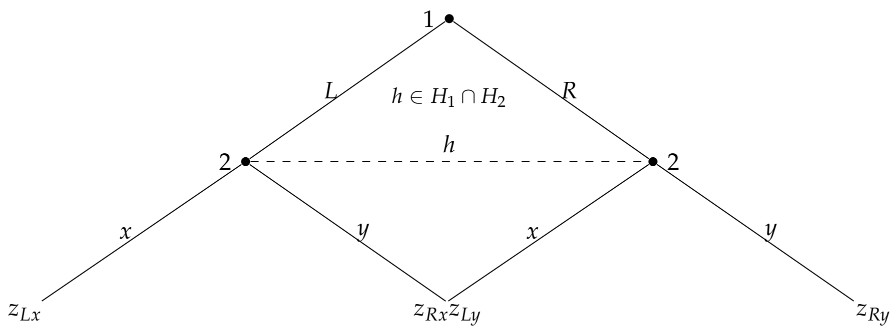


Game with $H_1 = H_2 = \{h\}$ and $h = \{n_L, n_R\}$.

# *Class 13: Strategies in Extensive Form Games*

### *Example 2: A simultaneous-move game and alternative visualizations*

Suppose we have a simultaneous-move game in which each player $i \in \{1, 2\}$ simultaneously chooses an action from

$$A_i = \{\square_i, \triangle_i\}.$$

*Formalization (as an extensive-form representation).* A convenient way to "formalize" a simultaneous-move game in extensive form is to use a single initial node and interpret the four branches as the four possible action profiles $(a_1, a_2) \in A_1 \times A_2$.

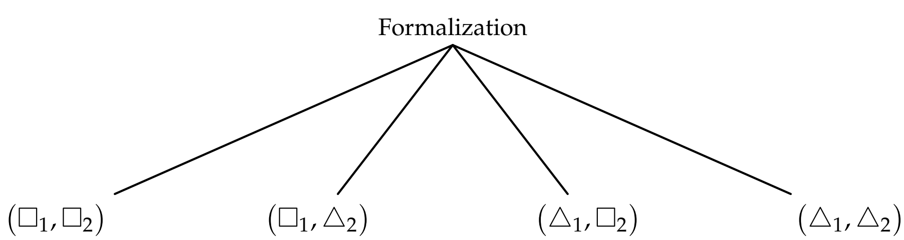


*Visualization (normal/strategic form).* The same simultaneous-move game can be visualized as a $2 \times 2$ normal-form table. (At this stage we only need the action labels; payoffs can be inserted later if needed.)

*Alternative visualizations ("from the world without simultaneous moves").* One can also represent the same simultaneous interaction by letting one player move first and the other move second, but placing the second mover's decision nodes in a single information set, so that the second mover does not observe the first mover's action.

Visualization

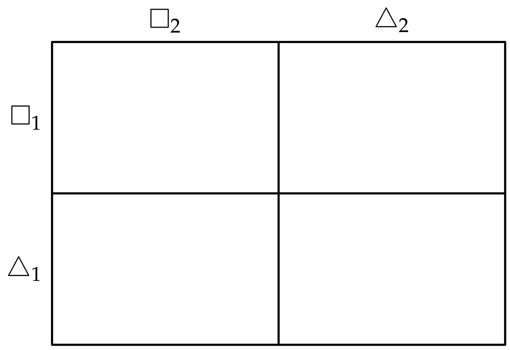


Alternative visualizations (from the world without simultaneous moves)

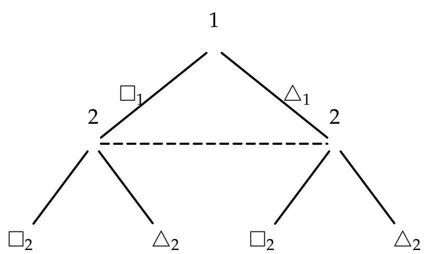


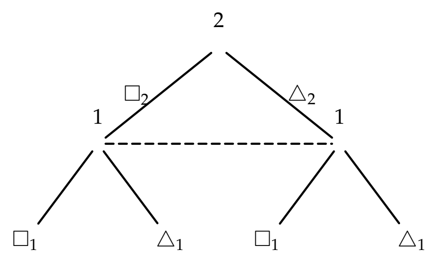


*Example 3: Strategies as contingent plans (and why beliefs require the full plan)*

In this example, player 2 can make different choices contingent on the information he has when he is called to move. The relevant point is that a *strategy* for player 2 is a *contingent plan*: it specifies what 2 would do after every history at which 2 might move.

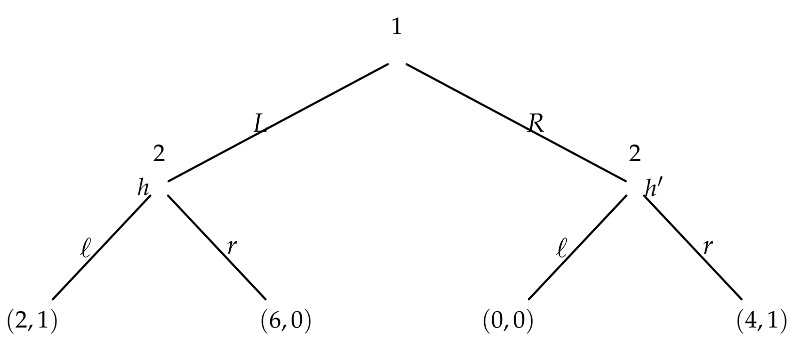


A strategy for player 2 must specify an action after both histories $h$

and $h'$. For instance, the contingent plan

"choose $\ell$ after $h$ and choose $r$ after $h'$"

is a single strategy for player 2 (a full contingent plan).

If player 1 anticipates this contingent plan, then the induced terminal history is the one associated with $(R,\ r \text{ after } h')$, i.e. the terminal node reached by $R$ followed by $r$.

It is tempting to think that one can ignore player 2's contingent plan at histories that are not reached, but that is wrong for equilibrium reasoning: for player 1 to know how to play, player 1 needs a belief about what player 2 would do after *both* $h$ and $h'$. If the solution concept is equilibrium-based, then player 1's belief is pinned down by what player 2 would actually play, which requires specifying player 2's *full* contingent plan in the first place.

### *Strategies in Extensive-Form Games*

**Definition.** A (pure) strategy of player $i$ is a mapping

$$s_i : H_i \longrightarrow A_i$$

such that

$$s_i(h_i) \in A_i(h_i) \quad \text{for all } h_i \in H_i.$$

That is, a strategy assigns an action to every information set at which player $i$ may be called to move.

Denote by $S_i$ the set of strategies of player $i$.

### *Example 3 (Revisited)*

From the previous tree:

*Strategies for Player 1.* Player 1 moves only at the root, so

$$S_1 = \{L, R\}.$$

*Convention.* If $A_i(h_i)$ is a singleton, we can ignore that node when writing the strategy, since no choice is involved.

*Strategies for Player 2.* Player 2 moves at two information sets: after $h$ and after $h'$.

Thus a strategy must specify an action at both nodes.

Hence:

$$S_2 = \{(\ell,\ell),(\ell,r),(r,\ell),(r,r)\}.$$

Think in words:

- $(\ell, \ell)$: choose $\ell$ after $h$ and $\ell$ after $h'$.
- $(\ell, r)$: choose $\ell$ after $h$ and $r$ after $h'$.
- $(r, \ell)$: choose $r$ after $h$ and $\ell$ after $h'$.
- $(r, r)$: choose $r$ after $h$ and $r$ after $h'$.

A strategy is a full contingent plan.

**Definition.** (Path Function) Given a strategy profile $s = (s_1, \ldots, s_I)$, define the path function

$$J : \prod_{i=1}^{I} S_i \longrightarrow Z,$$

where $Z$ is the set of terminal nodes. Thus $J(s)$ is the terminal node induced by the strategy profile $s$.

For example:

$$\begin{aligned} J(L,(\ell,\ell)) &= z_1 & \quad J(R,(\ell,\ell)) &= z_3 \\ J(L,(\ell,r)) &= z_1 & \quad J(R,(\ell,r)) &= z_4 \\ J(L,(r,\ell)) &= z_2 & \quad J(R,(r,\ell)) &= z_3 \\ J(L,(r,r)) &= z_2 & \quad J(R,(r,r)) &= z_4 \end{aligned}$$

Only the action that lies on the realized path matters for the terminal node, but the entire strategy must still be specified.

*Strategic Form Induced by an Extensive-Form Game*

**Definition.** The extensive-form game $\Gamma$ induces a strategic-form game

$$G = (S_1, \ldots, S_I, \mu_1, \ldots, \mu_I),$$

where:

- $S_i$ is the set of strategies of player $i$.
- $\mu_i = U_i \circ J$ is the payoff function of player $i$.

Thus,

$$\mu_i(s_1, \ldots, s_I) = U_i\big(J(s_1, \ldots, s_I)\big).$$

For Example 3, the induced normal form is:

| | $(\ell,\ell)$ | $(\ell,r)$ | $(r,\ell)$ | $(r,r)$ |
|---|---|---|---|---|
| $L$ | $(2,1)$ | $(2,1)$ | $(6,0)$ | $(6,0)$ |
| $R$ | $(0,0)$ | $(4,1)$ | $(0,0)$ | $(4,1)$ |

### *Analyzing the Game*

- We can always write down the induced strategic-form game and solve using Nash Equilibrium.
- However, as in earlier examples, Nash Equilibrium may give unreasonable predictions in dynamic settings.
- We would like some notion of sequential optimality.
- We want players to behave optimally at every stage of the game.
- We want players to "reason" that the other player behaves in a sequentially optimal manner.

This motivates refinements of Nash equilibrium.

### *Baseline Refinement: Subgame Perfect Equilibrium*

Idea: we look for an equilibrium throughout the entire extensive-form game.

*What is a subgame?* A subgame is a "mini game" that begins at a node and contains all successor nodes, without breaking any information sets.

Graphically, a subgame must have a single initial node.

*Formal requirement (no breaking of information sets).* Assuming no absent-mindedness, a subgame can be identified by a singleton information set $\{h\}$ such that:

If a node $h'$ follows $h$ and $h'$ belongs to some information set $H_i$, then all nodes in $H_i$ must follow $h$.

In other words, creating the subgame cannot split an information set.

We are not allowed to break up information sets when defining a subgame.

Subgame Perfect Equilibrium (SPE) will require Nash Equilibrium in every subgame.

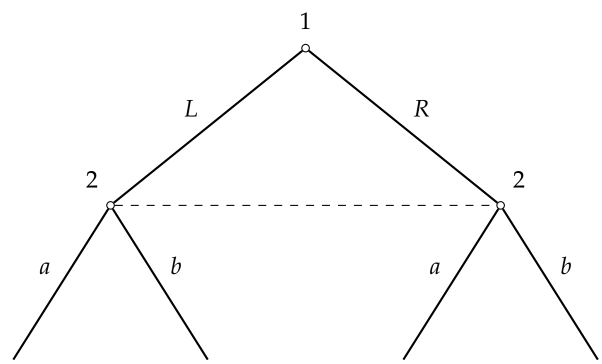


Game with imperfect information where there is NO subgame starting after $L$.

## *Appendix to Class 13: Subgames (How to Spot Them, Especially in Weird Cases)*

This appendix collects "odd" extensive-form configurations designed to train the mechanical identification of subgames.

### *Reminder: What counts as a subgame?*

A *subgame* is an extensive-form game that starts at a node $x$ and contains *all* successors of $x$, subject to the key restriction:

> A subgame cannot *break* an information set. Equivalently, if the subgame contains one node from an information set $H_i$, then it must contain *all* nodes in $H_i$.

Hence, candidate initial nodes for a subgame are: (i) singleton information sets (the player moving at that node knows exactly where they are), and (ii) nodes such that, once you include all successors, you do not end up including only part of some information set.

The *whole game* is always a subgame (starting at the root).

### *Case A: Perfect information (every decision node is a subgame root)*

Here there are no information sets with multiple nodes, so every decision node (singleton information set) starts a subgame.

### *Case B: One information set blocks the obvious subgames ("looks like" a subgame, but isn't)*

Although the left and right branches look like candidate "mini games," neither node reached after $L$ or $R$ is a singleton information set: they belong to the *same* information set of player 2. Starting a subgame at either one would include only part of that information set, which is forbidden.

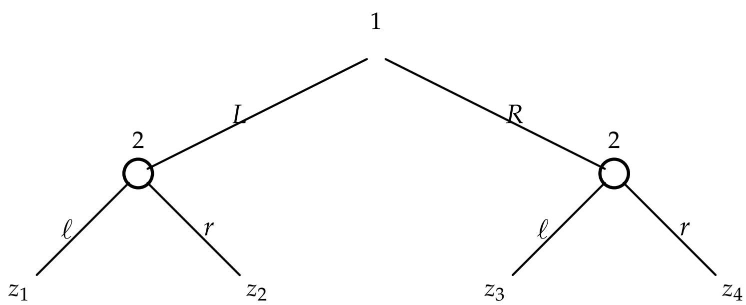


Subgames: whole game + subgame starting at each circled node.

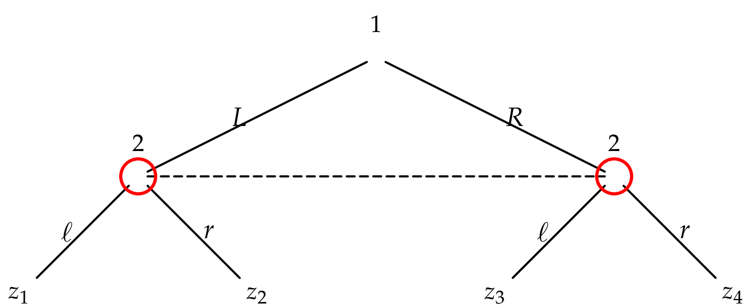


Only subgame: the whole game (root). Circles are *not* subgame roots.

*Case C: Subgame exists on one side but not the other (asymmetry due to information sets)*

Consider the following story.

Player 1 is a firm and Player 2 is an auditor. At the initial node, Player 1 chooses between: L (transparent behavior) or R (opaque behavior).

*Left branch (L).* If the firm behaves transparently, the auditor receives clear evidence and knows with certainty that $L$ was chosen. Thus, the auditor's node on the left is a singleton information set. The auditor then chooses $\ell$ or $r$, leading directly to terminal outcomes $z_1$ or $z_2$. Since this node is a singleton information set, it starts a subgame.

*Right branch (R).* If the firm behaves opaquely, the auditor receives a noisy or manipulated signal. Formally, there are two possible histories at which Player 2 might move, but the auditor cannot distinguish between them. These two nodes belong to the same information set. In the diagram, only one of them is drawn (the circled node), and the dashed line indicates that another node belongs to the same information set.

Because Player 2 does not know which of these two nodes has been reached, this node is *not* a singleton information set. Therefore, it cannot start a subgame: starting there would break the information set.

The subgames in this example are:

- The whole game (starting at the root),
- The subgame starting at the left 2-node only.

There is no subgame starting at the right 2-node because it belongs to a non-singleton information set.

The left 2-node is a singleton information set, so it starts a subgame. The right 2-node is *not* a singleton information set (it shares an information set with another node), so it cannot start a subgame.

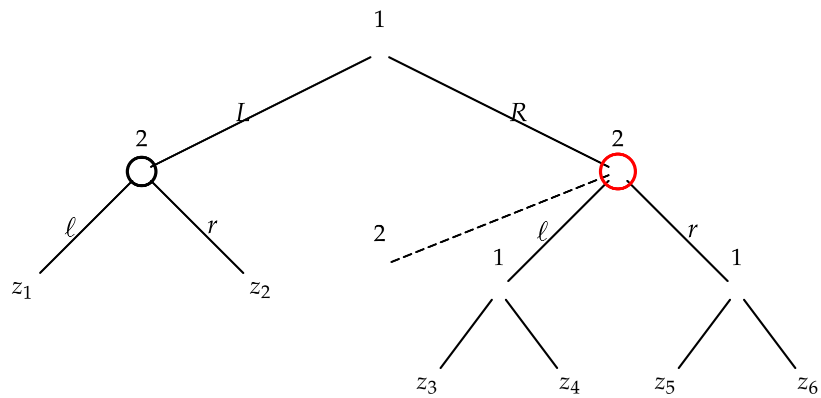


Subgames: whole game + the one starting at the circled left node only.

*Case D: Information set at the top ("no initial nodes" aside from the whole game)*

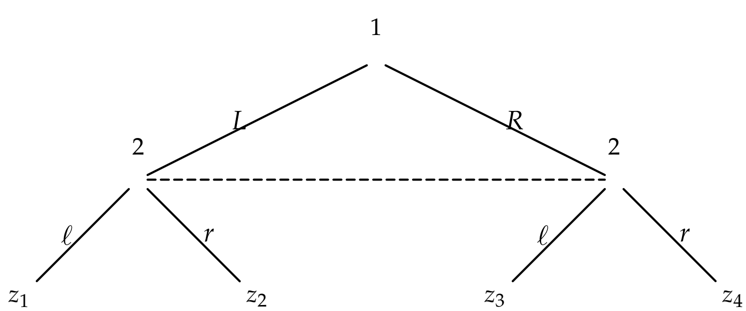


No proper subgames: any candidate start would break the information set.

This is the canonical "weird" case: the only subgame is the entire game. Any attempt to start a subgame after $L$ or after $R$ would necessarily select only one node of player 2's information set, which is not allowed.

*Case E: Singleton initial node, with a later information set fully contained (valid subgame)*

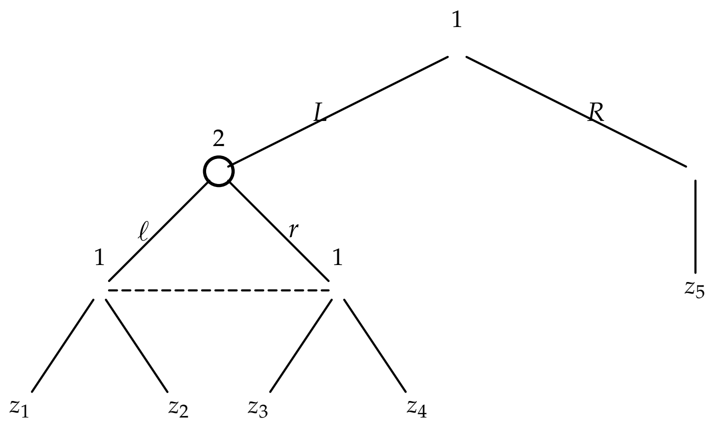


The subgame at the circled node is valid: it contains the entire later info set.

Starting at the node after $L$ is allowed because it is a singleton information set. Even though an information set appears later (the two 1-nodes connected by the dashed line), that information set is *entirely* contained in the continuation from $L$. So the subgame does not "split" it.

*Checklist for subgame detection (mechanical procedure)*

Given a candidate node $x$:

1. (Single initial node) Start with the subtree consisting of $x$ and all its successors.
2. (Info-set closure) For every information set that intersects the subtree, check whether the subtree contains *all* nodes in that information set.
3. If any information set is only partially included, then $x$ cannot start a subgame.
4. If the closure condition holds, then the subtree is a subgame.

The difficult cases are precisely those where step 2 fails: the candidate subtree would force you to include part of an information set but not the rest.

*Question 5 from PS5: Ann–Bob (Jumping, Imperfect Observation, and Recall)*

Ann moves first and chooses one of three actions:

$$a_A^0 \in \{N,\ J1,\ J2\},$$

where $N$ means "not jump," $J1$ means "jump once," and $J2$ means "jump twice."

*Bob's information and move.* If Ann chooses $N$, then Bob observes that she did not jump and chooses

$$a_B^N \in \{C,\ S,\ 0\},$$

interpreted as giving Ann chocolate $(C)$, slime $(S)$, or nothing $(0)$.

If Ann chooses $J1$ or $J2$, Bob observes only that Ann jumped, but not how many times. Thus Bob has an information set containing the two decision nodes corresponding to histories $J1$ and $J2$, and at that information set he chooses

$$a_B^J \in \{C,\ G\},$$

interpreted as giving Ann chocolate $(C)$ or gelato $(G)$.

*Ann's final move (perfect recall).* If Bob gives gelato, then Ann chooses whether to share it or not:

$$a_A^G \in \{Sh,\ NS\}.$$

Ann remembers how many times she jumped. Hence, the two Ann decision nodes after $(J1, G)$ and $(J2, G)$ are *not* in the same information set.

*Sets of pure strategies.* A pure strategy for Ann specifies her initial action and, if gelato is offered, her sharing decision after each possible jump history:

$$S_A = \Big\{ \big(a_A^0,\ a_A^G(J1),\ a_A^G(J2)\big) :\ a_A^0 \in \{N, J1, J2\},\ a_A^G(J1) \in \{Sh, NS\},\ a_A^G(J2) \in \{Sh, NS\} \Big\}.$$

Thus $|S_A| = 3 \cdot 2 \cdot 2 = 12$.

A pure strategy for Bob specifies what he would do after $N$ and what he would do after observing "jumped":

$$S_B = \Big\{ \big(a_B^N,\ a_B^J\big) :\ a_B^N \in \{C, S, 0\},\ a_B^J \in \{C, G\} \Big\}.$$

Thus $|S_B| = 3 \cdot 2 = 6$.

*Remark.* The key information structure is that Bob cannot distinguish $J1$ from $J2$ (single information set), while Ann has perfect recall and therefore distinguishes her own past action when deciding whether to share gelato.

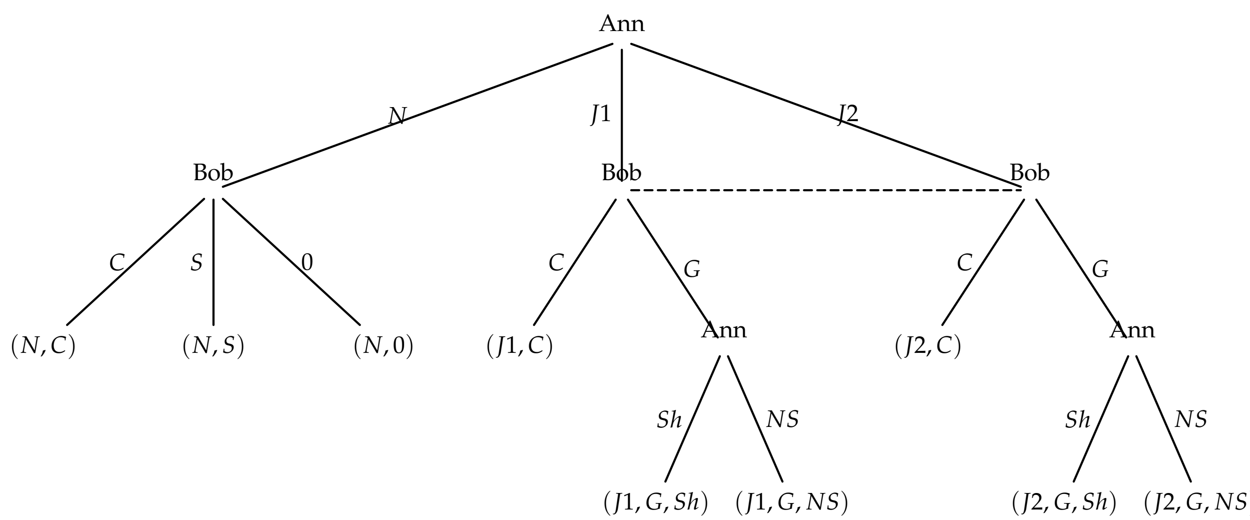


Bob observes "jumped" but not how many times; Ann remembers.

# *Class 14: Subgames and Subgame Perfect Equilibrium*

## *Subgames*

**Definition.** A subgame is identified with a singleton information set $\{h\}$ such that:

1. $h$ is the unique initial node of the subgame.
2. If a node $n'$ follows $h$ and $n' \in H_i$ for some information set $H_i$, then every node in $H_i$ must follow $h$.

The second condition ensures that we do not break information sets. Equivalently, once the subgame starts at $h$, it must contain whole information sets, never fragments of them.

*Remark (Minimal departure).* When analyzing a subgame starting at $h$, we modify strategies only insofar as needed to allow the path to pass through $h$. Outside of the subtree following $h$, strategies are kept unchanged. This will be formalized below.

Each subgame associated with $h$ induces a new strategic-form game, denoted $G(h)$.

## *Strategies in the Subgame*

For each player $i$, define the set of strategies that allow $h$ as:

$$S_i(h) = \{s_i \in S_i \ : \ \exists s_{-i} \text{ such that } J(s_i, s_{-i}) \text{ passes through a node in } h\}.$$

Thus $S_i(h)$ consists of strategies that are compatible with reaching $h$.

**Definition.** (Minimal modification of strategies) Fix an information set $h$. For each player $i$, let $S_i(h) \subseteq S_i$ denote the set of strategies that *allow* $h$ to be reached (for some profile of opponents' strategies).

Define the *minimal-departure operator*

Yes, I defined it literally as a $\cdot^h$ because that is what Amanda wanted to tell us.

$$\cdot^h : S_i \to S_i(h), \qquad s_i \mapsto s_i^h,$$

where $s_i^h$ is obtained from $s_i$ by modifying only those components of $s_i$ that lie on paths leading to $h$, and only insofar as necessary to make $h$ reachable.

Formally:

- If $s_i \in S_i(h)$, then $s_i^h = s_i$.
- If $s_i \notin S_i(h)$, then $s_i^h \in S_i(h)$ and $s_i^h$ coincides with $s_i$ at every information set of player $i$ that does not lie on a path to $h$.

Thus, $s_i^h$ represents a *minimal departure* from $s_i$. The operator is idempotent:

$$(s_i^h)^h = s_i^h.$$

The subgame $G(h)$ is then the strategic-form game induced by restricting attention to the subtree starting at $h$ and to the strategy sets $S_i(h)$ for each player $i$.

## *Example 4*

Consider the following game.

Player 1 first chooses between In and Out. If Out is chosen, the game ends with payoff $(0,3)$.

If In is chosen, Player 2 chooses between $L$ and $R$, and then Player 1 chooses between $U$ and $D$.

The terminal payoffs are:

| | $L$ | $R$ |
|---|---|---|
| $U$ | $(2,4)$ | $(0,0)$ |
| $D$ | $(0,0)$ | $(4,2)$ |

Let $h$ denote the information set reached after $(\text{In},\cdot)$.

Then:

$$S_1(h) = \{(\text{In}, U), (\text{In}, D)\}$$

since Player 1 must choose In in order to allow $h$.

For Player 2,

$$S_2(h) = \{L, R\} \text{ or } \{(\cdot, L), (\cdot, R)\}.$$

Observe how the minimal-departure operator works:

$$(\text{In}, U)^h = (\text{In}, U)$$
$$(\text{Out}, U)^h = (\text{In}, U)$$
$$(\text{In}, D)^h = (\text{In}, D)$$
$$(\text{Out}, D)^h = (\text{In}, D).$$

The only change required to allow $h$ is to replace Out by In. All other components of the strategy remain unchanged.

## Subgame Perfect Equilibrium

**Definition.** A strategy profile $s = (s_1, \dots, s_I)$ is a *subgame perfect equilibrium* (SPE) if for every subgame associated with $h$, the restricted profile

$$(s_1^h, \dots, s_I^h)$$

is a Nash equilibrium of $G(h)$.

## Analyzing Example 4

First compute the Nash equilibria of the subgame $G(h)$.

The normal form of $G(h)$ is:

| | $L$ | $R$ |
|---|---|---|
| $U$ | $(2,4)$ | $(0,0)$ |
| $D$ | $(0,0)$ | $(4,2)$ |

The Nash equilibria of $G(h)$ are:

$$((\text{In}, U), L) \quad \text{and} \quad ((\text{In}, D), R).$$

*Q1. Is $((In, U), L)$ an SPE?* No.

Although it is a Nash equilibrium of the subgame, it is not a Nash equilibrium of the whole game, because Player 1 prefers Out at the initial node:

$$u_1((\text{Out}, -), L) > u_1((\text{In}, U), L).$$

Thus it fails to be a Nash equilibrium of the entire game.

*Q2. Is $((In, D), R)$ an SPE?* Yes.

It is a Nash equilibrium of the subgame and a Nash equilibrium of the whole game.

*Q3. Is there an SPE in which Out is played?* Consider whether there exists $s_2$ such that $((\text{Out}, D), s_2)$ is an SPE.

For this to hold, the restriction to $G(h)$ must be a Nash equilibrium. That requires $s_2 = R$.

However,

$$((\text{Out}, D), R)$$

is not a Nash equilibrium of the whole game, since Player 1 would deviate to $(\text{In}, D)$.

Hence there is no SPE with $(\text{Out}, D)$.

Now consider $((\text{Out}, U), s_2)$.

For this to be an SPE, we need $((\text{Out}, U)^h, s_2^h)$ to be a Nash equilibrium of $G(h)$. This requires $s_2 = L$.

Indeed,

$$((\text{Out}, U), L)$$

is a Nash equilibrium of the whole game.

Therefore,

$$((\text{Out}, U), L)$$

is an SPE.

## *Recap: Subgame Perfect Equilibrium in Earlier Examples*

### *Example 1 (recall)*

Consider the following entry–deterrence game.

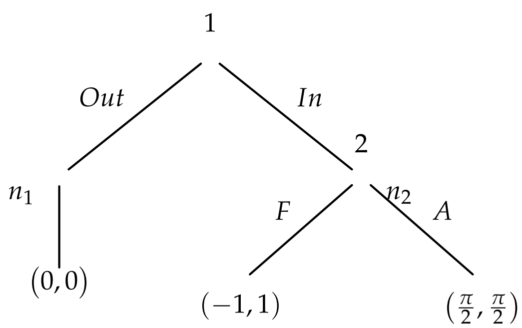


Example 1 (recall).

*

By backward reasoning, Player 2 prefers $A$ to $F$ at node $n_2$, so Player 1's best response at the initial node is to choose $In$. Thus the unique SPE is

$$\text{SPE: } (In, A).$$

### *Example 3 (recall)*

We return to Example 3 (from Class 13), now interpreted in terms of subgames $G(h)$ and $G(h')$.

For the subgames starting at $h$ and $h'$ we have:

$$NE\big(G(h)\big) = \{(L, \ell\ell), (L, \ell r)\}, \qquad NE\big(G(h')\big) = \{(R, \ell r), (R, rr)\}.$$

The unique strategy profile that is a Nash equilibrium in *both* subgames (and in the whole game) is

$$\text{SPE: } (R, \ell r).$$

This already "looks like" backward induction: we are enforcing best responses in every continuation of the game.

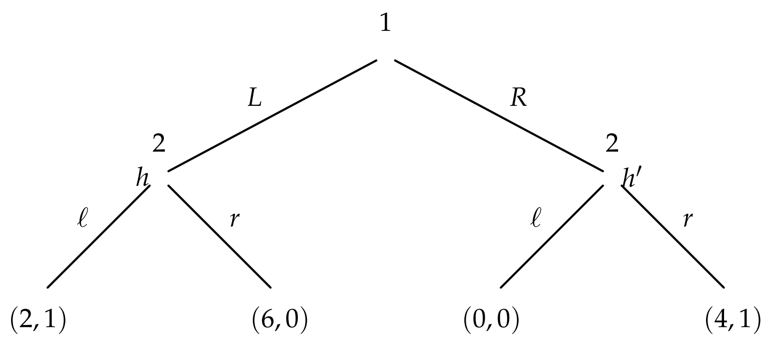


Example 3 (recall).

*

## *Backward Induction Algorithm (Perfect Information, Finite Horizon)*

We now state the backward induction (BI) algorithm for finite games of perfect information.

*Step 1.* Look at each last non-terminal node. At each such node, the player who moves chooses an action that is a best response *at that node*. (Notice: we are picking *actions* node by node.)

*Step 2.* Next, look at each non-terminal node that is followed only by Step 1 nodes or terminal nodes and that is followed by at least one Step 1 node. At each such node, choose an action that is a best response given the actions already chosen in Step 1.

Repeat this procedure, moving backwards in the tree, until we reach the initial node. Collect all selected actions into a full strategy profile. The result is the strategy profile induced by the backward induction algorithm.

## *Backward Induction and Subgame Perfect Equilibrium*

**Proposition 0.0.23.** *Fix a finite perfect-information game. If the backward induction algorithm induces the strategy profile* $(s_1, \ldots, s_I)$*, then* $(s_1, \ldots, s_I)$ *is a subgame perfect equilibrium. Conversely, if* $(s_1, \ldots, s_I)$ *is a subgame perfect equilibrium, then it is induced by the backward induction algorithm.*

## *Example 5*

We now illustrate the link between backward induction and SPE in a richer game with several subgames.

*Figure 5a: Extensive Form*

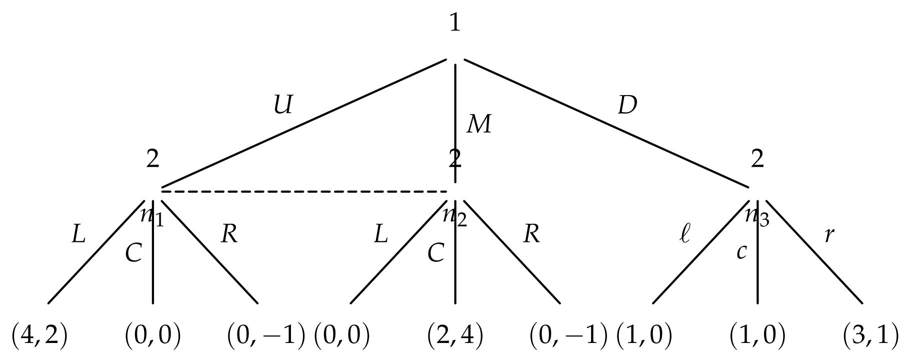


Figure 5a: Example 5.

Player 2 has three actions $L, C, R$ at her left information set $\{n_1, n_2\}$, and three actions $\ell, c, r$ at her right information set $n_3$. Hence her pure strategies can be indexed as

$$L\ell,\ Lc,\ Lr,\ C\ell,\ Cc,\ Cr,\ R\ell,\ Rc,\ Rr,$$

where, for instance, $Lc$ means "play $L$ at $\{n_1, n_2\}$ and $c$ at $n_3$," etc.

*Strategic Form and SPE*

The induced normal form has Player 1's strategy set $S_1 = \{U, M, D\}$ and Player 2's strategy set $S_2 = \{L\ell, Lc, Lr, C\ell, Cc, Cr, R\ell, Rc, Rr\}$. Entries follow mechanically from Figure 5a:

| | $L\ell$ | $Lc$ | $Lr$ | $C\ell$ | $Cc$ | $Cr$ | $R\ell$ | $Rc$ | $Rr$ |
|---|---|---|---|---|---|---|---|---|---|
| $U$ | $(4,2)$ | $(4,2)$ | $(4,2)$ | $(0,0)$ | $(0,0)$ | $(0,0)$ | $(0,-1)$ | $(0,-1)$ | $(0,-1)$ |
| $M$ | $(0,0)$ | $(0,0)$ | $(0,0)$ | $(1,4)$ | $(2,4)$ | $(1,4)$ | $(0,-1)$ | $(0,-1)$ | $(0,-1)$ |
| $D$ | $(1,0)$ | $(1,0)$ | $(3,1)$ | $(1,0)$ | $(1,0)$ | $(3,1)$ | $(1,0)$ | $(1,0)$ | $(3,1)$ |

Solving by backward induction (or, equivalently, imposing SPE) yields:

$$\text{SPE: } (U, L\ell), \quad (D, c\ell), \quad (D, R\ell).$$

But notice the last one:

$$(D, R\ell).$$

Here Player 1 expects that if he deviates to $U$, then Player 2 will play $R$. However, at both nodes $n_1$ and $n_2$, action $R$ is strictly worse for Player 2 than either $L$ or $C$. So why would Player 1 *reasonably* expect Player 2 to play $R$?

This illustrates again why we need subgame perfection and backward induction: they rule out equilibria that rely on such implausible (off-path) beliefs about future play.

# *Class 15: Subgames, SPE and Forward Induction*

## *Figure 5b: A Modified Version of Example 5*

We now slightly modify Example 5 to obtain a new game (Figure 5b). Player 1 first chooses between *In* and *D*. If he chooses *In*, he then chooses between *U* and *M*. After *U* or *M*, Player 2 moves at the information set $\{n_1, n_2\}$, exactly as in Figure 5a. If Player 1 chooses *D* at the initial node, the game moves to $n_3$, where Player 2 chooses between $\ell, c, r$.

Thus the *right-hand* part of the tree (node $n_3$) is unchanged relative to Figure 5a; the only change is that we have inserted an extra decision for Player 1 before reaching the left-hand part.

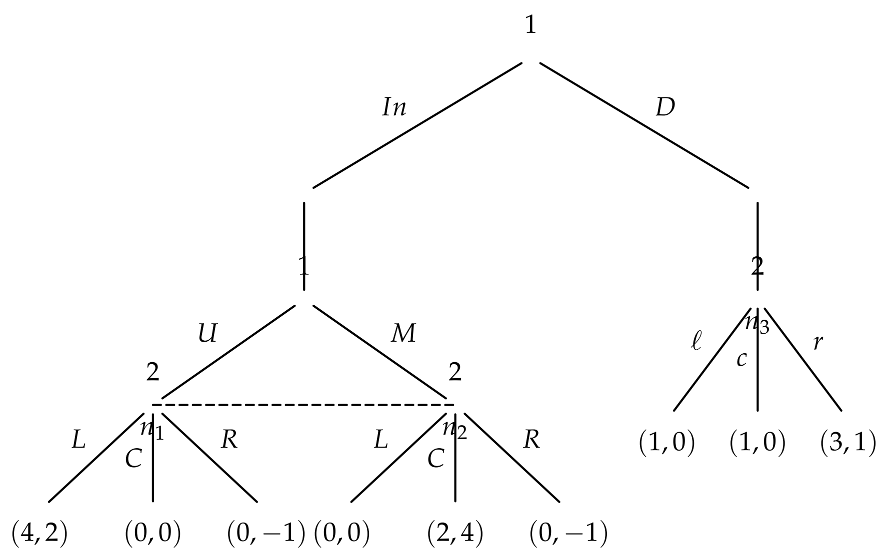


Figure 5b: Example 5 with an additional initial decision $In/D$.

The structure of payoffs is essentially the same as in Figure 5a. The key difference is that we have created a *new subgame*: the subtree starting at the node after *In*, where Player 1 chooses between *U* and *M* and Player 2 then moves at $\{n_1, n_2\}$. This new subgame will allow us to see more clearly what subgame perfection does and what it fails to do, and how it interacts with the "minimal departure" idea introduced in Class 14.

This modification does not change Player 2's strategy set: she still chooses an action at the left-hand information set $\{n_1, n_2\}$ $(L, C, R)$ and an action at the right-hand node $n_3$ $(\ell, c, r)$. For Player 1, however, we now have to specify both the initial move (*In* or *D*) and the continuation (*U* or *M* if *In*). Thus Player 1's pure strategies are

$$\{In\text{-}U,\ In\text{-}M,\ D\text{-}U,\ D\text{-}M\}.$$

Note that *D-U* and *D-M* are *realization-equivalent*: if Player 1 plays *D* at the first node, the continuation (*U* or *M*) is never reached.

**Definition.** Two strategies $s_1$ and $r_1$ of Player 1 are *realization-equivalent* if, for every strategy $s_2$ of Player 2, $(s_1, s_2)$ and $(r_1, s_2)$ induce the same terminal node.

In this game, *D-U* and *D-M* are realization-equivalent. If Player 1 chooses *D* at the initial node, the left-hand subtree is never reached, so the continuation plan (*U* or *M*) is irrelevant. Hence, for every $s_2$,

$$(D\text{-}U, s_2) \quad \text{and} \quad (D\text{-}M, s_2)$$

lead to the same outcome.

By contrast, *In-U* and *In-M* are *not* realization-equivalent, since once *In* is chosen, the choice between *U* and *M* affects the reached node and potentially the terminal payoff.[5]

[5] This illustrates that the strategic-form representation may contain distinct strategies that are behaviorally indistinguishable. In the present game, Player 1 has four pure strategies in strategic form, but only three distinct realization classes: $\{In\text{-}U\}$, $\{In\text{-}M\}$, and $\{D\text{-}U, D\text{-}M\}$. Realization equivalence identifies strategies that differ only at histories that are never reached.

### *Strategic Form of Figure 5b*

Player 2's pure strategies can be indexed as before by specifying an action at $\{n_1, n_2\}$ and an action at $n_3$:

$$L\ell,\ Lc,\ Lr,\ C\ell,\ Cc,\ Cr,\ R\ell,\ Rc,\ Rr,$$

where, for instance, *Lc* means "play *L* at $\{n_1, n_2\}$ and *c* at $n_3$".

The induced normal form has Player 1's strategy set

$$S_1 = \{In\text{-}U, In\text{-}M, D\text{-}U, D\text{-}M\}$$

and Player 2's strategy set

$$S_2 = \{L\ell, Lc, Lr, C\ell, Cc, Cr, R\ell, Rc, Rr\}.$$

The payoff table is:

| | $L\ell$ | $Lc$ | $Lr$ | $C\ell$ | $Cc$ | $Cr$ | $R\ell$ | $Rc$ | $Rr$ |
|---|---|---|---|---|---|---|---|---|---|
| *In-U* | (4,2) | (4,2) | (4,2) | (0,0) | (0,0) | (0,0) | (0,−1) | (0,−1) | (0,−1) |
| *In-M* | (0,0) | (0,0) | (0,0) | (1,4) | (2,4) | (1,4) | (0,−1) | (0,−1) | (0,−1) |
| *D-U* | (1,0) | (1,0) | (3,1) | (1,0) | (1,0) | (3,1) | (1,0) | (1,0) | (3,1) |
| *D-M* | (1,0) | (1,0) | (3,1) | (1,0) | (1,0) | (3,1) | (1,0) | (1,0) | (3,1) |

The first two rows correspond to entering and playing $U$ or $M$; the last two rows correspond to playing $D$ at the first node, with the continuation ($U$ or $M$) having no effect on realized payoffs.

### *The New Subgame and Its Nash Equilibria*

The node following *In* is the initial node of a proper subgame: from there, Player 1 chooses $U$ or $M$, and all subsequent nodes (including Player 2's information set $\{n_1, n_2\}$) lie entirely inside this subtree and no information set is broken.

In that subgame, Player 1's relevant actions are $U$ and $M$, and Player 2's relevant actions are $L, C, R$ at $\{n_1, n_2\}$. Reading from Figure 5b (or from the first two rows of the table), we obtain that the Nash equilibria of the subgame are:

$$(In\text{-}U, Lr) \quad \text{and} \quad (I\text{-}M, Cr),$$

where the notation $Lr$ and $Cr$ simply records that we are embedding the subgame strategies $L$ and $C$ into full-game strategies $Lr$ and $Cr$ (the action at $n_3$ is irrelevant for the subgame).

Intuitively:

- If Player 1 plans to play $I$-$U$, then at $\{n_1, n_2\}$ Player 2 is best off choosing $L$ (hence $Lr$ at the level of full strategies).
- If Player 1 plans to play $I$-$M$, then at $\{n_1, n_2\}$ Player 2 is best off choosing $C$ (hence $Cr$).
- The action $R$ at $\{n_1, n_2\}$ is always (weakly) dominated for Player 2.

### *Which Profiles are SPE?*

To obtain subgame perfect equilibria, we must check which full strategy profiles generate a Nash equilibrium in the *whole* game and, at the same time, induce a Nash equilibrium in the new subgame.

We have two subgame candidates:

1. $((I\text{-}U), Lr)$

   For Player 2, given $I$-$U$, $Lr$ is a best response. For Player 1, given $Lr$, the payoffs from the four strategies are:

   $$u_1(I\text{-}U, Lr) = 4, \quad u_1(I\text{-}M, Lr) = 0, \quad u_1(D\text{-}U, Lr) = u_1(D\text{-}M, Lr) = 3.$$

   Hence $I$-$U$ is a best response to $Lr$. Thus $((I\text{-}U), Lr)$ is a Nash equilibrium of the whole game, and it is a Nash equilibrium of the subgame starting after *In*. Therefore it is an SPE.

2. $((I\text{-}M), Cr)$

   In the subgame, $(M, C)$ is a Nash equilibrium, so $((I\text{-}M), Cr)$ looks like a candidate SPE. However, at the level of the whole game, if Player 2 plays $Cr$ then Player 1's payoffs are

$$u_1(I\text{-}U, Cr) = 0, \quad u_1(I\text{-}M, Cr) = 2, \quad u_1(D\text{-}U, Cr) = u_1(D\text{-}M, Cr) = 3.$$

   Player 1 would deviate to $D$ (either $D$-$U$ or $D$-$M$), obtaining 3 instead of 2. Thus $((I\text{-}M), Cr)$ is *not* a Nash equilibrium of the whole game and therefore cannot be an SPE.

   Hence, in Figure 5b, the unique SPE is

$$\text{SPE: } ((I\text{-}U), Lr).$$

### *Connection Back to Figure 5a and Minimal Departure*

In Figure 5a we had an equilibrium in which Player 1 chooses $D$ and Player 2 threatens to play $R$ at the left-hand information set. From the perspective of Figure 5b, such a profile corresponds to a strategy of the form $(D\text{-}U, Cr)$ or $(D\text{-}M, Cr)$.

To test whether a profile like $(D\text{-}U, Cr)$ is SPE, we must look at the subgame that starts after *In*. Following the "minimal departure" idea:

$$(D\text{-}U)^h = I\text{-}U,$$

i.e. we only modify Player 1's strategy as little as possible so that the subgame is actually reached (replace $D$ by *In*, keep $U$). Given $Cr$ for Player 2, the induced profile in the subgame is

$$(I\text{-}U, Cr),$$

which is *not* a Nash equilibrium of that subgame. Therefore $(D\text{-}U, Cr)$ fails the SPE condition. The extra initial move $In/D$ makes the problematic threat visible at the level of a subgame, and subgame perfection rules it out.

### *What SPE Still Does Not Rule Out*

The key point is that SPE ties optimality to *subgames*. It may still fail to rule out some incredible threats when a player receives information and must act at an information set that is *not* the start of a subgame.

This motivates stronger refinements, such as *sequential equilibrium*, which tie optimization to information sets rather than subgames.

### *Forward Induction Intuition*

There is also a forward-induction argument in this example. If Player 2's information set $\{n_1, n_2\}$ is reached, she knows that Player 1 chose *In* instead of *D*, thereby giving up the sure possibility of obtaining 3 (via *D* and *r*) in order to enter. Thus Player 2 should infer that Player 1 expects to obtain at least 3 by going *In*, which in this game points toward the outcome in which Player 1 plays *I-U* (the only case in which he can get 4).

Iterated elimination of weakly dominated strategies in the normal form is consistent with this forward-induction reasoning: weakly dominated columns for Player 2 (such as $R\ell$, *Rc*, etc.) can be eliminated, then weakly dominated rows for Player 1 (such as *In-M*) can be eliminated, and so on, until only the profile

$$((In\text{-}U), Lr)$$

survives.

### *Realization-Equivalent Strategies*

For Figure 5b, note that *D-U* and *D-M* are *realization-equivalent strategies* for Player 1:

$$u_1(D\text{-}U, s_2) = u_1(D\text{-}M, s_2) \quad \text{for all } s_2 \in S_2.$$

In iterated weak dominance we are free to collapse realization-equivalent strategies into a single row, since they behave identically in the strategic form.

However, for subgame perfection we cannot collapse them: SPE is defined in terms of full strategies and their induced behavior in each subgame, so different strategies that implement the same realized path may still matter when we consider off-path (subgame) deviations.

This distinction between realization equivalence (for payoff-based reasoning) and full strategies (for subgame perfection and, later on, sequential equilibrium) is one of the central conceptual points of Class 15.

## *Mixed Strategies and Behavioral Strategies*

### *Example 3 Revisited*

We now return to Example 3 and use it to clarify the distinction between *mixed* and *behavioral* strategies.

Player 2 has two actions at each information set, so her pure strategies are

$$S_2 = \{(\ell, \ell), (\ell, r), (r, \ell), (r, r)\}.$$

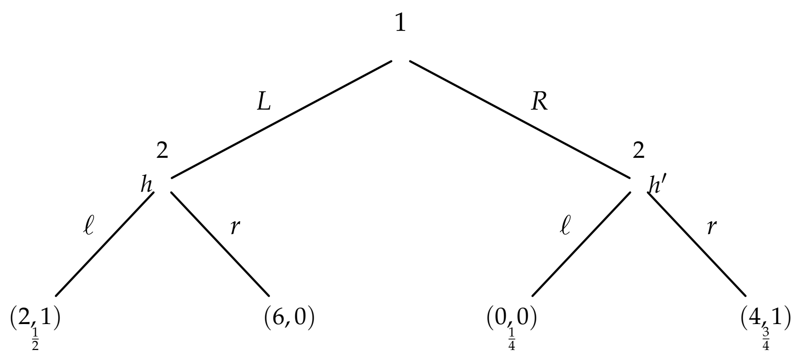


Example 3 revisited: Player 2 may mix differently at $h$ and $h'$.

*

The induced strategic form is

| | $(\ell,\ell)$ | $(\ell,r)$ | $(r,\ell)$ | $(r,r)$ |
|---|---|---|---|---|
| $L$ | $(2,1)$ | $(2,1)$ | $(6,0)$ | $(6,0)$ |
| $R$ | $(0,0)$ | $(4,1)$ | $(0,0)$ | $(4,1)$ |

A standard (ex-ante) mixed strategy for Player 2 is a probability distribution over $S_2$:

$$\sigma_2 \in \Delta(S_2), \qquad \sigma_2(\ell\ell),\ \sigma_2(\ell r),\ \sigma_2(r\ell),\ \sigma_2(rr) \geq 0,\ \sum_{s_2 \in S_2} \sigma_2(s_2) = 1.$$

However, in dynamic games it is often more natural to think that Player 2 *randomizes at each information set* $h$ and $h'$ separately, possibly in different ways. This leads to the notion of behavioral strategies.

## *Behavioral Strategies*

**Definition.** A *behavioral strategy* of player $i$ is a mapping

$$b_i : H_i \longrightarrow \Delta(A_i)$$

such that

$$b_i(h_i) \in \Delta\big(A_i(h_i)\big) \quad \text{for all } h_i \in H_i,$$

where $b_i(h_i)(a_i)$ is interpreted as the probability with which player $i$ chooses action $a_i \in A_i(h_i)$ when information set $h_i$ is reached.

Each behavioral strategy induces a mixed strategy over pure strategies.

*From behavioral to mixed.* Fix a behavioral strategy $b_i$ and recall that a pure strategy $s_i : H_i \to A_i$ specifies one action $s_i(h_i)$ at each information set. Define a mixed strategy $\sigma_i \in \Delta(S_i)$ by

$$\sigma_i(s_i) = \prod_{h_i \in H_i} b_i(h_i)\big(s_i(h_i)\big).$$

That is, $\sigma_i(s_i)$ is the probability that the behavioral strategy chooses, at every information set, the action prescribed by $s_i$.

For Example 3, suppose Player 2 uses the behavioral strategy

$$b_2(h)(\ell) = \frac{1}{2}, \quad b_2(h)(r) = \frac{1}{2}; \qquad b_2(h')(\ell) = \frac{1}{4}, \quad b_2(h')(r) = \frac{3}{4}.$$

Then the induced mixed strategy satisfies, for instance,

$$\sigma_2(\ell\ell) = b_2(h)(\ell)\, b_2(h')(\ell) = \frac{1}{2} \cdot \frac{1}{4} = \frac{1}{8}.$$

Thus every behavioral strategy induces a well-defined mixed strategy.

*But Not Every Mixed Strategy Comes from a Behavioral One*

Can we go the other way around? In general, no: a given mixed strategy over $S_2$ need not be generated by independent randomization at each information set.

Consider the mixed strategy

$$\sigma_2(\ell\ell) = \frac{1}{2}, \quad \sigma_2(rr) = \frac{1}{2}, \quad \sigma_2(\ell r) = \sigma_2(r\ell) = 0.$$

Suppose, for contradiction, that there exists a behavioral strategy $b_2 : H_2 \to \Delta(A_2)$ inducing $\sigma_2$.

To obtain $\sigma_2(\ell\ell) = \frac{1}{2}$ and $\sigma_2(rr) = \frac{1}{2}$ we need

$$b_2(h)(\ell) > 0, \quad b_2(h')(ell) > 0, \quad b_2(h)(r) > 0, \quad b_2(h')(r) > 0,$$

because otherwise one of the products $b_2(h)(\ell)b_2(h')(\ell)$ or $b_2(h)(r)b_2(h')(r)$ would be zero.

But then the probabilities of $(\ell r)$ and $(r\ell)$ are

$$\sigma_2(\ell r) = b_2(h)(\ell)b_2(h')(r) > 0, \qquad \sigma_2(r\ell) = b_2(h)(r)b_2(h')(\ell) > 0,$$

contradicting $\sigma_2(\ell r) = \sigma_2(r\ell) = 0$. Hence no behavioral strategy can induce this particular ex-ante mixed strategy.

The issue is that $\sigma_2$ correlates Player 2's actions at $h$ and $h'$: either she plays $(\ell, \ell)$ or $(r, r)$, never "cross" pairs. Behavioral strategies enforce independence across information sets.

*Equivalence at the Level of Terminal Histories*

Even though no behavioral strategy induces $\sigma_2$ *as a distribution over* $S_2$, it may still be possible to match the *distribution over terminal histories* generated by $(\sigma_1, \sigma_2)$.

Let $\sigma_2$ be as above and let Player 1 mix between $L$ and $R$ with

$$\sigma_1(L) = p, \qquad \sigma_1(R) = 1 - p.$$

Under $(\sigma_1, \sigma_2)$, each terminal node is reached with probability

$$Pr(z_1) = p \cdot \frac{1}{2}, \quad Pr(z_2) = p \cdot \frac{1}{2}, \quad Pr(z_3) = (1-p) \cdot \frac{1}{2}, \quad Pr(z_4) = (1-p) \cdot \frac{1}{2}.$$

Is there a behavioral-strategy profile that induces *the same* distribution over $(z_1, z_2, z_3, z_4)$?

Yes. Take

$$b_1(\text{root})(L) = p, \quad b_1(\text{root})(R) = 1 - p,$$

and let Player 2 randomize independently and symmetrically at each information set:

$$b_2(h)(\ell) = b_2(h)(r) = \frac{1}{2}, \qquad b_2(h')(\ell) = b_2(h')(r) = \frac{1}{2}.$$

A straightforward calculation shows that this behavioral profile generates exactly the same probabilities for $z_1, z_2, z_3, z_4$ as $(\sigma_1, \sigma_2)$. In other words, although the mixed strategy $\sigma_2$ cannot itself be written as the product of independent local randomizations, its effect on terminal outcomes can be replicated by a suitable behavioral strategy.

### *Kuhn's Theorem*

This example motivates the precise relationship between mixed and behavioral strategies captured by Kuhn's theorem.

**Theorem 0.0.24** (Kuhn)**.** *Fix a finite extensive-form game with no absent-mindedness and perfect recall. A probability distribution over terminal nodes can be induced by a mixed-strategy profile if and only if it can be induced by a behavioral-strategy profile.*

Thus, under perfect recall, for any profile of mixed strategies there exists a behaviorally-equivalent profile of behavioral strategies (and vice versa): they generate the same distribution over terminal histories, even though the underlying randomization devices—ex-ante mixing over $S_i$ versus independent randomization at information sets—need not coincide.

## *Appendix to Class 15: Game between Alice and Bob (Question 3, PS6)*

Consider a game between Alice and Bob. Alice moves first, choosing between $L$ and $R$. Bob then moves—without observing Alice's choice—and chooses between $\ell$ and $r$.

If Bob matches Alice's play (i.e., if either $(L, \ell)$ or $(R, r)$ is played), the game ends. If not, Bob observes Alice's choice and moves again:

- If $(L, r)$ was played, Bob chooses between $x$ and $y$.
- If $(R, \ell)$ was played, Bob chooses between $z$ and $q$.

The extensive-form payoffs (where the first number corresponds to Alice) are:

- If Alice plays $L$ and Bob plays $\ell$, payoffs are $(10, 9)$.
- If Alice plays $L$, Bob plays $r$ followed by $x$, payoffs are $(20, 0)$.
- If Alice plays $L$, Bob plays $r$ followed by $y$, payoffs are $(0, 10)$.
- If Alice plays $R$, Bob plays $\ell$ followed by $z$, payoffs are $(0, 20)$.
- If Alice plays $R$, Bob plays $\ell$ followed by $q$, payoffs are $(20, 10)$.
- If Alice plays $R$ and Bob plays $r$, payoffs are $(10, 0)$.

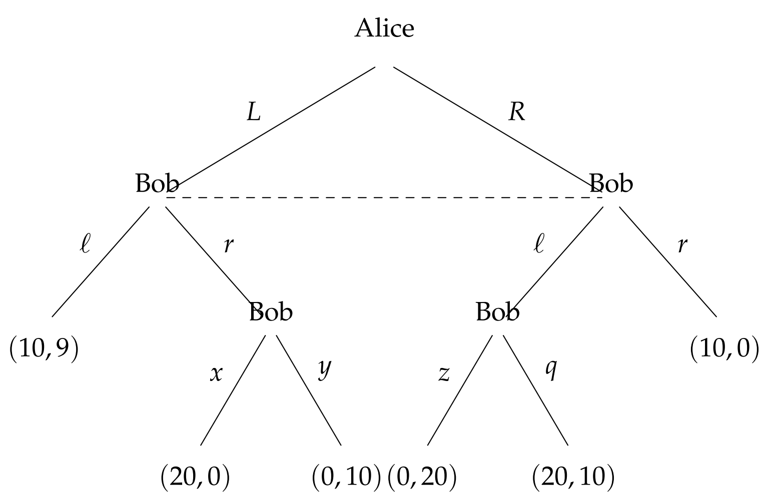


Extensive-form representation of the game

The game contains two proper subgames:

- The subgame that begins after the history $(L, r)$.
- The subgame that begins after the history $(R, \ell)$.

In the subgame following $(L, r)$, Bob chooses between $x$ and $y$, with payoffs

$$(20, 0) \text{ if } x, \qquad (0, 10) \text{ if } y.$$

Since $10 > 0$, Bob strictly prefers $y$. Hence, in any subgame perfect equilibrium, Bob must play $y$ after $(L, r)$.

In the subgame following $(R, \ell)$, Bob chooses between $z$ and $q$, with payoffs

$$(0, 20) \text{ if } z, \qquad (20, 10) \text{ if } q.$$

Since $20 > 10$, Bob strictly prefers $z$. Hence, in any subgame perfect equilibrium, Bob must play $z$ after $(R, \ell)$.

Therefore, continuation play in both proper subgames is uniquely determined: Bob plays $y$ after $(L, r)$ and $z$ after $(R, \ell)$.

Substituting these continuation strategies, the initial interaction reduces to the following strategic-form game:

| | $(\ell, y, z)$ | $(r, y, z)$ |
|---|---|---|
| $L$ | $(10, 9)$ | $(0, 10)$ |
| $R$ | $(0, 20)$ | $(10, 0)$ |

Best responses are:

- If Bob plays $\ell$, Alice prefers $L$.
- If Bob plays $r$, Alice prefers $R$.
- If Alice plays $L$, Bob prefers $r$.
- If Alice plays $R$, Bob prefers $\ell$.

No pure strategy profile satisfies mutual best responses. Therefore:

**There is no subgame perfect equilibrium in pure strategies.**

Let

- $q$ denote the probability that Alice plays $L$;
- $p$ denote the probability that Bob plays $\ell$ at his first information set.

Continuation strategies are fixed: Bob plays $y$ after $(L, r)$ and $z$ after $(R, \ell)$.

*Alice's indifference condition.* If Alice plays $L$:

$$u_A(L) = 10p.$$

If Alice plays $R$:

$$u_A(R) = 10(1 - p).$$

Indifference requires

$$10p = 10(1 - p) \quad \Longrightarrow \quad p = \frac{1}{2}.$$

*Bob's indifference condition.* If Bob plays $\ell$:

$$u_B(\ell) = 9q + 20(1-q) = 20 - 11q.$$

If Bob plays $r$:

$$u_B(r) = 10q.$$

Indifference requires

$$20 - 11q = 10q \quad \Longrightarrow \quad 21q = 20 \quad \Longrightarrow \quad q = \frac{20}{21}.$$

The unique subgame perfect equilibrium in behavioral strategies is:

- Alice plays $L$ with probability $\frac{20}{21}$ and $R$ with probability $\frac{1}{21}$.
- At his first information set, Bob plays $\ell$ with probability $\frac{1}{2}$ and $r$ with probability $\frac{1}{2}$.
- After $(L, r)$, Bob plays $y$ with probability 1.
- After $(R, \ell)$, Bob plays $z$ with probability 1.

This equilibrium is subgame perfect because continuation play is optimal in every proper subgame and the initial strategies constitute a Nash equilibrium of the reduced game.