Further
Solution Concepts

CMPUT 654: Modelling Human Strategic Behaviour

S&LB §3.4

Recap: Pareto Optimality

Definition: Outcome o **Pareto dominates** o' if

1. $\forall i \in N : o \succeq_i o',$ and
2. $\exists i \in N : o \succ_i o'.$

Equivalently, **action profile** a Pareto dominates a' if $u_i(a) \geq u_i(a')$ for all $i \in N$ and $u_i(a) > u_i(a')$ for some $i \in N$.

Definition: An outcome o^* is **Pareto optimal** if no other outcome Pareto dominates it.

Recap: Best Response and Nash Equilibrium

Definition:
The set of i's **best responses** to a strategy profile $s_{-i} \in S_{-i}$ is

$$BR_i(s_{-i}) \doteq \{s_i^* \in S_i \mid u_i(s_i^*, s_{-i}) \geq u_i(s_i, s_{-i}) \ \ \forall s_i \in S_i\}$$

Definition:
A strategy profile $s \in S$ is a **Nash equilibrium** iff

$$\forall i \in N: \quad s_i \in BR_{-i}(s_{-i})$$

- When at least one s_i is mixed, s is a **mixed strategy Nash equilibrium**
- When every s_i is deterministic, s is a **pure strategy Nash equilibrium**

Lecture Outline

1. Recap & Logistics
2. Maxmin Strategies
3. Dominated Strategies
4. Rationalizability
5. ϵ-Nash Equilibrium
6. Correlated Equilibrium

Maxmin Strategies

Question:

Why would an agent want to play a maxmin strategy?

What is the maximum amount that an agent can **guarantee** in expectation?

Definition:

A **maxmin strategy** for i is a strategy $\bar{s}_i$ that maximizes i's worst-case payoff:

$$\bar{s}_i = \arg\max_{s_i \in S_i} \left[\min_{s_{-i} \in S_{-i}} u_i(s_i, s_{-i}) \right]$$

Definition:

The **maxmin value** of a game for i is the value $\bar{v}_i$ guaranteed by a maxmin strategy:

$$\bar{v}_i = \max_{s_i \in S_i} \left[\min_{s_{-i} \in S_{-i}} u_i(s_i, s_{-i}) \right]$$

Minmax Strategies

Question:

Why would an agent want to play a minmax strategy?

The corresponding strategy for the other player is the minmax strategy: the strategy that **minimizes the other player's** payoff.

Definition: (two-player games)
In a two-player game, the **minmax strategy** for player i against player $-i$ is

$$\underline{s}_i = \arg\min_{s_i \in S_i} \left[\max_{s_{-i} \in S_{-i}} u_{-i}(s_i, s_{-i}) \right].$$

Definition: (n-player games)
In an n-player game, the **minmax strategy** for player i against player $j \neq i$ is i's component of the mixed strategy profile $\underline{s}_{(-j)}$ in the expression

$$\underline{s}_{(-j)} = \arg\min_{s_{-j} \in S_{-j}} \left[\max_{s_j \in S_j} u_j(s_j, s_{-j}) \right],$$

and the **minmax value** for player j is $\underline{v}_j = \min_{s_{-j} \in S_{-j}} \max_{s_j \in S_j} u_j(s_j, s_{-j})$.

Minimax Theorem

Theorem: [von Neumann, 1928]
In any finite, two-player, zero-sum game, in any Nash equilibrium $s^* \in S$, each player receives an expected utility v_i equal to both their maxmin and their minmax value.

Minimax Theorem Proof

Proof sketch:

1. Suppose that $v_i < \overline{v}_i$. But then i could guarantee a higher payoff by playing their maxmin strategy. So $v_i \geq \overline{v}_i$.

2. $-i$'s equilibrium payoff is $v_{-i} = \max_{s_{-i}} u_{-i}(s_i^*, s_{-i})$.

3. Equivalently, $v_i = \min_{s_{-i}} u_i(s_i^*, s_{-i})$. (**why?**)

4. So $v_i = \min_{s_{-i}} u_i(s_i^*, s_{-i}) \leq \max_{s_i} \min_{s_{-i}} u_i(s_i, s_{-i}) = \overline{v}_i$.

5. So $\overline{v}_i \leq v_i \leq \overline{v}_i$. ∎

Zero-sum game, so

$$v_{-i} = -v_i$$

$$\max_{s_{-i}} u_{-i}(s_i^*, s_{-i}) = \max_{s_{-i}} -u_i(s_i^*, s_{-i})$$

$$\max_{s_{-i}} -u_i(s_i^*, s_{-i}) = -\min_{s_{-i}} u_i(s_i^*, s_{-i})$$

Minimax Theorem Implications

In any **zero-sum** game:

1. Each player's maxmin value is equal to their minmax value.
   We call this the **value of the game**.

2. For both players, the maxmin strategies and the Nash equilibrium strategies are the same sets.

3. Any **maxmin strategy profile** (a profile in which both agents are playing maxmin strategies) is a Nash equilibrium. Therefore, each player gets the same payoff in every Nash equilibrium (namely, their value for the game).

Corollary: There is no **equilibrium selection** problem.

Dominated Strategies

When can we say that one strategy is **definitely** better than another, from an **individual's** point of view?

Definition: (domination)
Let $s_i, s_i' \in S_i$ be two of player i's strategies. Then

1. s_i **strictly dominates** s_i' if $\forall s_{-i} \in S_{-i} : u_i(s_i, s_{-i}) > u_i(s_i', s_{-i})$.

2. s_i **weakly dominates** s_i' if $\forall s_{-i} \in S_{-i} : u_i(s_i, s_{-i}) \geq u_i(s_i', s_{-i})$ and $\exists s_{-i} \in S_{-i} : u_i(s_i, s_{-i}) > u_i(s_i', s_{-i})$.

3. s_i **very weakly dominates** s_i' if $\forall s_{-i} \in S_{-i} : u_i(s_i, s_{-i}) \geq u_i(s_i', s_{-i})$.

Dominant Strategies

Definition:
A strategy is (strictly, weakly, very weakly) **dominant** if it (strictly, weakly, very weakly) dominates **every** other strategy.

Definition:
A strategy is (strictly, weakly, very weakly) **dominated** if is is (strictly, weakly, very weakly) dominated by **some** other strategy.

Definition:
A strategy profile in which every agent plays a (strictly, weakly, very weakly) dominant strategy is an **equilibrium in dominant strategies**.

Questions:

1. Are **dominant strategies** guaranteed to exist?
2. What is the maximum number of **weakly** dominant strategies?
3. Is an equilibrium in dominant strategies also a **Nash equilibrium**?

Prisoner's Dilemma again

	Coop.	Defect
Coop.	-1,-1	-5,0
Defect	0,-5	-3,-3

- *Defect* is a **strictly dominant** pure strategy in Prisoner's Dilemma.
    - *Cooperate* is **strictly dominated**.
- **Question:** Why would an agent want to play a **strictly dominant** strategy?
- **Question:** Why would an agent want to play a **strictly dominated** strategy?

Battle of the Sofas

	Ballet	Soccer	Home
Ballet	2,1	0,0	1,0
Soccer	0,0	1,2	0,0
Home	0,0	0,1	1,1

- What are the **dominated** strategies?
  - *Home* is a **weakly dominated** pure strategy in Battle of the Sofas.
- **Question:** Why would an agent want to play a **weakly dominated** strategy?

Fun Game:
Traveller's Dilemma

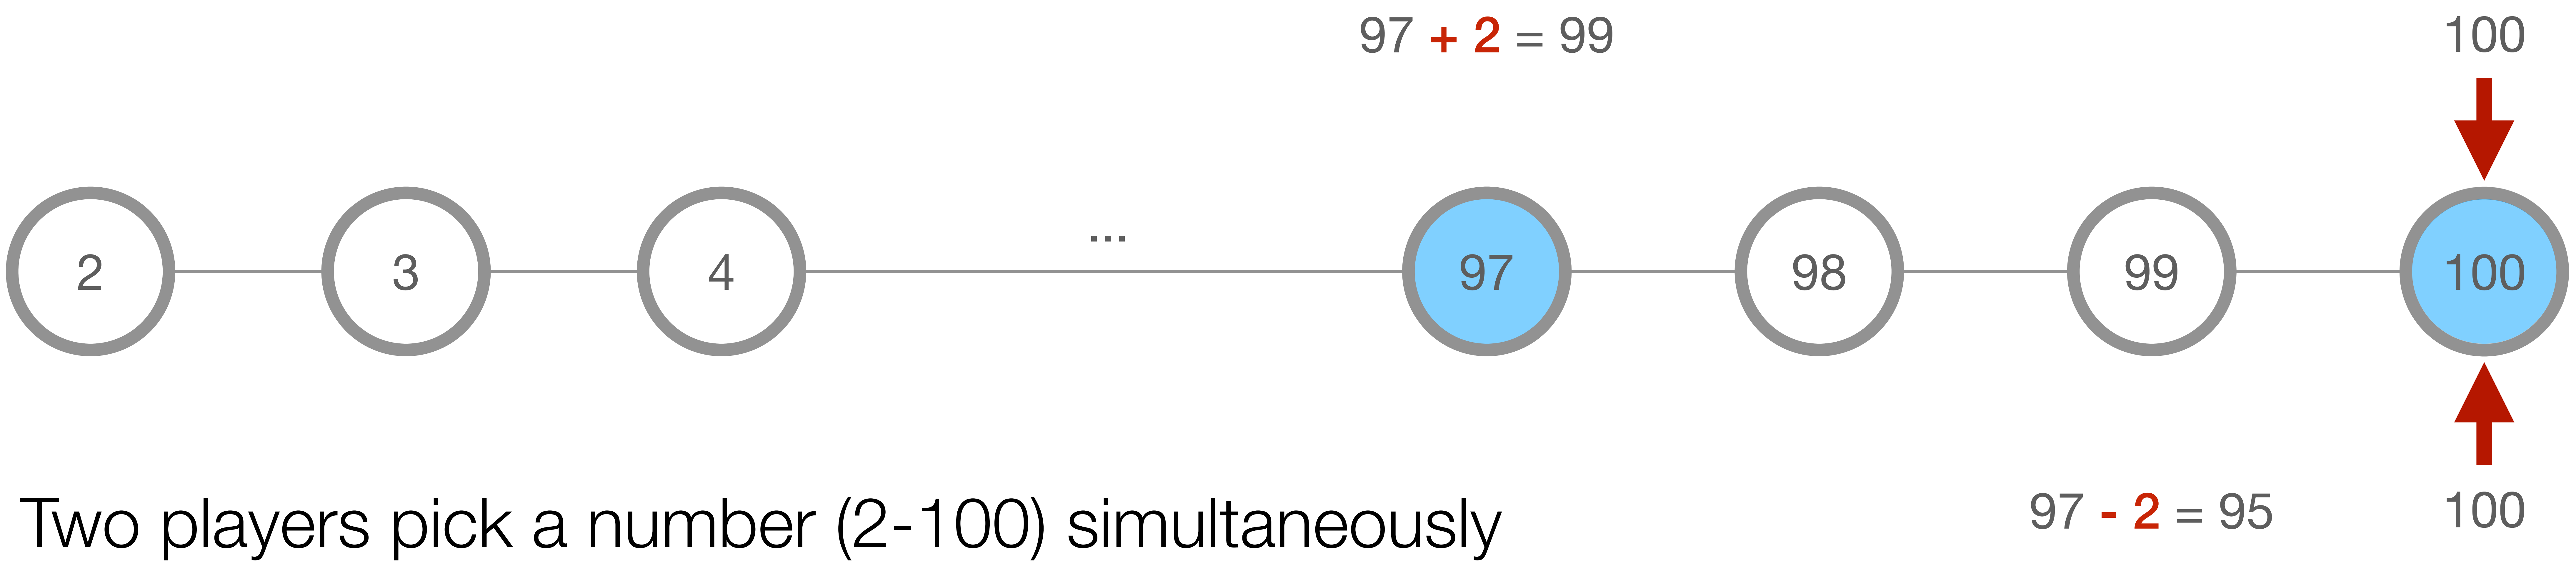

- Two players pick a number (2-100) simultaneously
- If they pick the same number *x*, then they both get $*x* payoff
- If they pick different numbers:
  - Player who picked lower number gets **lower** number, plus **bonus** of $2
  - Player who picked higher number gets **lower** number, minus **penalty** of $2
- Play against someone near you, three times in total. Keep track of your payoffs!

Traveller's Dilemma

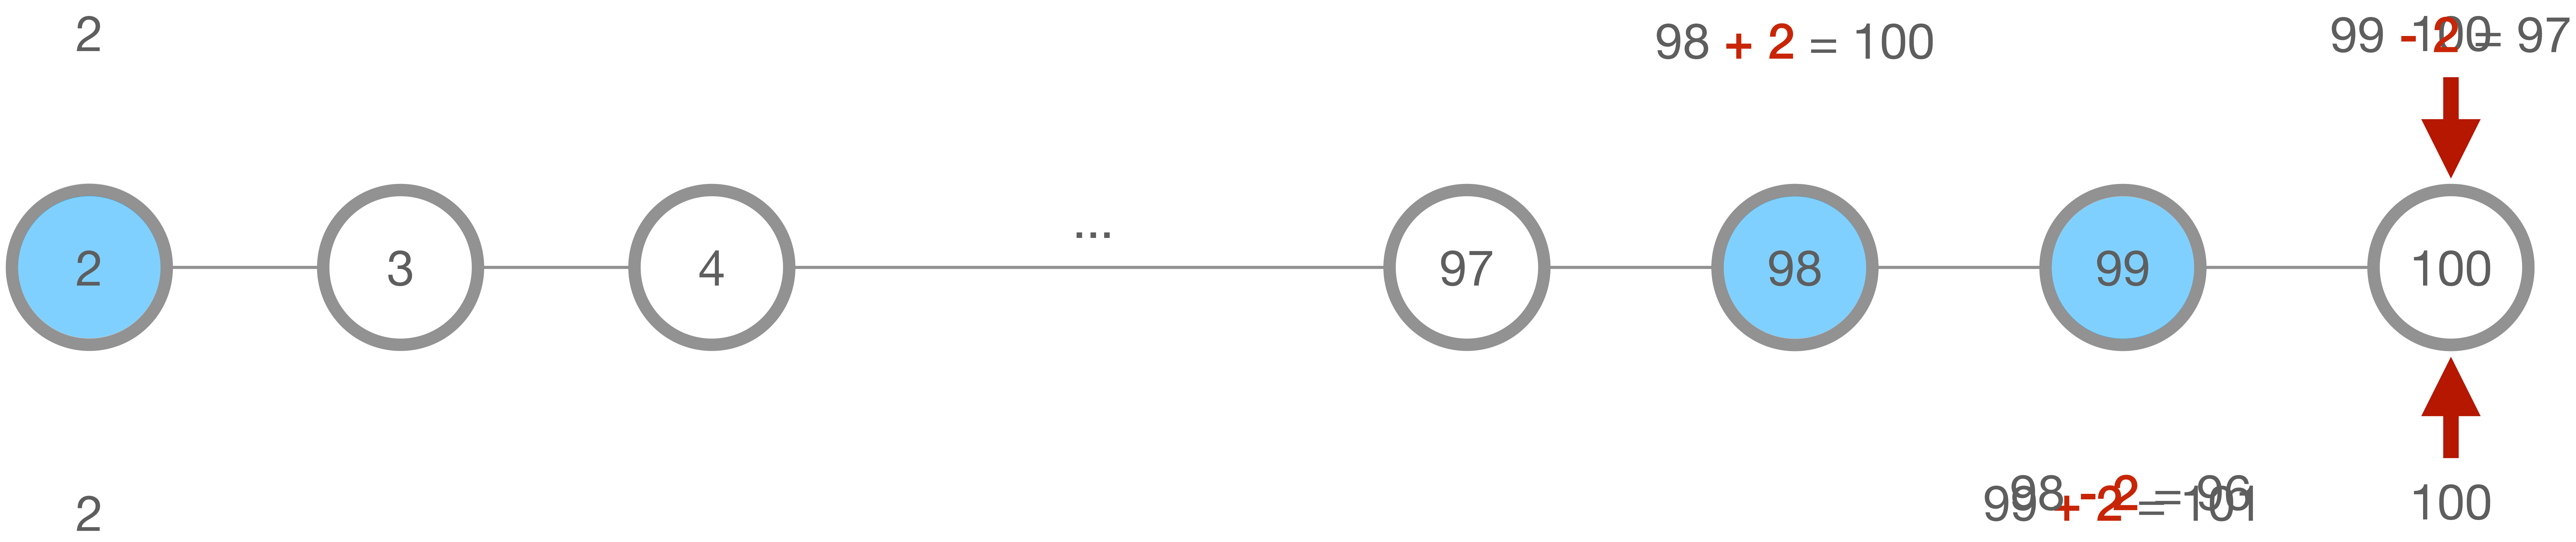

- Traveller's Dilemma has a unique Nash equilibrium

Iterated Removal of Dominated Strategies

- No **strictly dominated** pure strategy will ever be played by a fully rational agent.
- So we can remove them, and the game remains **strategically equivalent**
- But! Once you've removed a dominated strategy, another strategy that wasn't dominated before might **become dominated** in the new game.
  - It's safe to remove this newly-dominated action, because it's never a best response *to an action that the opponent would ever play*.
- You can repeat this process until there are no dominated actions left

Iterated Removal of Dominated Strategies

	Ballet	Soccer	Home
Ballet	2,1	0,0	1,0
Soccer	0,0	1,2	0,0
Home	0,0	0,1	1,1

- Removing **strictly dominated** strategies preserves **all equilibria.** (**Why?**)
- Removing weakly or very weakly dominated strategies **may not** preserve **all equilibria**. (**Why?**)
- Removing weakly or very weakly dominated strategies preserves **at least one equilibrium**. (**Why?**)
  - But because not all equilibria are necessarily preserved, the **order** in which strategies are removed can **matter**.

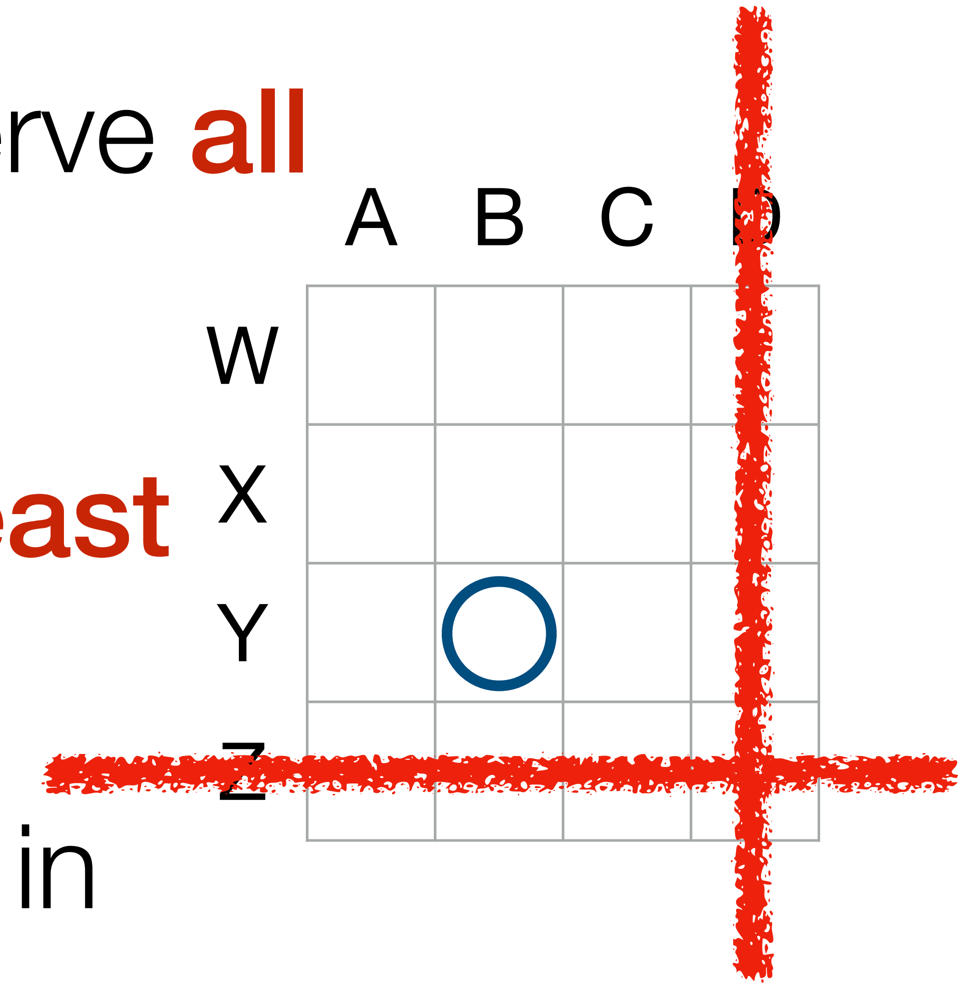

Nash Equilibrium Beliefs

One characterization of Nash equilibrium:

1. **Rational behaviour:**
   Agents maximize expected utility with respect to their beliefs.
2. **Rational expectations:**
   Agents have **accurate** probabilistic beliefs about the behaviour of the other agents.

Rationalizability

- We saw in the utility theory lecture that rational agents' **beliefs** need not be **objective** (or accurate)
- What strategies could possibly be played by:
  1. A **rational** player...
  2. ...with **common knowledge** of the rationality of **all players?**
- Any strategy that is a best response to **some beliefs consistent with** these two conditions is **rationalizable**.

Questions:

1. What kind of strategy definitely could **not** be played by a rational player with common knowledge of rationality?
2. Is a rationalizable strategy guaranteed to exist?
3. Can a game have more than one rationalizable strategy?

ε-Nash Equilibrium

- In a Nash equilibrium, agents best respond **perfectly**
- What if they are indifferent to **very small** gains in utility?
  - Could reflect modelling error (e.g., unmodelled cost of computational effort)

Definition:
For any $\varepsilon > 0$, a strategy profile s is an **ε-Nash equilibrium** if, for all agents i and strategies $s'_i \neq s_i$,

$$u_i(s_i, s_{-i}) \geq u_i(s'_i, s_{-i}) - \epsilon.$$

Questions:

For a given $\epsilon > 0$,

1. Is an ε-Nash equilibrium guaranteed to **exist**?
2. Is **more than one** ε-Nash equilibrium **guaranteed** to exist?

ε-Nash Equilibrium Example

	L	R
U	1, 1	0, 0
D	$1+(\varepsilon/2)$, 1	500, 500

Questions:

1. What are the **Nash equilibria** of this game?
2. What are the **ε-Nash equilibria** of this game?

- Every Nash equilibrium is surrounded by a **region** of ε-Nash equilibria
  - Every **numerical algorithm** for computing Nash equilibrium actually computes ε-Nash equilibrium
- However, the reverse is not true! Payoffs from an ε-Nash equilibrium can be **arbitrarily far** from Nash equilibrium payoffs.

Correlated Equilibrium Examples

	Ballet	Soccer
Ballet	2, 1	0, 0
Soccer	0, 0	1, 2

	Go	Wait
Go	-10, -10	1, 0
Wait	0, 1	-1, -1

- In the unique mixed strategy equilibrium of Battle of the Sexes, each player gets a utility of 2/3
- If the players could first observe a coin flip, they could coordinate on **which** pure strategy equilibrium to play
  - Each would get utility of 1.5
  - **Fairer** than either pure strategy equilibrium, and **Pareto dominates** the mixed strategy equilibrium
- **Correlated equilibrium** is a solution concept in which agents get private, potentially-correlated **signals** before choosing their action
  - In both of these example, each agent sees the **same signal** perfectly, but that is not necessary in general

Correlated Equilibrium

Definition:
Given an n-agent game $G = (N, A, u)$, a **correlated equilibrium** is a tuple (v, π, σ), where

$v = (v_1, \ldots, v_n)$ is a tuple of random variables with domains $(D_1, \ldots, D_n)$,

π is a joint distribution over v,

$\sigma = (\sigma_1, \ldots, \sigma_n)$ is a vector of mappings $\sigma_i : D_i \to A_i$, and

for every agent i and mapping $\sigma' : D_i \to A_i$,

$$\sum_{d \in D_1 \times \cdots \times D_n} \pi(d) u_i(\sigma_1(d_1), \ldots, \sigma_n(d_n)) \geq \sum_{d \in D_1 \times \cdots \times D_n} \pi(d) u_i(\sigma_1(d_1), \ldots, \sigma_i'(d_i), \ldots, \sigma_n(d_n))$$

Question: Why do the σ_i's map to **pure strategies** instead of mixed strategies?

Correlated Equilibrium Properties

Theorem:
For every **Nash equilibrium**, there exists a corresponding correlated equilibrium in which each action profile appears with the same frequency. (**how?**)

Theorem:
Any **convex combination** of correlated equilibrium payoffs can be realized in some correlated equilibrium. (**how?**)

Correlated Equilibrium
Another Example

	Ballet	Soccer
Ballet	2, 1	0, 0
Soccer	0, 0	1, 2

	Go	Wait
Go	-10, -10	1, 0
Wait	0, 1	-1, -1

- In our example correlated equilibria, each agent best-responded to the other **at every signal**
  - This is **not a requirement** of a correlated equilibrium
- Consider this correlated equilibrium, with $D_1 = \{x, y, z\}$ and $D_2 = \{m, r\}$:

$$\pi\left[(x,m)\right] = .25 \qquad \sigma_r(x) = X \qquad \sigma_c(m) = M$$
$$\pi\left[(y,m)\right] = .25 \qquad \sigma_r(y) = Y$$
$$\pi\left[(z,r)\right] = .5 \qquad \sigma_r(z) = Z \qquad \sigma_r(r) = R$$

	L	M	R
X	0,8	3,6	-9,1
Y	0,2	3,9	-12,10
Z	1,0	0,-2	7,7

  - **Question:** Does the column player best-respond at each signal?
  - **Question:** What are the **marginal probabilities** for each player's actions?
  - **Question:** What would happen if the agents played **mixed strategies** with those marginal probabilities?

Summary

- **Maxmin strategies** maximize an agent's **guaranteed payoff**
- **Minmax strategies** minimize the other agent's payoff as much as possible
- The **Minimax Theorem**:
  - Maxmin and minmax strategies are the **only** Nash equilibrium strategies in **zero-sum games**
  - Every Nash equilibrium in a zero-sum game has the **same payoff**
- **Dominated strategies** can be removed **iteratively** without strategically changing the game (too much)
- **Rationalizable** strategies are any that are a **best response** to some **rational belief**
- ϵ**-Nash equilibria:** stable when agents have no deviation that gains them more than ε
- **Correlated equilibria:** stable when agents have **signals** from a possibly-correlated randomizing device