# From Optimality to Equilibrium

(Based on Shoham and Leyton-Brown (2008). *Multiagent Systems: Algorithmic, Game-Theoretic, and Logical Foundations*, Cambridge.)

Leen-Kiat Soh

# Introduction

Why or why not?

- Recall Prisoner's Dilemma:

| | Player 2 No Betray | Player 2 Betray |
|---|---|---|
| Player 1 No Betray | **1,1** | -4,3 |
| Player 1 Betray | 3,-4 | **-3,-3** |

**Can an agent afford to pursue an "optimal" strategy?**

# Introduction

Why or why not?

- Recall Prisoner's Dilemma:

| | Player 2 No Betray | Player 2 Betray |
|---|---|---|
| **Player 1 No Betray** | **1,1** | -4,3 |
| **Player 1 Betray** | 3,-4 | **-3,-3** |

**Can an agent afford to pursue an "optimal" strategy?**

# Introduction

Implications for agent reasoning? How does this impact autonomy? How does this impact MAS design?

- In single-agent decision theory the key notion is that of an *optimal strategy*, that is, a strategy that maximizes the agent's expected payoff for a given environment in which the agent operates
  - Complications: incompleteness, non-deterministic, and dynamic (e.g., other agents at work)
- **Important. The notion of an optimal strategy for a given agent is not meaningful: the best strategy depends on the choices of others**
- Game theorists deal with this problem by identifying certain subsets of outcomes, called *solution concepts*:
  - **Two examples of most fundamental: Pareto optimality** and **Nash equilibrium**.

# Pareto Optimality

Implications for agent reasoning? How does this impact autonomy? How does this impact MAS design?

- **Definition 3.3.1 Pareto Domination.** *Strategy profile $s$* Pareto dominates *strategy profile $s'$ if for all $i \in N, u_i(s) \geq u_i(s')$, and there exists some $j \in N$ for which $u_j(s) > u_j(s')$.*
  - In other words, in a Pareto-dominated strategy profile **some player can be made better off without making any other player worse off**
  - Pareto domination gives us a *partial* ordering over strategy profiles
    - May not have a single "best" outcome; instead, we may have a set of noncomparable optima.
- **Definition 3.3.2 Pareto Optimality.** *Strategy profile $s$ is* Pareto optimal, *or* **strictly Pareto efficient**, *if there does not exist another strategy profile $s' \in S$ that Pareto dominates $s$.*

# Best Response

What is a **mixed strategy**? A probabilistic distribution over actions that an agent chooses to do in a strategy

- Let's look at the game from **an individual agent's point of view, rather than from the vantage point of an outside observer**
- ***Intuition***: If an agent knew how the others were going to play, his or her strategic problem would become simple
  - Specifically, he or she **would be left with the single-agent problem of choosing a utility-maximizing action**!
- Formally, define $s_{-i} = (s_1, \ldots, s_{i-1}, s_{i+1}, \ldots, s_n)$, a strategy profile $s$ without agent $i$'s strategy. Thus we can write $s = (s_i, s_{-i})$.
- If the agents other than $i$ (whom we denote $-i$) were to commit to play $s_{-i}$, a utility-maximizing agent $i$ would face the problem of determining his or her best response.
- **Definition 3.3.3 Best Response.** *Player $i$'s* best response *to the strategy profile $s_{-i}$ is a mixed strategy $s_i^* \in S_i$ such that $u_i(s_i^*, s_{-i}) \geq u_i(s_i, s_{-i})$ for all strategies $s_i \in S_i$.*

# Best Response

- **Not a solution concept because an agent does not know what other agents' choices are!**
- However, we can leverage the idea of best response to define what is arguably the most central notion in noncooperative game theory, the **Nash equilibrium**

👏👏👏👏👏👏👏👏

# Nash Equilibrium

- **Definition 3.3.4 Nash Equilibrium.** *A strategy profile* $s = (s_1, \ldots, s_n)$ *is a* Nash equilibrium *if, for **all** agents* $i$, $s_i$ *is a **best response** to* $s_{-i}$.
- **Intuitively, a Nash equilibrium is a *stable* strategy profile**
  - **No agent would want to change his or her strategy if he or she knew what strategies the other agents were following.**
- **Definition 3.3.5 (Strict Nash)** *A strategy profile* $s = (s_1, \ldots, s_n)$ *is a* strict Nash equilibrium *if, for all agents* $i$ *and for all strategies* $s_i' \neq s_i$, $u_i(s_i, s_{-i}) > u_i(s_i', s_{-i})$.
- **Definition 3.3.6 (Weak Nash)** *A strategy profile* $s = (s_1, \ldots, s_n)$ *is a* weak Nash equilibrium *if, for all agents* $i$ *and for all strategies* $s_i' \neq s_i$, $u_i(s_i, s_{-i}) \geq u_i(s_i', s_{-i})$.

# Connection to MAS?

Optimality may not be possible in a multi-agent environment → Pareto Optimality, Best Response

Nash equilibrium → solution stability

Complex! Simulations/multiple runs, applications to adversarial games (poker, trading, security)

Racecars drive around and around an oval track. What if when drivers switch lanes, they are required to turn on their signals? What would you do? Would there be a best response strategy? An equilibrium?