

CSCE 475/875 Multiagent Systems
Handout 6: From Optimality to Equilibrium
January 28, 2020
(Based on Shoham and Leyton-Brown 2011)

Introduction

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. The situation in the single-agent case can be fraught with uncertainty, since the environment might be stochastic, partially observable, and spring all kinds of surprises on the agent. However, the situation is even more complex in a multiagent setting.

Important. Thus 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*, that are interesting in one sense or another. Here we describe two of the most fundamental solution concepts: **Pareto optimality** and **Nash equilibrium**.

Pareto Optimality

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. Thus, we cannot generally identify 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 and Nash Equilibrium

Now we will look at games from an **individual agent's point of view, rather than from the vantage point of an outside observer**.

Intuition: Our first observation is that 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 $s_{-i} = (s_1, \dots, s_{i-1}, s_{i+1}, \dots, s_n)$, 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 $s_i \in S_i$.*

The best response is not necessarily unique. Further, when the support of a best response s^* includes two or more actions, the agent must be indifferent among them—otherwise, the agent would prefer to reduce the probability of playing at least one of the actions to zero. Thus, similarly, if there are two pure strategies that are individually best responses, any mixture of the two is necessarily also a best response.

Important: Of course, in general an agent will not know what strategies the other players plan to adopt. Thus, the notion of best response is not a solution concept—it does not identify an interesting set of outcomes in this general case.

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**.

Definition 3.3.4 (Nash equilibrium) *A strategy profile $s = (s_1, \dots, s_n)$ 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 $s = (s_1, \dots, s_n)$ 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 $s = (s_1, \dots, s_n)$ is a weak Nash equilibrium if, for all agents i and for all strategies $s_i' \neq s_i$, $u_i(s_i, s_{-i}) \geq u_i(s_i', s_{-i})$.*