Economics, AI, and Optimization
Lecture Note 2: Intro to Game Theory

Christian Kroer*

January 23, 2022

1 Nash Equilibrium

In this lecture we begin our study of Nash equilibrium. First we will learn the basic definitions, and then we will get started on regret minimization, which will be an essential tool for will computing Nash equilibrium later.

1.1 General-Sum Games

A *normal-form game* consists of:

- A set of players $N = \{1, \dots, n\}$
- A set of strategies $S = S_1 \times S_2 \times \cdots \times S_n$
- A utility function $u_i : S \to \mathbb{R}$

We will use the shorthand s_{-i} to denote the subset of a strategy vector s that does not include player i's strategy.

As a first solution concept we will consider *dominant-strategy equilibrium* (DSE). In DSE, we seek a strategy vector $s \in S$ such that each s_i is a best response *no matter what* s_{-i} *is.* A classic example is the *prisoner's dilemma*: two prisoners are on trial for a crime. If neither confesses (stay silent) to the crime then they will each get 1 year in prison. If one person confesses and the other does not, then the confessor gets no time, but their co-conspirator gets 9 years. If both confess then they both get 6 years.

	Silent	Confess
Silent	-1,-1	-9,0
Confess	0,-9	-6,-6

In this game, confessing is a DSE: it yields greater utility than staying silent no matter what the other player does. A DSE rarely exists in practice, but it can be useful in the context of mechanism design, where we get to decide the rules of the game. It is the idea underlying e.g. the second-price auction which we will cover later.

*Department of Industrial Engineering and Operations Research, Columbia University. Email: christian.kroer@columbia.edu.

Consider some strategy vector $s \in S$. We say that s is a *pure-strategy Nash equilibrium* if for each player i and each alternative strategy $s'_i \in S_i$:

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

where s_{-i} denotes all the strategies in s except that of i. A DSE is always a pure-strategy Nash equilibrium, but not vice versa. Consider the *Professor's dilemma*,[1] where the professor chooses a row strategy and the students choose a column strategy:

		Students	
		Listen	Sleep
Prof.	Prepare	$10^6, 10^6$	-10,0
	Slack off	0,-10	0,0

In this game there is no DSE, but there's clearly two pure-strategy Nash equilibria: the professor prepares and students listen, or the professor slacks off and students sleep. But these have quite different properties. Thus equilibrium selection is an issue for general-sum games. There are at least two reasons for this: first, if we want to predict the behavior of players then how do we choose which equilibrium to predict? Second, if we want to prescribe behavior for an individual player, then we cannot necessarily suggest that they player some particular strategy from a Nash equilibrium, because if the others player do not play the same Nash equilibrium then it may be a terrible suggestion.

Moreover, pure-strategy equilibria are not even guaranteed to exist, as we saw in the previous lecture with the rock-paper-scissors example.

To fix the existence issue we may consider allowing players to randomize over their choice of strategy (as in rock-paper-scissors where players should randomize uniformly). Let $\sigma_i \in \Delta^{|S_i|}$ denote player i's probability distribution over their strategy, this is called a *mixed strategy*. Let a strategy profile be denoted by $\sigma = (\sigma_1, \dots, \sigma_n)$. By a slight abuse of notation we may rewrite a player's utility function as

$$u_i(\sigma) = \sum_{s \in S} u_i(s) \prod_i \sigma_i(s_i)$$

A (mixed-strategy) Nash equilibrium is a strategy profile σ such that for all pure strategies σ'_i (σ'_i is pure if it puts probability 1 on a single strategy):

$$u_i(\sigma) \geq u_i(\sigma_{-i}, \sigma'_i).$$

Now, Nash's theorem says that

Theorem 1. *Any game with a finite set of strategies and a finite set of players has a mixed-strategy Nash equilibrium.*

Now, since our goal is the prescribe or predict behavior, we would also like to be able to compute a Nash equilibrium. Unfortunately this turns out to be computationally difficult:

Theorem 2. *The problem of computing a Nash equilibrium in general-sum finite games is PPAD-complete.*

[1]Example borrowed from Ariel Procaccia's slides

We won't go into detail on what the complexity class PPAD is for now, but suffice it to say that it is weaker than the class of NP-complete problems (it is not hard to come up with a MIP for computing a Nash equilibrium, for example), but still believed to take exponential time in the worst case.

As a sidenote, one may make the following observation about why Nash equilibrium does not "fit" in the class of NP-complete problems: typically in NP-completeness we ask questions such as "does there exist a satisfying assignment to this Boolean formula?" But given a particular game, we already know that a Nash equilibrium exists. Thus we cannot ask about the type of existence questions typically used in NP-complete problems, but rather it is only the task of finding one of the solutions that is difficult. This can be a useful notion to keep in mind when encountering other problems that have guaranteed existence. That said, once one asks for additional properties such as "does there exist a Nash equilibrium where the sum of utilities is at least v?" one gets an NP-complete problem.

Given a strategy profile σ, we will often be interested in measuring how "happy" the players are with the outcome of the game under σ. Most commonly, we are interested in the *social welfare* of a strategy profile (and especially for equilibria). The social welfare is the expected value of the sum of the player's utilities:

$$\sum_{i=1}^{n} u_i(\sigma) = \sum_{i=1}^{n} \sum_{s \in S} u_i(s) \prod_{i'=1}^{n} \sigma_{i'}(s_{i'}).$$

We already saw in the Professor's Dilemma that there can be multiple equilibria with wildly different social welfare: when the professor slacks off and the students sleep, the social welfare is zero; when the professor prepares and the students listen, the social welfare is $2 \cdot 10^6$.

1.2 Zero-Sum Games

In the special case of a two-player zero-sum game, we have $u_1(s) = -u_2(s) \forall s \in S$. In that case, we can represent our problem as the bilinear saddlepoint problem we saw in the last lecture:

$$\min_{x \in \Delta^n} \max_{y \in \Delta^m} \langle x, Ay \rangle.$$

A first observation one may make is that the minimization problem faced by the x-player is a convex optimization problem, since the max operation is convexity-preserving. This suggests that we should have a lot of algorithmic options to use. This turns out to be true: unlike the general case, we can compute a zero-sum equilibrium in polynomial time using linear programming (LP).

In fact, we have the following stronger statement, which is essentially equivalent to LP duality:

Theorem 3 (von Neumann's minimax theorem)**.** *Every two-player zero-sum game has a unique value v, called the* value of the game*, such that*

$$\min_{x \in \Delta^n} \max_{y \in \Delta^m} \langle x, Ay \rangle = \max_{y \in \Delta^m} \min_{x \in \Delta^n} \langle x, Ay \rangle = v.$$

We will prove a more general version of this theorem when we discus regret minimization.

Because zero-sum Nash equilibria are min-max solutions, they are the best that a player can do, given a worst-case opponent. This guarantee is the rationale for saying that a given game has been *solved* if a Nash equilibrium has been computed. Some games such as rock-paper-scissors are trivially solvable as we know that uniform distribution is the only equilibrium. However, this notion has also been applied to heads-up limit Texas hold'em, one of the smallest poker variants played by humans. In 2015, that games was *essentially solved* that game. Their notion of essentially solved is based on having computed a strategy that is statistically indistinguishable from a Nash equilibrium in a lifetime of human-speed play. The statistical notion was necessary because their solution was computed using iterative methods that only converge to an equilibrium in the limit (but in practice get quite close very rapidly). The same argument is also used in constructing AIs for even larger two-player zero-sum poker games where we can only try to approximate the equilibrium.

Note that this guarantee does not hold in general-sum games, where we have no payoff guarantees if our opponent does not play their part of the same Nash equilibrium that we play. Interestingly, the AI and optimization methods developed for two-player zero-sum poker turned out to still outperform top-tier human players in 6-player no-limit Texas hold'em poker. An AI based on these methods ended up beating professional human players, in spite of the methods having no guarantees on performance, nor even of converging to a general-sum Nash equilibrium.

Here is another interesting property of zero-sum Nash equilibrium: it is interchangeable. Meaning that if you take an equilibrium (x, y) and another equilibrium (x', y') then (x, y') and (x', y) are also equilibria. This is easy to see from the minimax formulation.

2 Historical Notes

[5] were the first to show that solving general games is a PPAD-complete problem. Their initial result was for four-player games. Chen et al. [3] showed that the result holds even for two-player general-sum games. NP-completeness of finding Nash equilibria with various properties was shown by Gilboa and Zemel [6] and Conitzer and Sandholm [4].

The result where Heads-up limit Texas hold'em was *essentially solved* was by Bowling et al. [1]. That paper also introduced the notion of "essentially solved." The strong performance against top-tier humans in 6-player poker was shown by Brown and Sandholm [2].

References

[1] Michael Bowling, Neil Burch, Michael Johanson, and Oskari Tammelin. Heads-up limit hold'em poker is solved. *Science*, 347(6218):145–149, 2015.

[2] Noam Brown and Tuomas Sandholm. Superhuman AI for multiplayer poker. *Science*, 365 (6456):885–890, 2019.

[3] Xi Chen, Xiaotie Deng, and Shang-Hua Teng. Settling the complexity of computing two-player nash equilibria. *Journal of the ACM (JACM)*, 56(3):1–57, 2009.

[4] Vincent Conitzer and Tuomas Sandholm. New complexity results about Nash equilibria. *Games and Economic Behavior*, 63(2):621–641, 2008.

[5] Constantinos Daskalakis, Paul W Goldberg, and Christos H Papadimitriou. The complexity of computing a Nash equilibrium. *SIAM Journal on Computing*, 39(1):195–259, 2009.

[6] Itzhak Gilboa and Eitan Zemel. Nash and correlated equilibria: Some complexity considerations. *Games and Economic Behavior*, 1(1):80–93, 1989.