

Normal-Form Games

CSCI 1440/2440

2025-09-03

We describe (complete-information) normal-form games (formally; no examples), and corresponding equilibrium concepts.

1 Introduction

Decision theory is a theory about how individuals—often called **agents**—make decisions. It relies on an underlying model of an agent’s knowledge and beliefs about the world, including its available actions, and its preferences over the consequences of its actions. Von Neumann and Morgenstern’s classic expected utility theory¹ argues that certain (reasonable: e.g., non-contradictory) preferences can be represented as **utility** functions, so that decision makers can be viewed as “expected² utility maximizers.” This point of view is taken in this class, where utilities are associated with agents’ actions, and more utility is preferred to less.

¹ J. Von Neumann and O. Morgenstern. *Theory of Games and Economic Behavior*. Princeton University Press, 1947

² This expectation can be with respect to randomness in the agents’ strategies, or uncertain beliefs about the world.

2 A Model of Interaction

Game theory might better be named “multiagent decision theory.” A **game** is a setting in which a set of players³ interact, after which they receive some utility. In a **normal-form** game (NFG), each player $i \in [n]$ selects an action a_i from a discrete set of actions A_i .⁴ Together, these actions form an action profile (i.e., a vector) $\mathbf{a} = (a_1, \dots, a_n)$, which is an element of the space of all possible action profiles $A = \prod_{i \in [n]} A_i$. Finally, each player $i \in [n]$ is endowed with a utility function $u_i : A \rightarrow \mathbb{R}$, which depends on the action profile taken.

³ We tend to use the terms players and agents interchangeably.

⁴ In this lecture, we assume finite games, meaning finitely many players and actions.

One key assumption is that all of the above is **common knowledge** among the players. This means that all players know what they know; they know what they don’t know; they don’t know anything that isn’t true; and they know all implications of what they know. One such implication might be “I know that you know that I know that you know ... all the details of the game we are playing.”

A **strategy** s_i in a NFG is either **pure**, meaning simply an action, or **mixed**, meaning a probability distribution over actions. Define $S_i = \Delta(A_i)$ as player i ’s set of—in general, mixed—⁵ strategies. Assuming players employ (mixed⁶) strategies, they must reason about **expected** utility, where the expectation is taken over their mixture.

⁵ Note that pure strategies are degenerate mixed strategies.

⁶ Hereafter, “mixed” is implicit in the term strategy.

Rational players are those who maximize their expected utility.

For today (at least), we assume rationality on the part of the players.

Let $\mathbf{s} \in S = \prod_{i \in [n]} S_i$ denote a strategy profile for all players; and let $\mathbf{s}_{-i} = (s_1, \dots, s_{i-1}, s_{i+1}, \dots, s_n) \in S_{-i} = \prod_{j \neq i} S_j$ denote the strategies of all players except player i . We use corresponding notation for action and other profiles throughout.

A strategy $s_i \in S_i$ for player $i \in [n]$ is **dominant** if it is (weakly) optimal, regardless of the other players' actions: i.e.,

$$u_i(s_i, \mathbf{a}_{-i}) \geq u_i(s'_i, \mathbf{a}_{-i}), \quad \forall s'_i \in S_i, \forall \mathbf{a}_{-i} \in A_{-i} \quad (1)$$

A strategy profile $\mathbf{s} \in S$ is a **dominant strategy equilibrium** (DSE) if all players play dominant strategies.⁷ DSE is a worst-case concept; it does not rely on the assumption that the other players are rational.

Remark 2.1. If a strategy is dominant regardless of other players' actions (i.e., pure strategies), it is also dominant regardless of other players' mixed strategies.

Proof. Fix a player i and assume $s_i \in S_i$ is a dominant strategy for i . Given an arbitrary (mixed) strategy profile \mathbf{s}_{-i} , let $\pi(\mathbf{a}_{-i}) = \prod_{j \neq i} s_j(a_j)$ denote the joint probability of action profile \mathbf{a}_{-i} . Then:

$$\begin{aligned} u_i(s_i, \mathbf{s}_{-i}) &= \sum_{\mathbf{a}_{-i}} u_i(s_i, \mathbf{a}_{-i}) \pi(\mathbf{a}_{-i}) \\ &\geq \sum_{\mathbf{a}_{-i}} u_i(s'_i, \mathbf{a}_{-i}) \pi(\mathbf{a}_{-i}), \quad \forall s'_i \in S_i \\ &= u_i(s'_i, \mathbf{s}_{-i}), \quad \forall s'_i \in S_i. \end{aligned}$$

The inequality follows from the definition of dominant strategy. \square

A strategy $s_i^* \in S_i$ (mixed or pure) is a **best response** for player i to a fixed other-player strategy profile $\mathbf{s}_{-i} \in S_{-i}$ if it is expected-utility maximizing for i :

$$s_i^* \in \arg \max_{s_i \in S_i} u_i(s_i, \mathbf{s}_{-i}). \quad (2)$$

A strategy profile $\mathbf{s} = (s_i, \mathbf{s}_{-i}) \in S$ is a **Nash equilibrium** (NE) if all players strategies are best responses to one another. In other words, no player can increase her expected utility by unilaterally changing her strategy: i.e.,

$$u_i(s_i, \mathbf{s}_{-i}) \geq u_i(s'_i, \mathbf{s}_{-i}), \quad \forall i \in [n], \forall s'_i \in S_i. \quad (3)$$

It turns out that not all normal-form games have a pure strategy Nash equilibrium. Consider the childhood game of *Rochambeau* (i.e., rock-paper-scissors). Rock is preferred to (i.e., a best-response to) scissors, which in turn is preferred to paper, which in turn is preferred to rock. This game does, however, have a mixed strategy Nash equilibrium, which is to play all three actions uniformly at random.

⁷ This condition is sufficient for DSE, but not necessary. It is also possible to discover DSE via the iterative deletion of dominated strategies.

Perhaps the most important feature of the Nash equilibrium is its guaranteed existence in mixed strategies in finite games,⁸ a fact which was established by Nobel laureate John Nash in 1951.

Theorem 2.2 (Nash, 1951). *Every finite game has a mixed strategy Nash equilibrium.*

Unfortunately, despite their guaranteed existence, Nash equilibria are hard to compute. It turns out that their computation is complete for a complexity class called **PPAD**.

The PPAD complexity class differs from other well-known complexity classes such as P and NP in that it concerns **function problems**, rather than decision problems: i.e., given a polynomial-time computable (total) function⁹ $P(x, y)$ s.t. for all x there exists a y s.t. $P(x, y)$ holds, and given a particular x , find y s.t. $P(x, y)$ holds. Problems in the PPAD class specialize problems of this sort with one further commonality: at their core lies a fixed point computation.

Analogous to NP-completeness, a problem in the PPAD class is called PPAD-complete if all problems in PPAD can be reduced to that problem. In other words, a PPAD-complete problem is one that is at least as hard to solve as every other problem in the PPAD class.

Theorem 2.3 (2000s). *(Chen et al. 2007; Daskalakis et al. 2009) Computing a Nash equilibrium in two-player, finite-action games is PPAD-complete.*

A von Neumann & Morgenstern's Expected Utility Theory

It is natural to express preferences as comparisons: e.g., “I prefer apples to bananas.” We can also compare lotteries (i.e., randomized outcomes): “I prefer a banana with probability 90% to an apple with probability 50%.” Let Ω denote an outcome space, and let σ and τ denote lotteries. We write $\sigma \succ \tau$ to indicate that σ is preferred to τ .

Theorem A.1 (vNM, 1944). *Assuming an agent's preferences over lotteries satisfy certain axioms, there exists a unique¹⁰ utility function $u : \Omega \rightarrow \mathbb{R}$ s.t. $\sigma \succ \tau$ iff $\mathbb{E}[u(\sigma)] > \mathbb{E}[u(\tau)]$. The axioms are completeness, transitivity, continuity, and independence of irrelevant attributes (IIA).*

Completeness means *all* lotteries are comparable (i.e., $\sigma \succeq \tau$ or $\tau \succeq \sigma$), although, as the notation suggests, the agent may be indifferent between lotteries. Transitivity means if $\sigma \succeq \tau$ and $\tau \succeq v$, then $\sigma \succeq v$.

Continuity is defined as follows: if $\sigma \succeq \tau \succeq v$, then there exists $p \in [0, 1]$ s.t. the agent is indifferent between τ and σ with probability p , and τ and v with probability $1 - p$.

Finally, IIA can be understood as follows: if σ is the preferred outcome among a set of outcomes T , and if $\sigma \in S \subseteq T$, then σ should

⁸ A finite game has a finite number of players and actions.

⁹ Technically, a predicate.

¹⁰ up to affine transformations

be the preferred outcome in S as well. In other words, removing outcomes from a set that are not preferred should not change the preferred outcome. On the surface, IIA sounds innocuous enough. But Nobel Laureate Daniel McFadden's example with three transportation alternatives, a car, a red bus, and a blue bus, shows that it does not properly account for substitutes, like red and blue buses.

References

- [1] J.F. Nash. Non-cooperative games. *Annals of Mathematics*, 54(2):286–295, 1951.
- [2] J. Von Neumann and O. Morgenstern. *Theory of Games and Economic Behavior*. Princeton University Press, 1947.