Game Theory – theory used to analyze games of simultaneous and continuous moves.

- **Agent Design** – determining the best strategy to take against a rational player as well as the expected return for each player.
- Components
  - *Players*
  - *Actions* each player can choose.
  - **Payoff Matrix** – gives the utility for each player in each combination of actions the players can take.
- **strategy** – a policy for taking actions in a given situation.
  - **pure strategy** – there is a predetermined action for each situation.
  - **mixed strategy** – a randomized policy choosing actions from a distribution; action a_i chosen w/ probability p_i: $[p_1 : a_1; \ldots; p_n : a_n]$
- **strategy profile** – an assignment of a strategy to each player.
  - **solution** – a strategy profile where each player adopts a rational strategy.
  - **strongly dominates** – a strategy *s* strongly dominates strategy *s'* if the outcome for *s* is better than the outcome for *s'* (with respect to player *p*).
  - **weakly dominates** – a strategy *s* weakly dominates strategy *s'* if *s* is better than *s'* on at least one strategy profile and is no worse on any other profile.
  - **dominant strategy** – a strategy that dominates all others.
- **outcome** – a numeric value for each player based on the results of the game.
  - **Pareto optimal** – an outcome preferred by all players over any other.
  - **Pareto dominated** – one outcome is pareto dominated by a second if all players would prefer the second outcome.
- **Nash equilibrium** – a property of a strategy profile such that no player can benefit from changing strategies.
  - **dominant strategy equilibrium** – each player has a dominant strategy.
  - *Every game has a Nash equilibrium* (although not necessarily dominant).
  - *When there a multiple acceptable solutions (equilibria), if each player chooses a different solution, the resulting strategy profile may not be a solution and all agents will suffer.*
    - could use Pareto-optimal Nash Equilibrium if one exists.
    - **coordination games** – games in which players need to communicate.

- **Maximin Equilibrium** – a Nash Equilibrium for mixed strategies.
  - **zero-sum game** – game in which payoffs in each cell of the payoff matrix sum to 0.
  - Algorithm
    - Assume that the 1st player goes 1st. The strategy for the 2nd player is now a pure strategy since the expected utility becomes a convex combination and thus no mixed strategy can do better than a pure.
      - This can be thought of as a minimax tree with a branch for each of the 1st player's possible strategies, each of which has 2 branches for the 2nd player.
      - The result is a hyperplane in the space defined by the probability of each action and its expected utility. For *n*-actions, *n* such hyperplanes are created.
      - dominated strategies for the 2nd player are removed
      - the optimal choice is at the intersection of the hyperplanes (a maximum), which can be found by linear programming.
    - This process is repeated for each player.
  - *Every two-player zero-sum game has a maximin equilibrium for mixed strategies.*
  - *Every equilibrium in a zero-sum game is a maximin for both players.*
  - Non-zero-sum games:
    1. Enumerate all possible subsets of actions that might form mixed strategies.
    2. For each strategy profile enumerated, check to see if it is an equilibrium.
- **prisoner's dilemma** – a game in which two thieves are being interrogated seperatly. If both *refuse* to confess, they will get 1 year each. If *testify*, both will get 5 years. But if 1 *refuses* and the other *testifies*, the former gets 10 years and the later gets 0.
  - optimal strategy for both is to *testify*.
- **repeated game** – player's face the same choice repeatedly but each time with the knowledge of the history of all players' previous choices.
  - If the number of repetitions (meetings) is known, the outcome can be inductively determined by the optimal strategy for the last meeting.
  - More cooperative behavior is possible if the chance that the player's will meet again is probabilistic.
    - **perpetual punishment** – equilibrium strategy to be nice to other player unless the other player has ever betrayed you.
    - **tit-for-tat** – start with *refuse* action and mimic other player's previous move from that point on.
    - **ignorance is bliss** – having other player think you are ignorant.
- **games of partial information** – repeated games with partially-observablity.
- **Bayes-Nash equilibrium** – an equilibrium with respect to a player's prior probability distribution over the other players' strategies.
  - addresses the fact that the other player might not play an equilibrium strategy thereby allowing an improvement.

Mechanism Design– how to define rules of the environment so that the collective good of all agents is maximized when each agent adopts the game-theoretic solution to maximize its own utility. Alternatively, a way to design multiagent systems that solve problems in a distributed fashion without each agent needing to know what problem is being solved.

- **mechanism** – consists of (1) a language for describing the strategies an agent may use and (2) an outcome rule *G* that determines the payoffs to the agents given a strategy profile.
- **tragedy of commons** – situation in which individuals acting for individual good create global bad (Farmers overgraze commons – shared field).
    - Must ensure that all *externalities* (effects on global utility not recognized by agents) are made explicit.
- **strategy-proof mechanism** – a mechanism where players have a dominant strategy that ultimately reveals their true incentives.
- Auctions – (1) there is a single good (2) each bidder has a utility value v_i for the good (3) the value is only known to the bidder. The bidders make bids b_i and the highest bid wins the goods.
    - **English Auction** – auctioneer increments the price of the goods until only 1 bidder remains.
        - simple dominant strategy, bid until your personal value is exceeded, results in bidder with highest value getting the goods.
        - Requires high bandwidth secure communication.
    - **Sealed-Bid Auction** – each bidder makes a single bid communicated to the auctioneer and the highest bid wins.
        - Player with highest value may not get the goods.
        - Players must spend effort considering other player's strategies.
    - **Sealed-Bid Second-Price (Vickrey) Auction** – winner pays the price of the second highest bid.
    - dominant strategy is to bid player's actual value and player with the highest value wins the goods.