Introduction

- What if we didn't always repeat back to the same stage game?
- A stochastic game is a generalization of repeated games
  - agents repeatedly play games from a set of normal-form games
  - the game played at any iteration depends on the previous game played and on the actions taken by all agents in that game
- A stochastic game is a generalized Markov decision process
  - there are multiple players
  - one reward function for each agent
  - the state transition function and reward functions depend on the action choices of both players

Formal Definition

Definition

A stochastic game is a tuple (Q, N, A, P, R), where

- Q is a finite set of states,
- N is a finite set of n players,
- $A = A_1 \times \cdots \times A_n$, where A_i is a finite set of actions available to player i,
- $P : Q \times A \times Q \mapsto [0, 1]$ is the transition probability function; $P(q, a, \hat{q})$ is the probability of transitioning from state q to state $\hat{q}$ after joint action a, and
- $R = r_1, \ldots, r_n$, where $r_i : Q \times A \mapsto \mathbb{R}$ is a real-valued payoff function for player i.

Remarks

- This assumes strategy space is the same in all games
  - otherwise just more notation
- Again we can have average or discounted payoffs.
- Interesting special cases:
  - zero-sum stochastic game
  - single-controller stochastic game
    - transitions (but not payoffs) depend on only one agent

Strategies

- What is a pure strategy?

Strategies

- What is a pure strategy?
  - pick an action conditional on every possible history
  - of course, mixtures over these pure strategies are possible too!
- Some interesting restricted classes of strategies:
  - behavioral strategy: $s_i(h_t, a_{i_j})$ returns the probability of playing action a_{i_j} for history h_t.
    - the substantive assumption here is that mixing takes place at each history independently, not once at the beginning of the game
  - Markov strategy: s_i is a behavioral strategy in which $s_i(h_t, a_{i_j}) = s_i(h'_t, a_{i_j})$ if $q_t = q'_t$, where q_t and q'_t are the final states of h_t and h'_t, respectively.
    - for a given time t, the distribution over actions only depends on the current state
  - stationary strategy: s_i is a Markov strategy in which $s_i(h_{t_1}, a_{i_j}) = s_i(h'_{t_2}, a_{i_j})$ if $q_{t_1} = q'_{t_2}$, where q_{t_1} and q'_{t_2} are the final states of h_{t_1} and h'_{t_2}, respectively.
    - no dependence even on t

Equilibrium (discounted rewards)

- Markov perfect equilibrium:
  - a strategy profile consisting of only Markov strategies that is a Nash equilibrium regardless of the starting state
  - analogous to subgame-perfect equilibrium

Theorem

Every n-player, general sum, discounted reward stochastic game has a Markov perfect equilibrium.

Equilibrium (average rewards)

- Irreducible stochastic game:
  - every strategy profile gives rise to an irreducible Markov chain over the set of games
    - irreducible Markov chain: possible to get from every state to every other state
  - during the (infinite) execution of the stochastic game, each stage game is guaranteed to be played infinitely often—for any strategy profile
  - without this condition, limit of the mean payoffs may not be defined

Theorem

Every 2-player, general sum, average reward, irreducible stochastic game has a Nash equilibrium.

Folk Theorems for Stochastic Games

Theorem

For every 2-player, general sum, irreducible stochastic game, and every feasible outcome with a payoff vector r that provides to each player at least his minmax value, there exists a Nash equilibrium with a payoff vector r. This is true for games with average rewards, as well as games with large enough discount factors (i.e. with players that are sufficiently patient).