Models of Bounded Rationality
Lecture 2: Game Theory

Thomas Icard

August 16, 2016

Motivating Example (Ben-Sasson et al. 2007)

Imagine the following two-player game:

- Player 1 presents player 2 with an m-bit number N.
- Player 2 responds with a list $p_1, \ldots, p_k$.
- If $p_1, \ldots, p_k$ are the prime factors of N, then Player 1 pays Player 2 \$10; otherwise Player 2 pays Player 1 \$10.

1. Quick review of game theory
2. Games where the Players are Automata
3. Interlude on Probabilistic Computation
4. Games Played by Probabilistic Turing Machines

Basic Elements of a Game

- Set of players: $N = \{1, 2, \ldots, n\}$
- Possible actions A_i for each player $i \in N$
  We write $A = (A_1, \ldots, A_n)$
- Payoff function $u_i : A \to \mathbb{R}$ for each player

Games in Normal Form

	C	D
C	$3,3$	$0,4$
D	$4,0$	$1,1$

Games in Normal Form

	R	P	S
R	$0, 0$	$-1, 1$	$1, -1$
P	$1, -1$	$0, 0$	$-1, 1$
S	$-1, 1$	$1, -1$	$0, 0$

"Rational strategy choice by a given player in any game always amounts to choosing a strategy maximizing his expected payoff in terms of a subjective probability distribution over the strategy combinations available to the other players. But this immediately poses the question of *how this probability distribution is to be chosen* by a rational player—more specifically, how this distribution is to be chosen by a player who expects the *other players to act rationally*, and also expects these other players to entertain similar expectations about him and about each other."

—Harsanyi 1982

Strategies

Definition

A *strategy* σ_i for player i is an element of $\Delta(A_i)$, the set of distributions over the action space A_i.

Let us write

$$U_i^\sigma = \sum_{a \in A} \left(\sigma_1(a_1) \times \cdots \times \sigma_n(a_n)\right) u_i(a)$$

for a vector $\sigma = (\sigma_1, \ldots, \sigma_n)$ of strategies.

Nash Equilibrium

$$U_i^\sigma = \sum_{a \in A} \left(\sigma_1(a_1) \times \cdots \times \sigma_n(a_n)\right) u_i(a)$$

Definition

A vector of strategies $\sigma = (\sigma_1, \dots, \sigma_n)$ is a *Nash equilibrium* if for all $i \in N$:

$$U_i^\sigma \geq U_i^{\sigma'}$$

for any σ' that differs only on σ_i.

	C	D
C	$3, 3$	$0, 4$
D	$4, 0$	$1, 1$

$$(D, D)$$

	R	P	S
R	$0, 0$	$-1, 1$	$1, -1$
P	$1, -1$	$0, 0$	$-1, 1$
S	$-1, 1$	$1, -1$	$0, 0$

$$(\langle \frac{1}{3}R, \frac{1}{3}P, \frac{1}{3}S \rangle, \langle \frac{1}{3}R, \frac{1}{3}P, \frac{1}{3}S \rangle)$$

Theorem (Nash 1950)

Every finite game has a Nash equilibrium.

Repeated Games

Given a game $G = (N, A, u)$, we can imagine playing this game infinitely many times. An *outcome* is a sequence $a = (a^1, \dots, a^n)$. Let $\mathcal{H}$ be the set of all finite sequences of outcomes. Elements $h \in \mathcal{H}$ are called *histories*.

Definition

A *strategy* is now a function $f_i : \mathcal{H} \to \Delta(A_i)$.

Imagine an infinite sequence of outcomes $\vec{a} = (a^1, \dots, a^k, \dots)$. We can assess the utility for player i using a discount $\gamma < 1$:

$$\bar{U}_i(\vec{a}) = \sum_{k=1}^{\infty} \gamma^{k-1} u_i(a^k)$$

The utility for i given strategy profile f is:

$$\bar{U}_i^f = \mathbb{E}_{\vec{a}} \bar{U}_i(\vec{a})$$

Nash equilibrium is again defined analogously.

Theorem (Folk)

Suppose there is a set of (deterministic) strategies σ for a single-shot game such that U_i^σ is better than the minmax payoff for each $i \in N$.

Then there is a set of strategies $f = (f_1, \dots, f_n)$ for the repeated game, such that $\bar{U}_i^f = \sum_{k=1}^{\infty} \gamma^{k-1} U_i^\sigma$.

In particular this means that cooperation is an equilibrium of the infinite prisoners dilemma!

Automata in Game Theory

The study of iterated games by finite automata has been explored by a number of authors: Aumann, Radner, Rubinstein, Neyman, Kalai, and many others. Here we largely follow Rubinstein (1998) and Kalai & Stanford (1988).

Moore Automata

$$(Q_i, q_i^0, o_i, \tau_i)$$

- Q_i a finite set of states
- q_i^0 a distinguished initial state
- $o_i : Q_i \to A_i$ an output function
- $\tau_i : Q_i \times A \to Q_i$ a transition function

Example: "Always Cooperate"

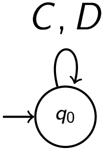

$$o_i(q_0) = C$$

Example: “Always Defect”

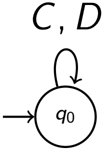

$$o_i(q_0) = D$$

Example: "Trigger"

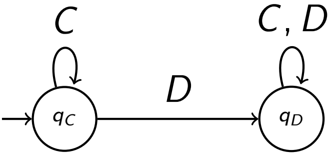

$$o_i(q_C) = C \qquad o_i(q_D) = D$$

Example: “Tit-for-Tat”

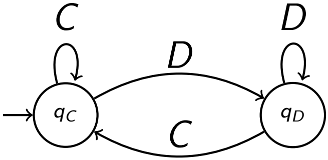

$$o_i(q_C) = C \qquad o_i(q_D) = D$$

Example: "3-Punisher"

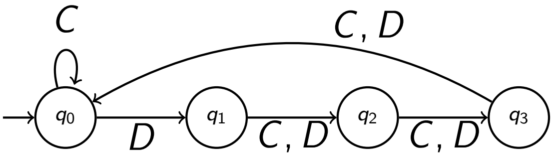

$$o_i(q_0) = C \quad o_i(q_1) = o_i(q_2) = o_i(q_3) = D$$

Fact

Every (deterministic) strategy f_i can be represented as a (possibly infinite) automaton.

Proof.

Let $\mathcal{H}$ be the set of states, ϵ the initial state, define $\tau_i(h, a) = h \cdot a$, and let $o_i(h) = f_i(h)$. □

Many strategies can be played by finite automata.

Given a strategy f_i and a history h, let us write $f_i \upharpoonright h$ for the function

$$f_i \upharpoonright h\,(h') = f_i(h \cdot h').$$

Definition

The *complexity* of a strategy f_i is given by the cardinality of the set

$$\{f_i \upharpoonright h : h \in \mathcal{H}\}.$$

Definition

The *complexity* of a strategy f_i, written $C(f_i)$, is given by the cardinality of the set

$$\{f_i \upharpoonright h : h \in \mathcal{H}\}$$

i.e., the number of equivalence classes on $\mathcal{H}$, where $h_1 \equiv h_2$ iff $f_i \upharpoonright h_1 = f_i \upharpoonright h_2$.

Theorem (Kalai & Stanford)

The smallest f_i automaton has exactly $C(f_i)$ states.

Definition

The cost-adjusted utility of strategy f_i is given by

$$\mathbf{U}_i^f \;=\; \bar{U}_i^f - \lambda C(f_i)$$

for some $\lambda \geq 1$.

Example

When $\gamma = \frac{3}{4}, \lambda = 3$, and the other player is playing Tit-for-Tat (Cooperate, Trigger, 3-Punisher), the utility of Tit-for-Tat is $\left(\sum_{k=1}^{\infty}\left(\frac{3}{4}\right)^{k-1} 3\right) - 6 = 6$.

Question: Does the folk theorem still hold?

Answer: No. Just consider Tit-for-Tat . . .

Notice also that the pair of machines playing Tit-for-Tat are no longer in equilibrium!

Funny Example (from Rubinstein 1998)

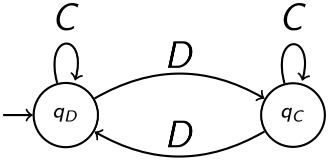

$$o_i(q_D) = D \qquad o_i(q_C) = C$$

Further Questions

- Characterize when machine players are in equilibrium (Rubinstein).
- Extend to other types of equilibria, e.g., subgame perfect (Kalai & Stanford).
- Extend to other types of games, e.g., congestion games (Bar-Sasson et al.).
- Study reinforcement learning in this context (Bar-Sasson et al.).
- Investigate computational conditions for cooperation (Anderlini).

When considering computational devices, it is useful to allow probabilistic computers. In fact, these can take the place of mixed strategies:

Theorem (Kuhn 1953)

Under the assumption of perfect recall, probabilistic strategies and mixed strategies are equivalent.

But instead of moving to probabilistic automata, we consider more powerful probabilistic machines.

Probabilistic Turing Machines

- Add to TMs ability to read a random bit tape.
- The bits on the random bit tape can be thought of flips of a fair coin.
- In this way, a PTM can be thought of as defining a distribution over outputs.

Definition (Turing)

A real number $r \in [0, 1]$ is *computable* if there is a (deterministic) TM M, which on input n outputs a number r' such that

$$|r - r'| < \frac{1}{2^n}$$

Definition

A computable distribution on $\{0, 1\}^*$ is one for which the probability of each string is a computable real number, uniformly in the string.

Theorem (Universality)

The distributions definable by a PTM are exactly the computable distributions.

Example

Build a machine to output 1, 11, and 111, each with probability $\frac{1}{3}$.

Note that $\frac{1}{3} = \sum_{k=1}^{\infty} 2^{-(2k+1)}$, so its binary representation is $0.01010101\ldots$.

Bayesian Machine Games (Halpern, Pass)

- Agents are represented by PTMs
- Nature chooses *types* $t_i \in \{0,1\}^*$ for each player, with distribution Pr over $\left(\{0,1\}^*\right)^n$
- A machine M_i takes in $t = (t_1, \ldots, t_n)$ and a random bit r, outputs $a_i \in A_i = \{0,1\}^*$
- Machines $M_i \in \mathcal{M}$ with inputs $t; r$ are associated with costs:

$$C_i : \mathcal{M} \times \{0,1\}^*; \{0,1\}^\infty \to \mathbb{N}$$

- $u_i : T \times \left(\{0,1\}^*\right)^n \times \mathbb{N}^m \to \mathbb{R}$

Bayesian Machine Games (Halpern, Pass)

$$\mathcal{G} = (N, \mathcal{M}, T, Pr, C_1, \dots, C_n, u_1, \dots, u_n)$$

$$u_i^{\mathcal{G},M}(t, r) \;=\; u_i\Big(t, M_1(t_1, r_1), \dots, M_n(t_n, r_n), C_1(M_1, t_1; r_1), \dots, C_n(M_n, t_n; r_n)\Big)$$

$$U_i^{\mathcal{G}}[M] = \mathbb{E}_{Pr^*}[U_i^{\mathcal{G},M}]$$

Definition

A list of choices $M = (M_1, \ldots, M_n)$ of PTMs is a *Nash equilibrium* of $\mathcal{G}$ if for all $i \in N$, M_i is a best response to M_{-i}. I.e., if

$$U_i^{\mathcal{G}}[(M_i, M_{-i})] \geq U_i^{\mathcal{G}}[(M_i', M_{-i})]$$

for all $M_i' \in \mathcal{M}$.

Nash Equilibria do not Always Exist

	R	P	S
R	$0, 0$	$-1, 1$	$1, -1$
P	$1, -1$	$0, 0$	$-1, 1$
S	$-1, 1$	$1, -1$	$0, 0$

$$C_i(M_i) = \begin{cases} 1 & \text{if } M_i \text{ is deterministic} \\ 2 & \text{if } M_i \text{ involves randomization} \end{cases}$$

	R	P	S
R	$0, 0$	$-1, 1$	$1, -1$
P	$1, -1$	$0, 0$	$-1, 1$
S	$-1, 1$	$1, -1$	$0, 0$

$$C_i(M_i) = \begin{cases} 1 & \text{if } M_i \text{ is deterministic} \\ 2 & \text{if } M_i \text{ involves randomization} \end{cases}$$

- If M_1 used randomization, player 1 would do better by playing best response to highest probability action for M_2.
- Likewise, M_2 cannot use randomization. Yet there are no deterministic equilibria!

	R	P	S
R	$0, 0$	$-1, 1$	$1, -1$
P	$1, -1$	$0, 0$	$-1, 1$
S	$-1, 1$	$1, -1$	$0, 0$

$$C_i(M_i) = \begin{cases} 0 & \text{if } M_i \text{ uses fewer than 10,000 steps} \\ 1 & \text{otherwise} \end{cases}$$

- Note that M_i cannot with probability 1 simulate the uniform distribution on $\{R, P, S\}$!
- Thus, there can be cases with no equilibria even when all constant time strategies are costless!

Another Difference

If standard game theory, any deterministic strategy in the support of a mixed strategy equilibrium is as good as the mixed strategy.

E.g., in Rock-Paper-Scissors R is as good as $\langle \frac{1}{3}R, \frac{1}{3}P, \frac{1}{3}S \rangle$, given the other plays $\langle \frac{1}{3}R, \frac{1}{3}P, \frac{1}{3}S \rangle$.

This is not so in Bayesian machine games!

The type t_1 of player 1 is distributed uniformly among all odd numbers between 2^{100} and 2^{101}. The goal is to determine whether t_1 is prime.

$$U_1^M = \begin{cases} 2 & \text{if correct, using fewer than } K \text{ steps} \\ 0 & \text{if correct, using at least } K \text{ steps} \\ -1000 & \text{if wrong} \\ 1 & \text{if abstain} \end{cases}$$

There may well be probabilistic primality tests that are very accurate and run in time less than K, while no deterministic algorithm does.

Existence of Nash Equilibria?

- Are there natural conditions guaranteeing the existence of Nash equilibria in Bayesian machine games?
- Halpern & Pass (2008) show that, in a certain sense, if randomization is free, then equilibria will always exist.
- Assumptions: game is *computable*, meaning that $Pr[t]$ and $u_i(t, a, c)$ are all computable real numbers.

Lemma

In a computable game, if there is no charge for computation, then an equilibrium exists.

This is just a standard Bayesian game, so it has a solution. But how do we know this solution can be implemented by PTMs?

Lemma (Halpern & Pass)

In a computable game, if there is no charge for computation, then an equilibrium exists.

Proof Sketch.

First show that there is a mixed strategy solution with *computable* probabilities. Then by universality of PTMs, we know we can find machines that have this behavior. The first step uses the Tarski-Seidenberg Theorem. □

Theorem (Tarski-Seidenberg)

If $R \subseteq R'$ are real closed fields, and P is a finite set of polynomials with coefficients in R, then P has a solution in R iff it has one in R'.

Note that σ is a Nash equilibrium iff for all t_i, a_i, a_i':

$$\sigma(t_i, a_i) \geq 0 \qquad \sum_{a_i \in A_i} \sigma(t_i, a_i) = 1$$

$$\sum_{t_{-i}} \sum_{a} Pr(t)\, u_i(t, a) \prod_{j \in N} \sigma_j(t_j, a_j)$$

$$\geq \sum_{t_{-i}} \sum_{a_{-i}} Pr(t)\, u_i(t, a_i', a_{-i}) \prod_{j \in N \setminus i} \sigma_j(t_j, a_j)$$

Now replace each $\sigma_j(t_j, a_j)$ with a variable x_{j,t_j,a_j}.

Corollary (Halpern & Pass)

Under the same assumptions, if $\mathcal{M}$ is the computable convex closure of some finite set $\mathcal{M}_0$ of TMs, then the game has an equilibrium.

Proof Sketch.

Given $\mathcal{M}_0$, we can consider this a Bayesian game with payoffs cost-adjusted. This has an equilibrium, and using the same trick as before, we can find PTMs that simulate mixtures. □

Many Questions about Games with PTMs

- Sequential games
- Games with communication
- Cryptographic applications
- See Halpern & Pass for more

Conclusions

- When taking computational costs seriously, important theorems (Nash, Folk) may fail.
- Once we take computational considerations into account, is (standard) game theory still the right tool to analyze multiagent interaction?
- How might this work be brought closer to empirical work on social reasoning? (Cf. Halpern & Pass for some ideas.)
- How does all of this relate to bounded optimality?