Bounding Rationality by Discounting Time

Lance Fortnow

Rahul Santhanam

Plan of the Talk

- Introduction
- The Model
- Results
- Future Directions

Plan of the Talk

- *Introduction*
- The Model
- Results
- Future Directions

Perfect Rationality

- Perfect rationality in a strategic situation
  - Each player is rational (knows its payoff, and wishes to maximize it)
  - Each player knows that the other player is rational
  - Each player can derive all consequences of common rationality

Bounded Rationality

- Herbert Simon – "Boundedly rational agents experience limits in formulating and solving complex problems and in processing (receiving, storing, retrieving, transmitting) information"
- In particular, boundedly rational agents are subject to *computational constraints*

Games

- Simultaneous-move (eg., Prisoner's Dilemma) or Sequential-move (eg., chess)
- Simultaneous-move
  - Action spaces: A_1, A_2
  - Strategy spaces: $P(A_1)$, $P(A_2)$
  - Payoff functions: $A_1 \times A_2 \rightarrow \mathbb{R}$
- Sequential-move (one-shot)
  - Strategy spaces: $P(A_1)$, $P(A_2)$^A_1
  - Payoff functions: $A_1 \times A_2 \rightarrow \mathbb{R}$

Nash Equilibrium

- A pair of strategies (S_1,S_2) is an NE if
  - For all T_2, $u_2(S_1,S_2) >= u_2(S_1,T_2)$
  - For all T_1, $u_1(S_1,S_2) >= u_1(T_1, S_2)$
- Theorem [Nash]: Every finite game has an NE

Almost-Nash Equilibrium

- A pair of strategies (S_1,S_2) is a γ-NE if
  - For all T_2, $u_2(S_1,S_2) >= u_2(S_1,T_2) - \gamma$
  - For all T_1, $u_1(S_1,S_2) >= u_1(T_1, S_2) - \gamma$

The Largest Number Game

Alice	Bob	Payoff (to Alice)
M (Integer)	N (Integer)	100 if M > N, 50 if M=N, 0 otherwise

Largest Number game does not have an NE, or even an almost-NE if $\gamma < 50$

The Factoring Game (sequential-move)

Alice	Bob	Payoff (to Bob)
M (Integer)	X,Y (Integers)	100 if M is prime or if $1 < X,Y < M$ and $M = X * Y$, 1 otherwise

Factoring Game has infinitely many Nash equilibria, in each of which Bob gets payoff 100 and Alice gets payoff 1 (Bob's strategy is simply to factor Alice's number)

Plan of the Talk

- Introduction
- *The Model*
- Results
- Future Directions

Time is Money

- The *time* it takes to implement a strategy is relevant
- Payoffs should decrease with time
- Exponential discounting: Let $\varepsilon < 1$ be a discount factor. Then payoff decreases by a factor $(1-\varepsilon)^t$ after t steps

Asymmetric Discounting

- In general, different players have different discount factors
  - The players might have different roles in the game
  - Even if the game is symmetric, the players themselves might not be equally patient
- ε: Alice's discount factor
- δ: Bob's

Discounting and Computational Power

- By “time” we mean *computational time*
- Suppose Alice and Bob are equally patient with respect to real time but Alice’s computer is 100 times as powerful as Bob’s. Then $\delta \sim 100\,\varepsilon$
- Discount factor is not just an index of patience, but also of computational power

The Discounted Game

- Let $G = (A_1, A_2, u_1, u_2)$ be a game
- The (ε, δ)-discounted version of G has
  - Actions: Probabilistic machines which take as input ε and δ, and output actions in A_1 (resp. A_2)
  - Payoffs: Alice's payoff corresponding to machines M_1 (Alice) and M_2 (Bob) outputting $a_1 \in A_1$ and $a_2 \in A_2$ resp. is $u_1(a_1,a_2)(1 - \varepsilon)^t$, where t is time taken for M_1 to output a_1

Uniform Equilibria

- A pair of strategies (S_1, S_2) for the discounted game is a uniform NE if neither player can gain *in the limit* as $\varepsilon, \delta \rightarrow 0$ by playing a different strategy
- Limit case interesting because
  - ε, δ are typically small
  - As computational power increases, ε and δ get smaller

Plan of the Talk

- Introduction
- The Model
- *Results*
- Future Directions

Finite Games

- Theorem: Let G be a finite game. For every NE of G, the discounted version of G has a uniform NE with the same payoffs in the limit

Infinite Games

- Theorem: Every countable game with bounded computable payoffs has a uniform NE
- Note that such games do not always have an NE or even an almost-NE (eg., Largest Number Game)

The Largest Number Game, Revisited

Alice	Bob	Payoff (to Alice)
M (Integer)	N (Integer)	100 if M > N, 50 if M=N, 0 otherwise

Largest Number game does not have an NE, or even an almost-NE if $\gamma < 50$

The Largest Number Game, Revisited

- All the uniform equilibria of Largest Number game yield payoff 0 for both players
- Example: both players play 2^{1/ε^2 + 1/δ^2}
- If more is known about relationship between ε and δ, eg., $\varepsilon >> \delta$, then there might be other equilibria yielding non-zero payoffs

The Factoring Game, Revisited

Alice	Bob	Payoff (to Bob)
M (Integer)	X,Y (Integers)	100 if M is prime or if 1 < X,Y < M and M = X * Y, 1 otherwise

Factoring Game has infinitely many Nash equilibria, in each of which Bob gets payoff 100 and Alice gets payoff 1 (Bob's strategy is simply to factor Alice's number)

Complexity Through Game Theory

- Tight connection between computational complexity of Factoring and uniform equilibrium payoffs of discounted Factoring game
- Let $\delta = \varepsilon^c$, for some $c > 1$, wlog
- Theorem: If Factoring is in time $o(n^c)$ on average, then every uniform NE of discounted game gives payoff 1 to Alice and 100 to Bob

Complexity Through Game Theory

- Theorem: Suppose there is no algorithm which runs in time n^c polylog(n) and solves Factoring on average for infinitely many input lengths. Then there is a uniform NE of discounted game giving payoff 100 to Alice and 1 to Bob.
- Proof idea: Consider strategy for Alice of outputting random number of size ~ 1/ε. Show that any strategy for Bob yielding payoff more than 1 in the limit yields factoring algorithm

A Spurious Equilibrium

- In the case where Factoring is hard, there is still a uniform NE where Bob wins
- This corresponds to Bob playing a brute-force Factoring algorithm
- However, in practice, we wouldn't expect this to happen – Bob's threat is not credible

Plan of the Talk

- Introduction
- The Model
- Results
- *Future Directions*

Future Directions: Refining the Model

- Defining a notion of subgame-perfection for discounted games
- An approach based on preference relations rather than real-number payoffs
- Capture bounded rationality not just in implementation but also in design

Future Directions: Applications of the Model

- Using discounting in choice situations (“flexible” or “anytime” algorithms)
- Perspective on foundations of cryptography, where protocol is treated as a game and adversary is modelled as bounded-rational
- Bounded rationality in extensive-form games