CS109 Great Ideas of Computer Science: Game Theory and Simulation

---

Announcements:

No HW, No Lab this week!

Lab WILL be held next Monday, and you will have another homework!

Lecture Outline

Part I. The Prisoner's Dilemma: Cooperate or Defect

- Being nice isn't rational!

Part II. The Iterated Prisoner's Dilemma

- Nice guys CAN finish first!

Part III. Is the Iterated Prisoner's Dilemma Realistic?

- Being nice is sometimes complicated…..

Part IV. Conclusions and Fun

Game Theory

A GAME is "a competitive activity involving skill, chance, or endurance on the part of two or more persons who play according to a set of rules, usually for their own amusement or for that of spectators" (Webster's)

Game THEORY is a branch of mathematics that studies multi-player skill and chance games. It attempts to understand how STRATEGIES adopted by individuals interact. Even simple choices made by individuals sometimes have complex outcomes…….

Let's Try a Game…..

Game theory provides interesting insights into behavior…..

Let's try this game:

- Close your eyes;
- When I say "NOW!" either remain still or raise your right hand;
- When I say "OK," open your eyes.

Outcome: If no one in the room raises his or her hand, I will give you all $100; but if anyone raises his or her hand, I will give the hand raisers $9 and the rest get nothing.

Prisoner's Dilemma

"Two suspects, A and B, are arrested. The police have insufficient evidence for a conviction, and, having separated both prisoners, visit each of them to offer the same deal:

- If one testifies for the prosecution against the other and the other remains silent, the betrayer goes free and the silent accomplice receives the full 10-year sentence.
- If both remain silent, both prisoners are sentenced to only six months in jail for a minor charge.
- If each betrays the other, each receives a five-year sentence.

Each prisoner must make the choice of whether to betray the other or to remain silent. However, neither prisoner knows for sure what choice the other prisoner will make. So this dilemma poses the question: How should the prisoners act?" (Wikipedia)

Let's make the PD a mathematical game….

Rules of the Game:

- There are two players, A and B;
- A move is C (cooperate/help) or D (defect/trick);
- Each player chooses a move without knowing the other player's move in advance;
- There are four possible outcomes (C & C, C & D, D & C, D & D);
- Rewards and punishments are determined as follows….

Let's make the PD a mathematical game….

Listed are MY rewards or fines; you have the same!

		What You do: Cooperate (Help)	What You do: Defect (Trick)
What I do	Cooperate	Pretty Good! Reward for mutual cooperation: **$300**	Horrible! Fine for being a sucker: **-$100**
	Defect	Excellent! Reward for fooling the sucker: **$500**	Pretty bad! Punishment for mutual defection: **-$10**

Since we're in MCS 109, we will use a spreadsheet to keep track of our winnings…… Any volunteers?

Prisoner's Dilemma

		What You do: Cooperate	What You do: Defect
What I do	Cooperate	Pretty Good! Reward for mutual cooperation: $300	Horrible! Fine for being a sucker: -$100
	Defect	Excellent! Reward for fooling the sucker: $500	Pretty bad! Punishment for mutual defection: -$10

There are two deterministic strategies (i.e., not based on chance):

- Selfish = Defect
- Altruistic = Cooperate

Note that the game is only one move!

Why is this interesting?

A RATIONAL player trying to maximize his own profit never cooperates!

Why?

- If your partner Cooperates, your best payoff is with Defect;
- If your partner Defects, your least fine is with Defect;

Ergo…. No matter what your partner does, you should defect. Defect is the only rational strategy!

		What I do: Cooperate	What I do: Defect
What you do	Cooperate	Fairly Good : Reward for mutual cooperation: $300	Very bad! Fine for being a sucker: -$100
	Defect	Very good! Reward for fooling the sucker: $500	Fairly bad! Punishment for mutual defection: -$10

So don't blame the hand raisers, they're just being rational……

Why is this interesting?

A RATIONAL player trying to maximize his profit will never cooperate!

Thus:

Two RATIONAL players will both defect and get -$10 each time;

Two RANDOM players will each get, on average, ($500+$300-$100-$10)/4 = $172.50.

Two IRRATIONAL players who decide to Cooperate each time each get $300 .

So we have proved mathematically that being irrational or acting randomly is better in this game than acting rationally! At least once…..

Why is this interesting?

The Prisoner's Dilemma is a familiar paradigm in modern life:

- Should nations build up their store of weapons or agree to disarm?
- Should nations curb their emission of greenhouse gasses?
- Should athletes take steroids or performance-enhancing drugs?
- Should you cheat on your schoolwork?
- Should organisms share food and mates?

Iterated Prisoner's Dilemma

But altruism exists widely at all levels (from genes to societies)! How can this be? Maybe our model is not complex enough!

- We live in communities, and such games are rarely played just once and for all, but are repeated over and over;
- If you repeat the game, you can CHANGE your choices based on what the other players did in the past; you can have a more complex strategy!

Iterated Prisoner's Dilemma = Play some number of rounds and total your winnings/losses at the end.

Again, a spreadsheet tallies our wins and losses……

Strategies for the IPD

A STRATEGY is now

- What to do the first time;
- What to do given the history of past rounds in this one game.

Iterated Prisoner's Dilemma

What strategy would YOU use?

Any volunteers?

Strategies for the IPD

Angel = Always cooperate Golden Rule: "Do unto others as you….."

Devil = Always defect "Do it to the other guy before he does it to you."

Trickster = Cooperate the first time, thereafter defect

Random = Flip a coin each time "Whatever….."

Grim Retaliator = Cooperate unless other player defects, then defect forever after

Tit-for-Tat (TFT) = Cooperate on first move; thereafter copy your opponents last move Silver rule: "Repay kindness with kindness but evil with justice….Eye for an eye…"

Accountant = Start randomly; if you win, keep same move; if you lose, change your move: "win-stay, lose-shift". "Watch the bottom line."

Exploitative TFT = Just like TFT, but if on last move you Defect and he Cooperates (the sucker!), then Defect again! "Never give a sucker an even break."

Tit-for-Two-Tats = Like TFT, but allow two defections before retaliate.

"If your neighbor striketh your left cheek, giveth him your right cheek as well…"

Strategies for the IPD

OK, so which is the best? "Well, it depends"

If we know exactly how many rounds will be played, say N, then Devil is the only rational strategy:

- Round N is last, so no retaliation will occur afterwards, so might as well Defect (same reasoning as for one single round);
- But then both players assume other will defect on last round, so only round N-1 is uncertain…. But he is going to retaliate on last round, so I might as well retaliate on next-to-last, and so on.

Only interesting case is when neither player knows how many rounds will be played. This is quite realistic, in fact.

Strategies for the IPD

It is instructive to consider how various strategies compare with each other:

Angel + Angel: Everybody wins

Angel + Devil: Devil wins big, Angel loses big

TFT + TFT: Everybody wins (same for any strategy that never spontaneously defects)

TFT + Devil: Fooled once, but not twice….

Accountant + Angel: Accountant might exploit an Angel big time…..

Accountant + Devil: Devil wins big, Accountant keeps changing his mind

Accountant + Accountant: Eventually will settle down and cooperate, but perhaps not optimally.

Strategies for the IPD

But which strategies do well in a complex, large, and varied population?

Robert Axelrod, a political scientist, performed an experiment in late 1970s to answer this question, and reported it in a famous article "The Evolution of Cooperation" in the journal Nature:

- Iterated Prisoner's Dilemma run by computer;
- 63 strategies (some written as very complex computer programs); each play was 200 rounds, and it ran for 1000 generations
- Winner was one of simplest: Tit-for-Tat!

Reason: Cooperating with other cooperators, but not being fooled forever by nasty strategies, is the best strategy!

Strategies for the IPD: TFT Wins Big!

TFT, a.k.a., Reciprocal Altruism, a.ka., The Brazen Rule: "If thy enemy inclines toward peace, do thou also incline toward peace" (from the Koran, quoted by Hillary's husband at the Israeli/Palestinian peace accords on September 13th, 1993):

Strategies for the IPD

But in a more complex model, TFT is NOT optimal!

Suppose strategies sometimes make mistakes, or behave with some probability (mistakes do happen!);

- Then TFT can get caught in an endless series of backbiting…..
- Tit-for-two-tats is the winner when mistakes occur.

Also, perhaps you cooperate or defect only with some probability:

- Probabilistic TFT: Cooperate on first move, then copy opponents move with 95% probability (might make a mistake!).
- Probabilistic Accountant: Win-stay, lose-shift, but only 90% of the time (think of it as an Accountant, but didn't go to SMG…)

Strategies for the IPD

In a probabilistic setting, with rewards being given as offspring, with mutation and crossover, after 200,000 rounds, Accountant wins out:

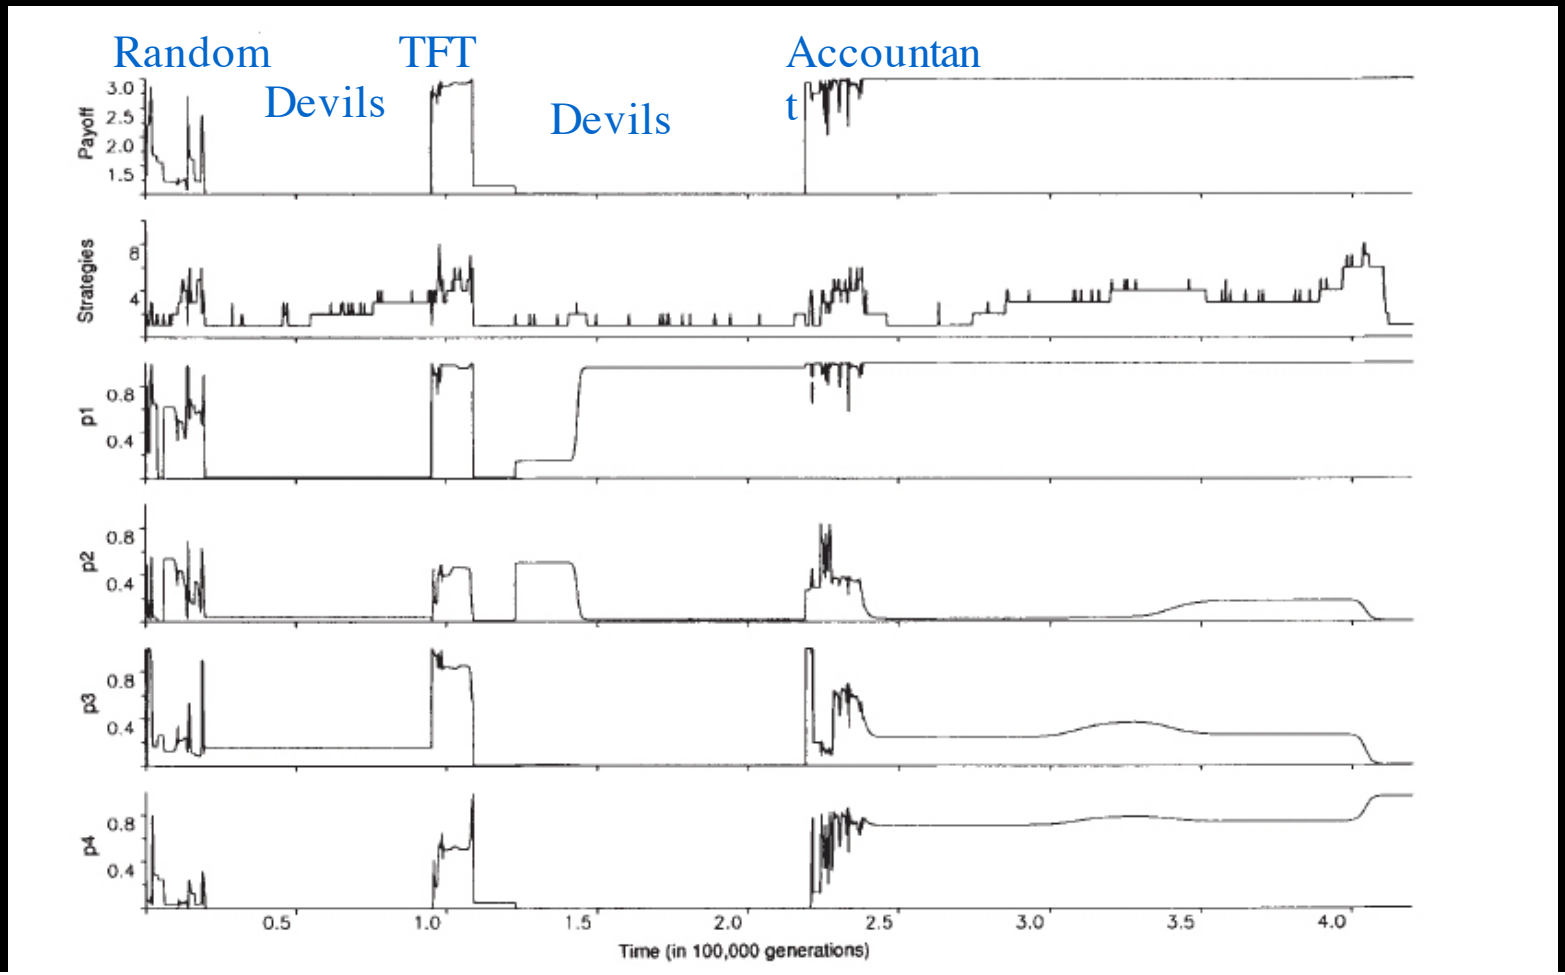

Conclusions

In the short term, being selfish pays;

Conclusions

In the short term, being selfish pays;

In the long term, being altruistic pays;

Conclusions

In the short term, being selfish pays;

In the long term, being altruistic pays;

But in an uncertain world, nothing beats knowing a good accountant….

Just for fun…..

The Prisoner's Dilemma has become a paradigm for quite a number of TV game shows…..

- Deal or No Deal (Australian)
- Big Brother (British)
- Friend or Foe (US)
- Golden Balls (British)