Leigh Tesfatsion/Econ 308

Notes on Axelrod's Iterated Prisoner's Dilemma (IPD) Tournaments

Basic References:

[**1**] L. Tesfatsion, **Game Theory: Basic Concepts and Terminology**
http://www.econ.iastate.edu/classes/econ308/tesfatsion/gamedef.308.pdf

[**2**] C. Cook, **Axelrod Tournament Demonstration Software**
http://www.econ.iastate.edu/tesfatsi/demos/axelrod/axelrodt.htm

ONE-STAGE SIMULTANEOUS-MOVE TWO-PLAYER PRISONER'S DILEMMA GAME

EMPLOYER

WORKER		C	D
	C	(40,40)	(10,60)
	D	(60,10)	(20,20)

Illustrative modeling of a work-site interaction between a Worker and Employer as a "Prisoner's Dilemma Game", D = Defect (Shirk), C = Cooperate (Work Hard), (P1,P2) = (Worker Payoff, Employer Payoff)

Summary of Axelrod's Iterated Prisoner's Dilemma (IPD) Tournaments

- *First Tournament:* Fourteen entries (computerized IPD strategies) in a round-robin IPD, including RANDOM introduced by Axelrod. Every entry played every other entry (plus RANDOM and a clone of itself) 200 times. Tournament was run five times to smooth out random effects.

- *Second Tournament:* Sixty-two entries, plus RANDOM, in same kind of tournament as first, except that every submitter had full information about the structure and results of the first tournament.

- *Ecological Tournament:* Entries (plus RANDOM) from the second tournament used as the initial conditions of an "evolutionary" tournament consisting of 1000 "generations." The number of strategies of type T in the population pool at the beginning of generation G was set equal to the total number of points won by strategies of type T in the previous generation G-1.

AND THE WINNERS WERE......??

- *First Tournament:* Fourteen entries (computerized IPD strategies) in a round-robin IPD, including RANDOM introduced by Axelrod. Every entry played every other entry (plus RANDOM and a clone of itself) 200 times. Tournament was run five times to smooth out random effects.

  Winner = Tit-For-Tat Strategy: *Start by cooperating. Then do whatever your partner did on the previous iteration.*

- *Second Tournament:* Sixty-two entries, plus RANDOM, in same kind of tournament as first, except that every submitter had full information about the structure and results of the first tournament.

  Winner = Tit-For-Tat Strategy.

- *Ecological Tournament:* Entries (plus RANDOM) from the second tournament used as the initial conditions of an "evolutionary" tournament consisting of 1000 "generations." The number of strategies of type T in the population pool at the beginning of generation G was set equal to the total number of points won by strategies of type T in the previous generation G-1.

  Winner = Tit-for-Tat Strategy.

What Properties Characterize Successful IPD Strategies?

A **(PURE) STRATEGY** for a player in a particular game is a complete contingency plan, i.e., a plan describing what move that player should take in each possible situation ("information state") that might arise for him.

In Axelrod's IPD tournaments, strategies exhibiting the following four properties tended to be more successful (i.e., to accumulate higher total payoffs), with the clear-cut winner being the Tit-for-Tat strategy.

- **Niceness:** Never be the first to defect.
- **Provocability:** Get mad quickly at defectors and retaliate.
- **Forgiveness:** Do not hold a grudge once you have vented your anger.
- **Clarity:** Act in ways that are straightforward for others to understand.

WHY Did These Properties Lead to Success in the Axelrod IPD Tournaments?

First Observation:

In any IPD game with FINITELY many iterations (known to all players), the only Nash equilibrium is (AllD,AllD), where AllD is the strategy of always choosing to defect regardless of what your rival does.

Second Observation:

This implies that AllD is the best response to AllD.

However, for any IPD game with at least two iterations, AllD is NOT a *dominant strategy*, i.e., AllD is NOT a best response to EVERY possible strategy the other player might choose.

Third Observation:

More generally, for any IPD game with at least two iterations, there is NO single best strategy S^* against ALL possible types of rivals.

For example, what would be your "best" choice of strategy in a 12-iteration IPD game played with a rival having each of the following strategies:

RANDOM: In each iteration I will flip an unbiased coin to decide whether I will cooperate with you (heads) or defect against you (tails).

TIT-FOR-TWO-TATS: I will start by cooperating with you in the first two iterations of the game. Starting in the third iteration, I will defect against you if you have defected against me in each of the previous two iterations; otherwise I will cooperate with you.

TRIGGER: I will start by cooperating with you and I will continue to cooperate with you until you defect against me. Once you defect against me, I will defect against you in all subsequent iterations.

In Axelrod's IPD tournaments, the pool of strategies was not known to participants in advance.

Thus, to be successful OVERALL, a strategy had to be capable of doing REASONABLY well with many DIFFERENT types of strategies.

Axelrod summarizes two major requirements for attaining this OVERALL success, as follows:

- **MINIMIZE NEGATIVE ECHO EFFECTS**

- **INDUCE COOPERATIVE BEHAVIOR**

More Precisely....

Take into account that any unprovoked defections on your part might lead to retaliatory defections by your rival. A good tactic is to be *NICE* (don't defect first).

Don't be a chump who lets others freely defect against you with no fear of punishment. That is, be *PROVOCABLE* in the sense that you retaliate quickly against defections.

However, your retaliation should be measured so you don't get into a vicious cycle of endless recriminations. You should therefore be *FORGIVING*, i.e., willing to return to cooperation whenever your rival does.

Also, make sure your intentions are communicated with *CLARITY* to your rivals. If your behavior is too complicated, you will appear to be RANDOM to your rivals – and the best response to RANDOM is ALLD!

Last but not least, for long-run success in the ECOLOGICAL tournament, you had better be able to play well with agents of your own type!

Do the lessons learned in Axelrod's IPD tournaments carry over to other forms of games?

Is there any general lesson here for real-world social and economic policy makers?

Is there any general lesson here for the design of computational agents?

PLAYER 2

		C	D
PLAYER 1	**C**	**(R,R)**	**(S,T)**
	D	**(T,S)**	**(P,P)**

General 2×2 **Symmetric Game**

Twenty-four different strategic games are obtained under the twenty-four possible orderings of the four payoffs R, S, T, and P by value.

Most of these games are not of great interest, but four stand out: the Prisoner's Dilemma Game, the Deadlock Game, the Chicken Game, and the Stag Hunt Game.

Question: Do any of the lessons learned from Axelrod's IPD tournaments carry over to these four types of games?

This issue can be examined using the Axelrod Tournament Demonstration Software by Chris Cook; see Ref.[2].

The Prisoner's Dilemma (PD) Game

$T > R > P > S \quad ([T+S]/2 < R)$

EXAMPLE:

:	C	D
C	(R=3, R=3)	(S=0, T=5)
D	(T=5, S=0)	(P=1, P=1)

R – Reward for cooperation;

T – Temptation to defect;

S – Sucker's payoff;

P – Punishment for defection

The Prisoner's Dilemma (PD) Game

T > R > P > S ([T+S]/2 < R)

EXAMPLE:

	C	D
C	(R=3, R=3)	(S=0, T=5)
D	(T=5, S=0)	(P=1, P=1)

ANALYSIS:

(D,D) is the unique Nash equilibrium, and D is a dominant strategy choice for each player. But (C,C) Pareto-***dominates*** (D,D). The three choice pairs (C,C), (C,D), and (D,C) are all Pareto optimal, but (C,C) is the most ***socially*** efficient choice pair.

The Dead Lock Game

T > P > R > S

EXAMPLE:

Choice:	C	D
C	(R=1, R=1)	(S=0, T=3)
D	(T=3, S=0)	(P=2, P=2)

The Dead Lock Game

T > P > R > S

EXAMPLE:

Choice:	C	D
C	(R=1, R=1)	(S=0, T=3)
D	(T=3, S=0)	(P=2, P=2)

ANALYSIS:

(D,D) is a Nash equilibrium and D is a dominant strategy choice for each player, as in the PD game. However, here (D,D) Pareto dominates (C,C), and indeed (D,D) is the most socially efficient outcome. The three choice pairs (D,D), (C,D), (D,C) are all Pareto optimal.

The Chicken Game

T > R > S > P

EXAMPLE:

Choice:	C	D
C	(R=2, R=2)	(S=1, T=3)
D	(T=3, S=1)	(P=0, P=0)

Example: Two drivers play a game of Chicken and see who chickens out and swerves.

“Rebel Without a Cause” (1955 movie) starring James Dean

C = Swerve; D = Drive Straight

The Chicken Game

T > R > S > P

EXAMPLE:

Choice:	C	D
C	(R=2, R=2)	(S=1, T=3)
D	(T=3, S=1)	(P=0, P=0)

ANALYSIS:

(C,D) and (D,C) are ***both*** Nash Equilibria, but ***neither*** Pareto dominates the other. ***Neither*** player has a dominant strategy choice. The three choice pairs (C,C), (D,C), and (C,D) are ***all*** Pareto optimal and equally socially efficient.

The Stag Hunt Game

R > T > P > S

EXAMPLE:

Choice:	C	D
C	(R=3, R=3)	(S=0, T=2)
D	(T=2, S=0)	(P=1, P=1)

ORIGINAL STORY (Jean Jacques Rousseau, French Philosopher):

Each hunter chooses either C (stay in position to hunt a stag - an adult deer) or D (go after a running rabbit). Hunting stags is quite challenging – to be successful it requires BOTH hunters to choose C and not be tempted by the running rabbit.

ANOTHER STORY: Next to the last day of the school, you and your friend decided to do something cool and show up on the last day of school with a crazy haircut. A night of indecision follows ....

The Stag Hunt Game

R > T > P > S

EXAMPLE:

Choice:	C	D
C	(R=3, R=3)	(S=0, T=2)
D	(T=2, S=0)	(P=1, P=1)

ANALYSIS:

(C,C) and (D,D) are "Pareto-ranked" Nash equilibria, in the following sense. Both are Nash equilibria, but (C,C) Pareto-dominates (D,D). Neither player has a dominant strategy choice – here each player is better off doing whatever the other is doing. The only Pareto optimal choice is (C,C), which is also socially efficient.