Leigh Tesfatsion/19 February 2003

Notes on 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 player played every other player (including 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 **STRATEGY** for a player in a particular game is a complete contingency plan, i.e., a plan describing what decision that player should make under each possible situation that might arise.

In Axelrod's 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 Tournaments?

First Observation:

In any IPD game with FINITELY many iterations (known to all players), the only Nash equilibrium is (AllD,AllD).

Second Observation:

This implies that AllD is the best response to AllD.

However, 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, there is NO single best strategy S^* for playing the IPD against ALL possible types of rivals.

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

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

TIT-FOR-TWO-TATS: I will start by cooperating in the first two iterations of the game. Starting in the third iteration, if you have cooperated in either of the two immediately preceding iterations, I will cooperate in the next iteration, nd if you have defected in both of the two immediately preceding iterations, I will defect in the next iteration.

TRIGGER: I will start by cooperating and I will continue to cooperate until you defect, after which I will defect 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!

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

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