CS3191 Section 5

Game Models

Andrea Schalk

School of Computer Science, University of Manchester

Game models

Section 5 investigates the use of games in modelling various scenarios.

Game models

Section 5 investigates the use of games in modelling various scenarios.

The game we will concentrate on is the Prisoner's Dilemma (also abbreviated as PD) game.

Game models

Section 5 investigates the use of games in modelling various scenarios.

The game we will concentrate on is the Prisoner's Dilemma (also abbreviated as PD) game.

We will look at variants of this game before we consider its repeated version.

Game models

Section 5 investigates the use of games in modelling various scenarios.

The game we will concentrate on is the Prisoner's Dilemma (also abbreviated as PD) game.

We will look at variants of this game before we consider its repeated version.

Lastly we explore the many strategies available when playing repeated Prisoner's Dilemma.

The Prisoner's Dilemma Revisited

The Prisoner's Dilemma revisited

Recall the Prisoner's Dilemma game.

		Joe	
		defect	cooperate
Fred	defect	$(-8,-8)$	$(0,-10)$
	cooperate	$(-10,0)$	$(-2,-2)$

The Prisoner's Dilemma revisited

Recall the Prisoner's Dilemma game.

		Joe	
		defect	cooperate
Fred	defect	$(-8, -8)$	$(0, -10)$
	cooperate	$(-10, 0)$	$(-2, -2)$

We know that $(1, 1)$, or ('defect','defect'), is the only equilibrium point.

The Prisoner's Dilemma revisited

		Joe	
		defect	cooperate
Fred	defect	$(-8,-8)$	$(0,-10)$
	cooperate	$(-10,0)$	$(-2,-2)$

We know that $(1,1)$, or ('defect','defect'), is the only equilibrium point.

Note that the defect strategy dominates the cooperate strategy.

The Prisoner's Dilemma revisited

		Joe	
		defect	cooperate
Fred	defect	$(-8, -8)$	$(0, -10)$
	cooperate	$(-10, 0)$	$(-2, -2)$

We know that $(1, 1)$, or ('defect','defect'), is the only equilibrium point.

Note that the defect strategy dominates the cooperate strategy.

Even if Joe knew in advance that Fred was going to defect, the best he could do is to defect in turn.

The Prisoner's Dilemma revisited

		Joe	
		defect	cooperate
Fred	defect	$(-8,-8)$	$(0,-10)$
	cooperate	$(-10,0)$	$(-2,-2)$

We know that $(1,1)$, or ('defect','defect'), is the only equilibrium point.

Note that the defect strategy dominates the cooperate strategy.

Even if Joe knew in advance that Fred was going to defect, the best he could do is to defect in turn.

Hence game theory's recommendation for this game is to defect.

The Prisoner's Dilemma revisited

		Joe	
		defect	cooperate
Fred	defect	$(-8,-8)$	$(0,-10)$
	cooperate	$(-10,0)$	$(-2,-2)$

We know that $(1,1)$, or ('defect','defect'), is the only equilibrium point.

Note that the defect strategy dominates the cooperate strategy.

Even if Joe knew in advance that Fred was going to defect, the best he could do is to defect in turn.

Hence game theory's recommendation for this game is to defect.

However, both parties would much prefer the outcome they get from both cooperating.

Generalizing the game

Generalization

There is nothing specific to the numbers that appear in our Prisoner's Dilemma game matrix. So let's assume we define a more general game.

Generalization

Assume that the game we are talking about is symmetric, and that both sides have the possibility to defect or to cooperate.

Generalization

Assume that the game we are talking about is symmetric, and that both sides have the possibility to defect or to cooperate.

We name the possible pay-offs as follows.

		Player 2	
		defect	cooperate
Player 1	defect	P	T
	cooperate	S	R

Generalization

We name the possible pay-offs as follows.

		Player 2	
		defect	cooperate
Player 1	defect	P	T
	cooperate	S	R

Symmetry means that the pay-off for Player 2 can be described by the transposed matrix

		Player 2	
		defect	cooperate
Player 1	defect	P	S
	cooperate	T	R

Generalization

We name the possible pay-offs as follows.

		Player 2	
		defect	cooperate
Player 1	defect	P	T
	cooperate	S	R

We would therefore normally give the pay-off matrix as

		Player 2	
		defect	cooperate
Player 1	defect	(P, P)	(T, S)
	cooperate	(S, T)	(R, R)

but information in second version can be derived from the first.

Generalization

We name the possible pay-offs as follows.

		Player 2	
		defect	cooperate
Player 1	defect	P	T
	cooperate	S	R

The numbers are typically read as

P: the Punishment for mutual defection;

Generalization

We name the possible pay-offs as follows.

		Player 2	
		defect	cooperate
Player 1	defect	P	T
	cooperate	S	R

The numbers are typically read as

P: the Punishment for mutual defection;

T: the Temptation to defect on a cooperating fellow player;

Generalization

We name the possible pay-offs as follows.

		Player 2	
		defect	cooperate
Player 1	defect	P	T
	cooperate	S	R

The numbers are typically read as

- P: the Punishment for mutual defection;
- T: the Temptation to defect on a cooperating fellow player;
- S: the 'Sucker's pay-off' for cooperating while the other defects;

Generalization

We name the possible pay-offs as follows.

		Player 2	
		defect	cooperate
Player 1	defect	P	T
	cooperate	S	R

The numbers are typically read as

- P: the Punishment for mutual defection;
- T: the Temptation to defect on a cooperating fellow player;
- S: the ‘Sucker’s pay-off’ for cooperating while the other defects;
- R: the Reward for mutual cooperation.

Generalization

We name the possible pay-offs as follows.

		Player 2	
		defect	cooperate
Player 1	defect	P	T
	cooperate	S	R

Assume that $T > R > P > S$.

Generalization

We name the possible pay-offs as follows.

		Player 2	
		defect	cooperate
Player 1	defect	P	T
	cooperate	S	R

Assume that $T > R > P > S$.

Then the game retains all its properties:

Generalization

We name the possible pay-offs as follows.

		Player 2	
		defect	cooperate
Player 1	defect	P	T
	cooperate	S	R

Assume that $T > R > P > S$.

Then the game retains all its properties:

- the defect strategy dominates the cooperate one;

Generalization

We name the possible pay-offs as follows.

		Player 2	
		defect	cooperate
Player 1	defect	P	T
	cooperate	S	R

Assume that $T > R > P > S$.

Then the game retains all its properties:

- the defect strategy dominates the cooperate one;
- the only equilibrium point is $(1, 1)$;

Generalization

We name the possible pay-offs as follows.

		Player 2	
		defect	cooperate
Player 1	defect	P	T
	cooperate	S	R

Assume that $T > R > P > S$.

Then the game retains all its properties:

- the defect strategy dominates the cooperate one;
- the only equilibrium point is $(1, 1)$;
- both players would prefer the outcome of $(2, 2)$.

Repeated games

Repeated games

So what happens if both sides know they will play the game more than once? Does the possibility of the other side cooperating in the future (given one's own present action) make the protagonists 'play nice'?

Repeated games

Here is a tree for two round PD.

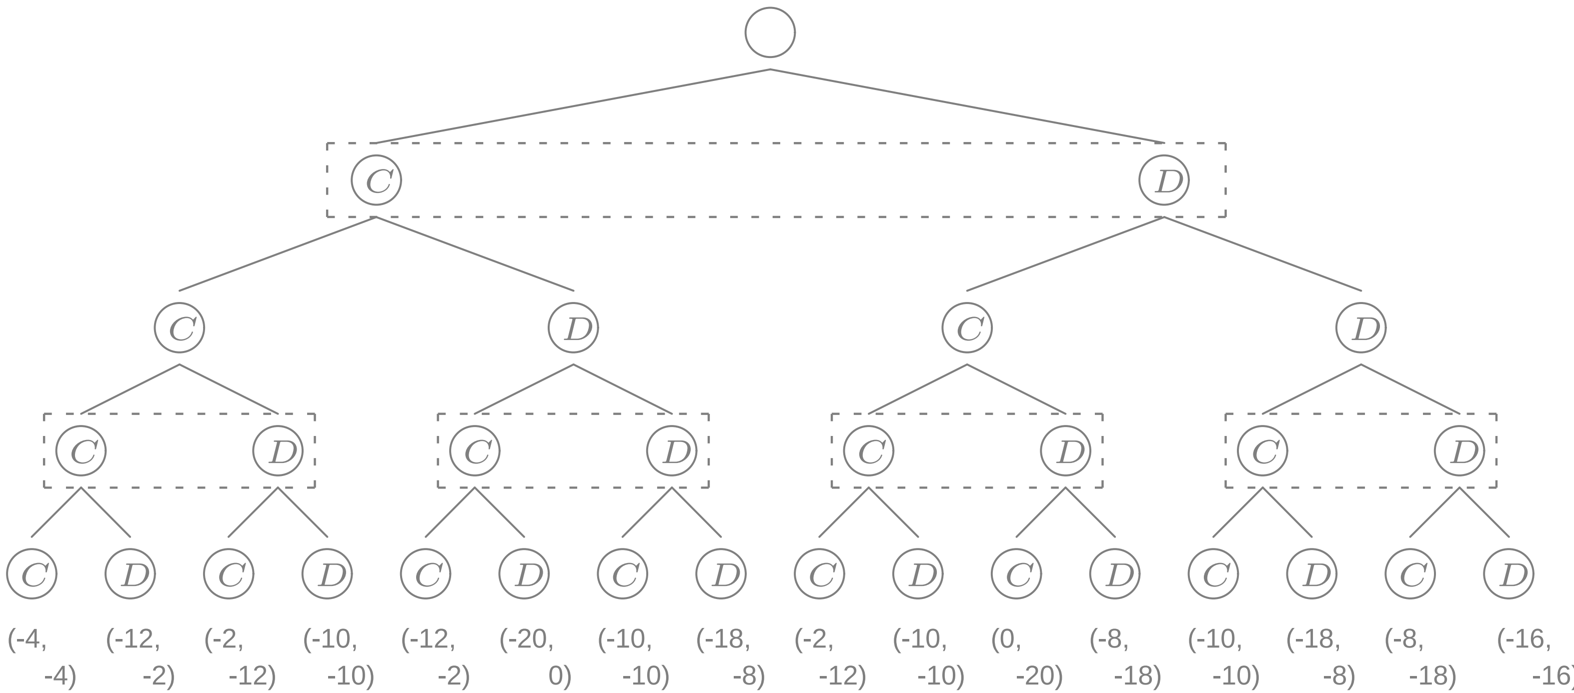

Repeated games

And here is the corresponding matrix game.

	$C(C\|C)$	$C(C\|D)$	$C(D\|C)$	$C(D\|D)$	$D(C\|C)$	$D(C\|D)$	$D(D\|C)$	$D(D\|D)$
$C\|C)$	(-4,-4)	(-4,-4)	(-12,-2)	(-12,-2)	(-12,-2)	(-12,-2)	(-20,0)	(-20,0)
$C\|D)$	(-4,-4)	(-4,-4)	(-12,-2)	(-12,-2)	(-10,-10)	(-10,-10)	(-18,-8)	(-18,-8)
$D\|C)$	(-2,-12)	(-2,-12)	(-10,-10)	(-10,-10)	(-12,-2)	(-12,-2)	(-20,0)	(-20,0)
$D\|D)$	(-2,-12)	(-2,-12)	(-10,-10)	(-10,-10)	(-10,-10)	(-10,-10)	(-18,-8)	(-18,-8)
$C\|C)$	(-2,-12)	(-10,-10)	(-2,-12)	(-10,-10)	(-10,-10)	(-18,-8)	(-10,-10)	(-18,-8)
$C\|D)$	(-2,-12)	(-10,-10)	(-10,-10)	(-10,-10)	(-8,-18)	(-16,-16)	(-8,-18)	(-16,-16)
$D\|C)$	(0,-20)	(-8,-18)	(0,-20)	(-8,-18)	(-10,-10)	(-18,-8)	(-10,-10)	(-18,-8)
$D\|D)$	(0,-20)	(-8,-18)	(0,-20)	(-8,-18)	(-8,-18)	(-16,-16)	(-8,-18)	(-16,-16)

Repeated games

And here is the corresponding matrix game.

	$C(C\|C)$	$C(C\|D)$	$C(D\|C)$	$C(D\|D)$	$D(C\|C)$	$D(C\|D)$	$D(D\|C)$	$D(D\|D)$
$C\|C)$	(-4,-4)	(-4,-4)	(-12,-2)	(-12,-2)	(-12,-2)	(-12,-2)	(-20,0)	(-20,0)
$C\|D)$	(-4,-4)	(-4,-4)	(-12,-2)	(-12,-2)	(-10,-10)	(-10,-10)	(-18,-8)	(-18,-8)
$D\|C)$	(-2,-12)	(-2,-12)	(-10,-10)	(-10,-10)	(-12,-2)	(-12,-2)	(-20,0)	(-20,0)
$D\|D)$	(-2,-12)	(-2,-12)	(-10,-10)	(-10,-10)	(-10,-10)	(-10,-10)	(-18,-8)	(-18,-8)
$C\|C)$	(-2,-12)	(-10,-10)	(-2,-12)	(-10,-10)	(-10,-10)	(-18,-8)	(-10,-10)	(-18,-8)
$C\|D)$	(-2,-12)	(-10,-10)	(-10,-10)	(-10,-10)	(-8,-18)	(-16,-16)	(-8,-18)	(-16,-16)
$D\|C)$	(0,-20)	(-8,-18)	(0,-20)	(-8,-18)	(-10,-10)	(-18,-8)	(-10,-10)	(-18,-8)
$D\|D)$	(0,-20)	(-8,-18)	(0,-20)	(-8,-18)	(-8,-18)	(-16,-16)	(-8,-18)	(-16,-16)

Note that the number of strategy increases.

Repeated games

And here is the corresponding matrix game.

	$C(C\|C)$	$C(C\|D)$	$C(D\|C)$	$C(D\|D)$	$D(C\|C)$	$D(C\|D)$	$D(D\|C)$	$D(D\|D)$
$C\|C)$	(-4,-4)	(-4,-4)	(-12,-2)	(-12,-2)	(-12,-2)	(-12,-2)	(-20,0)	(-20,0)
$C\|D)$	(-4,-4)	(-4,-4)	(-12,-2)	(-12,-2)	(-10,-10)	(-10,-10)	(-18,-8)	(-18,-8)
$D\|C)$	(-2,-12)	(-2,-12)	(-10,-10)	(-10,-10)	(-12,-2)	(-12,-2)	(-20,0)	(-20,0)
$D\|D)$	(-2,-12)	(-2,-12)	(-10,-10)	(-10,-10)	(-10,-10)	(-10,-10)	(-18,-8)	(-18,-8)
$C\|C)$	(-2,-12)	(-10,-10)	(-2,-12)	(-10,-10)	(-10,-10)	(-18,-8)	(-10,-10)	(-18,-8)
$C\|D)$	(-2,-12)	(-10,-10)	(-10,-10)	(-10,-10)	(-8,-18)	(-16,-16)	(-8,-18)	(-16,-16)
$D\|C)$	(0,-20)	(-8,-18)	(0,-20)	(-8,-18)	(-10,-10)	(-18,-8)	(-10,-10)	(-18,-8)
$D\|D)$	(0,-20)	(-8,-18)	(0,-20)	(-8,-18)	(-8,-18)	(-16,-16)	(-8,-18)	(-16,-16)

Question. How does the number of strategy increase with the number of rounds?

Repeated games

And here is the corresponding matrix game.

	$C(C\|C)$	$C(C\|D)$	$C(D\|C)$	$C(D\|D)$	$D(C\|C)$	$D(C\|D)$	$D(D\|C)$	$D(D\|D)$
$C\|C)$	(-4,-4)	(-4,-4)	(-12,-2)	(-12,-2)	(-12,-2)	(-12,-2)	(-20,0)	(-20,0)
$C\|D)$	(-4,-4)	(-4,-4)	(-12,-2)	(-12,-2)	(-10,-10)	(-10,-10)	(-18,-8)	(-18,-8)
$D\|C)$	(-2,-12)	(-2,-12)	(-10,-10)	(-10,-10)	(-12,-2)	(-12,-2)	(-20,0)	(-20,0)
$D\|D)$	(-2,-12)	(-2,-12)	(-10,-10)	(-10,-10)	(-10,-10)	(-10,-10)	(-18,-8)	(-18,-8)
$C\|C)$	(-2,-12)	(-10,-10)	(-2,-12)	(-10,-10)	(-10,-10)	(-18,-8)	(-10,-10)	(-18,-8)
$C\|D)$	(-2,-12)	(-10,-10)	(-10,-10)	(-10,-10)	(-8,-18)	(-16,-16)	(-8,-18)	(-16,-16)
$D\|C)$	(0,-20)	(-8,-18)	(0,-20)	(-8,-18)	(-10,-10)	(-18,-8)	(-10,-10)	(-18,-8)
$D\|D)$	(0,-20)	(-8,-18)	(0,-20)	(-8,-18)	(-8,-18)	(-16,-16)	(-8,-18)	(-16,-16)

Question. How does the number of strategy increase with the number of rounds?

The number of strategies is exponential in the number of rounds.

Equilibrium points

	$C(C\|C)$	$C(C\|D)$	$C(D\|C)$	$C(D\|D)$	$D(C\|C)$	$D(C\|D)$	$D(D\|C)$	$D(D\|D)$
$C\|C)$	(-4,-4)	(-4,-4)	(-12,-2)	(-12,-2)	(-12,-2)	(-12,-2)	(-20,0)	(-20,0)
$C\|D)$	(-4,-4)	(-4,-4)	(-12,-2)	(-12,-2)	(-10,-10)	(-10,-10)	(-18,-8)	(-18,-8)
$D\|C)$	(-2,-12)	(-2,-12)	(-10,-10)	(-10,-10)	(-12,-2)	(-12,-2)	(-20,0)	(-20,0)
$D\|D)$	(-2,-12)	(-2,-12)	(-10,-10)	(-10,-10)	(-10,-10)	(-10,-10)	(-18,-8)	(-18,-8)
$C\|C)$	(-2,-12)	(-10,-10)	(-2,-12)	(-10,-10)	(-10,-10)	(-18,-8)	(-10,-10)	(-18,-8)
$C\|D)$	(-2,-12)	(-10,-10)	(-10,-10)	(-10,-10)	(-8,-18)	(-16,-16)	(-8,-18)	(-16,-16)
$D\|C)$	(0,-20)	(-8,-18)	(0,-20)	(-8,-18)	(-10,-10)	(-18,-8)	(-10,-10)	(-18,-8)
$D\|D)$	(0,-20)	(-8,-18)	(0,-20)	(-8,-18)	(-8,-18)	(-16,-16)	(-8,-18)	(-16,-16)

Equilibrium points

	$C(C\|C)$	$C(C\|D)$	$C(D\|C)$	$C(D\|D)$	$D(C\|C)$	$D(C\|D)$	$D(D\|C)$	$D(D\|D)$
$C\|C)$	(-4,-4)	(-4,-4)	(-12,-2)	(-12,-2)	(-12,-2)	(-12,-2)	(-20,0)	(-20,0)
$C\|D)$	(-4,-4)	(-4,-4)	(-12,-2)	(-12,-2)	(-10,-10)	(-10,-10)	(-18,-8)	(-18,-8)
$D\|C)$	(-2,-12)	(-2,-12)	(-10,-10)	(-10,-10)	(-12,-2)	(-12,-2)	(-20,0)	(-20,0)
$D\|D)$	(-2,-12)	(-2,-12)	(-10,-10)	(-10,-10)	(-10,-10)	(-10,-10)	(-18,-8)	(-18,-8)
$C\|C)$	(-2,-12)	(-10,-10)	(-2,-12)	(-10,-10)	(-10,-10)	(-18,-8)	(-10,-10)	(-18,-8)
$C\|D)$	(-2,-12)	(-10,-10)	(-10,-10)	(-10,-10)	(-8,-18)	(-16,-16)	(-8,-18)	(-16,-16)
$D\|C)$	(0,-20)	(-8,-18)	(0,-20)	(-8,-18)	(-10,-10)	(-18,-8)	(-10,-10)	(-18,-8)
$D\|D)$	(0,-20)	(-8,-18)	(0,-20)	(-8,-18)	(-8,-18)	(-16,-16)	(-8,-18)	(-16,-16)

The only equilibrium points of this game is $(8,8)$.

Equilibrium points

	$C(C\|C)$	$C(C\|D)$	$C(D\|C)$	$C(D\|D)$	$D(C\|C)$	$D(C\|D)$	$D(D\|C)$	$D(D\|D)$
$C\|C)$	(-4,-4)	(-4,-4)	(-12,-2)	(-12,-2)	(-12,-2)	(-12,-2)	(-20,0)	(-20,0)
$C\|D)$	(-4,-4)	(-4,-4)	(-12,-2)	(-12,-2)	(-10,-10)	(-10,-10)	(-18,-8)	(-18,-8)
$D\|C)$	(-2,-12)	(-2,-12)	(-10,-10)	(-10,-10)	(-12,-2)	(-12,-2)	(-20,0)	(-20,0)
$D\|D)$	(-2,-12)	(-2,-12)	(-10,-10)	(-10,-10)	(-10,-10)	(-10,-10)	(-18,-8)	(-18,-8)
$C\|C)$	(-2,-12)	(-10,-10)	(-2,-12)	(-10,-10)	(-10,-10)	(-18,-8)	(-10,-10)	(-18,-8)
$C\|D)$	(-2,-12)	(-10,-10)	(-10,-10)	(-10,-10)	(-8,-18)	(-16,-16)	(-8,-18)	(-16,-16)
$D\|C)$	(0,-20)	(-8,-18)	(0,-20)	(-8,-18)	(-10,-10)	(-18,-8)	(-10,-10)	(-18,-8)
$D\|D)$	(0,-20)	(-8,-18)	(0,-20)	(-8,-18)	(-8,-18)	(-16,-16)	(-8,-18)	(-16,-16)

The only equilibrium points of this game is $(8, 8)$.

It leads to a play where both sides keep defecting, and the pay-off gained is -16.

Equilibrium points

	$C(C\|C)$	$C(C\|D)$	$C(D\|C)$	$C(D\|D)$	$D(C\|C)$	$D(C\|D)$	$D(D\|C)$	$D(D\|D)$
$C\|C)$	(-4,-4)	(-4,-4)	(-12,-2)	(-12,-2)	(-12,-2)	(-12,-2)	(-20,0)	(-20,0)
$C\|D)$	(-4,-4)	(-4,-4)	(-12,-2)	(-12,-2)	(-10,-10)	(-10,-10)	(-18,-8)	(-18,-8)
$D\|C)$	(-2,-12)	(-2,-12)	(-10,-10)	(-10,-10)	(-12,-2)	(-12,-2)	(-20,0)	(-20,0)
$D\|D)$	(-2,-12)	(-2,-12)	(-10,-10)	(-10,-10)	(-10,-10)	(-10,-10)	(-18,-8)	(-18,-8)
$C\|C)$	(-2,-12)	(-10,-10)	(-2,-12)	(-10,-10)	(-10,-10)	(-18,-8)	(-10,-10)	(-18,-8)
$C\|D)$	(-2,-12)	(-10,-10)	(-10,-10)	(-10,-10)	(-8,-18)	(-16,-16)	(-8,-18)	(-16,-16)
$D\|C)$	(0,-20)	(-8,-18)	(0,-20)	(-8,-18)	(-10,-10)	(-18,-8)	(-10,-10)	(-18,-8)
$D\|D)$	(0,-20)	(-8,-18)	(0,-20)	(-8,-18)	(-8,-18)	(-16,-16)	(-8,-18)	(-16,-16)

The only equilibrium points of this game is $(8,8)$.

It leads to a play where both sides keep defecting, and the pay-off gained is -16. If they both kept cooperating they would only lose 4 apiece!

Equilibrium points

	$C(C\|C)$	$C(C\|D)$	$C(D\|C)$	$C(D\|D)$	$D(C\|C)$	$D(C\|D)$	$D(D\|C)$	$D(D\|D)$
$C\|C)$	(-4,-4)	(-4,-4)	(-12,-2)	(-12,-2)	(-12,-2)	(-12,-2)	(-20,0)	(-20,0)
$C\|D)$	(-4,-4)	(-4,-4)	(-12,-2)	(-12,-2)	(-10,-10)	(-10,-10)	(-18,-8)	(-18,-8)
$D\|C)$	(-2,-12)	(-2,-12)	(-10,-10)	(-10,-10)	(-12,-2)	(-12,-2)	(-20,0)	(-20,0)
$D\|D)$	(-2,-12)	(-2,-12)	(-10,-10)	(-10,-10)	(-10,-10)	(-10,-10)	(-18,-8)	(-18,-8)
$C\|C)$	(-2,-12)	(-10,-10)	(-2,-12)	(-10,-10)	(-10,-10)	(-18,-8)	(-10,-10)	(-18,-8)
$C\|D)$	(-2,-12)	(-10,-10)	(-10,-10)	(-10,-10)	(-8,-18)	(-16,-16)	(-8,-18)	(-16,-16)
$D\|C)$	(0,-20)	(-8,-18)	(0,-20)	(-8,-18)	(-10,-10)	(-18,-8)	(-10,-10)	(-18,-8)
$D\|D)$	(0,-20)	(-8,-18)	(0,-20)	(-8,-18)	(-8,-18)	(-16,-16)	(-8,-18)	(-16,-16)

The only equilibrium points of this game is $(8, 8)$.

If we play more rounds then there are more equilibrium points, see Exercise 22. In the five round Prisoner's Dilemma game there are at least 3 equilibrium points, all of which lead to play where both sides keep defecting.

Removing undesirable equ. points

We can get rid of these artificial equilibrium points by demanding that an equilibrium point strategy work for every sub-tree of the original game tree.

Removing undesirable equ. points

We can get rid of these artificial equilibrium points by demanding that an equilibrium point strategy work for every sub-tree of the original game tree.

Definition 11 *Let a a sub-game of a given game be obtained by choosing any position in the game tree and considering it as the root, with the rest of the tree given by the part of the original tree which is below it.*

A sub-game equilibrium point in a j-person game consists of a j-tuple of strategies, one for each player, such that this tuple gives an equilibrium point for each sub-game of the game whose root is reachable when playing these strategies.

Removing undesirable equ. points

We can get rid of these artificial equilibrium points by demanding that an equilibrium point strategy work for every sub-tree of the original game tree.

Definition 11 *Let a a sub-game of a given game be obtained by choosing any position in the game tree and considering it as the root, with the rest of the tree given by the part of the original tree which is below it.*

A sub-game equilibrium point in a j-person game consists of a j-tuple of strategies, one for each player, such that this tuple gives an equilibrium point for each sub-game of the game whose root is reachable when playing these strategies.

We will see an example of an equilibrium point which is not a sub-game equilibrium point in Exercise 22.

Removing undesirable equ. points

We can get rid of these artificial equilibrium points by demanding that an equilibrium point strategy work for every sub-tree of the original game tree.

Definition 11 *Let a a sub-game of a given game be obtained by choosing any position in the game tree and considering it as the root, with the rest of the tree given by the part of the original tree which is below it.*

A sub-game equilibrium point in a j-person game consists of a j-tuple of strategies, one for each player, such that this tuple gives an equilibrium point for each sub-game of the game whose root is reachable when playing these strategies.

We will see an example of an equilibrium point which is not a sub-game equilibrium point in Exercise 22.

Proposition 5.1 *The only sub-game equilibrium point in a game of Prisoner's Dilemma with finitely many repetitions consists of choosing the 'always defect' strategy for each player.*

Whither now?

We have seen that even with higher stakes in the repeated PD game, it makes sense for both sides to keep defecting.

Whither now?

We have seen that even with higher stakes in the repeated PD game, it makes sense for both sides to keep defecting.

Here is an argument (known as backwards induction): The last game in a series is just like ordinary PD, so both players should defect. But if they know that this will happen in the last round, they should also both defect in the last but one round—but this argument works backwards all the way until the first round is reached.

Whither now?

We have seen that even with higher stakes in the repeated PD game, it makes sense for both sides to keep defecting.

This is not how people behave in every-day life (compare Hofstadter's PD game with 20 players).

Whither now?

We have seen that even with higher stakes in the repeated PD game, it makes sense for both sides to keep defecting.

The PD game has turned out to be very useful in modelling real-world interactions, and we are not done with it yet by a long margin.

Whither now?

We have seen that even with higher stakes in the repeated PD game, it makes sense for both sides to keep defecting.

When playing the repeated version we want to make sure that the players cannot profit from deciding to alternate defecting and cooperating, so we assume that

$$\frac{T+S}{2} < R.$$

This means that, over two rounds getting the temptation pay-off combined with the sucker's one a player is worse off than getting the reward twice.

Whither now?

We have seen that even with higher stakes in the repeated PD game, it makes sense for both sides to keep defecting.

When playing the repeated version we want to make sure that the players cannot profit from deciding to alternate defecting and cooperating, so we assume that

$$\frac{T+S}{2} < R.$$

This means that, over two rounds getting the temptation pay-off combined with the sucker's one a player is worse off than getting the reward twice.

There are a great number of strategies when playing the iterated PD game: For round $n+1$ our decision is allowed to depend on the 4^n possible histories of the game!

Experiment

	defect	cooperate
defect	$(1, 1)$	$(5, 0)$
cooperate	$(0, 5)$	$(3, 3)$

A computer tournament

Successful strategies?

The political scientist Robert Axelrod was wondering which strategies would be successful when playing the iterated PD game.

Successful strategies?

The political scientist Robert Axelrod was wondering which strategies would be successful when playing the iterated PD game.

In 1979 he had the obvious(?) idea that it would be helpful if he had lots of people thinking about this, and cleverly attracted them into taking part in a computer tournament.

Successful strategies?

The political scientist Robert Axelrod was wondering which strategies would be successful when playing the iterated PD game.

In 1979 he had the obvious(?) idea that it would be helpful if he had lots of people thinking about this, and cleverly attracted them into taking part in a computer tournament.

He invited game theorists from a number of different academic fields such as economics, psychology, political science, mathematics and sociology to send him a computer program to play repeated Prisoner's Dilemma in a round robin tournament.

Successful strategies?

The political scientist Robert Axelrod was wondering which strategies would be successful when playing the iterated PD game.

In 1979 he had the obvious(?) idea that it would be helpful if he had lots of people thinking about this, and cleverly attracted them into taking part in a computer tournament.

At the same time he gave them the results of a preliminary tournament played with fairly simple strategies.

Successful strategies?

The political scientist Robert Axelrod was wondering which strategies would be successful when playing the iterated PD game.

In 1979 he had the obvious(?) idea that it would be helpful if he had lots of people thinking about this, and cleverly attracted them into taking part in a computer tournament.

In the proposed contest each program would play each of the other programs, as well as a clone of itself, for 200 rounds. Whichever program had the most points at the end would win. The entire tournament would be run five times.

Successful strategies?

The political scientist Robert Axelrod was wondering which strategies would be successful when playing the iterated PD game.

In 1979 he had the obvious(?) idea that it would be helpful if he had lots of people thinking about this, and cleverly attracted them into taking part in a computer tournament.

In the proposed contest each program would play each of the other programs, as well as a clone of itself, for 200 rounds. Whichever program had the most points at the end would win. The entire tournament would be run five times.

The following pay-off matrix was used:

	defect	cooperate
defect	1	5
cooperate	0	3

Successful strategies?

The political scientist Robert Axelrod was wondering which strategies would be successful when playing the iterated PD game.

In 1979 he had the obvious(?) idea that it would be helpful if he had lots of people thinking about this, and cleverly attracted them into taking part in a computer tournament.

In the proposed contest each program would play each of the other programs, as well as a clone of itself, for 200 rounds. Whichever program had the most points at the end would win. The entire tournament would be run five times.

The following pay-off matrix was used:

	defect	cooperate
defect	1	5
cooperate	0	3

Question. What would you have submitted? Do you think your strategy might have won?

Submissions

He had fourteen submissions of programs from 4 to 77 lines (of Basic). He himself introduced another entry, the strategy which plays totally randomly, by 'tossing a coin' before each game. It is referred to as RANDOM.

Submissions

He had fourteen submissions of programs from 4 to 77 lines (of Basic). He himself introduced another entry, the strategy which plays totally randomly, by ‘tossing a coin’ before each game. It is referred to as RANDOM.

Note that here strategies are competing with each other for points, which is a different point of view from the one employed when considering equilibria. Now it is important how a strategy does when compared to others, whereas before we only worried about what would happen if one player unilaterally moved away from an equilibrium point.

Submissions

He had fourteen submissions of programs from 4 to 77 lines (of Basic). He himself introduced another entry, the strategy which plays totally randomly, by 'tossing a coin' before each game. It is referred to as RANDOM.

Note that here strategies are competing with each other for points, which is a different point of view from the one employed when considering equilibria. Now it is important how a strategy does when compared to others, whereas before we only worried about what would happen if one player unilaterally moved away from an equilibrium point.

It is clear that given the conditions imposed by Axelrod there is no 'best' strategy, that is one that will win such a tournament against all comers. Instead which strategy is best depends on the population, that is the other strategies present and their numbers.

The winner

To everybody's surprise, the shortest program won. It had been submitted by the psychologist and philosopher Anatol Rapoport. It looks quite simple:

- On the first move, cooperate;
- then do whatever the other player did last.

The winner

To everybody's surprise, the shortest program won. It had been submitted by the psychologist and philosopher Anatol Rapoport. It looks quite simple:

- On the first move, cooperate;
- then do whatever the other player did last.

This strategy is known as TITFORTAT, because it cooperates until the other player defects, and then it retaliates by defecting in the next round. What happens subsequently depends on the other player: If the other player goes back to cooperating, so does TITFORTAT. A very simple idea, but actually more subtle than you might think at first.

The winner

To everybody’s surprise, the shortest program won. It had been submitted by the psychologist and philosopher Anatol Rapoport. It looks quite simple:

- On the first move, cooperate;
- then do whatever the other player did last.

This strategy is known as TitForTat, because it cooperates until the other player defects, and then it retaliates by defecting in the next round. What happens subsequently depends on the other player: If the other player goes back to cooperating, so does TitForTat. A very simple idea, but actually more subtle than you might think at first.
If over 200 games both strategies always cooperate, they can each get 600 points. If both strategies always defect they each get 200 points.

The winner

To everybody's surprise, the shortest program won. It had been submitted by the psychologist and philosopher Anatol Rapoport. It looks quite simple:

- On the first move, cooperate;
- then do whatever the other player did last.

This strategy is known as TITFORTAT, because it cooperates until the other player defects, and then it retaliates by defecting in the next round. What happens subsequently depends on the other player: If the other player goes back to cooperating, so does TITFORTAT. A very simple idea, but actually more subtle than you might think at first.
If over 200 games both strategies always cooperate, they can each get 600 points. If both strategies always defect they each get 200 points.

TITFORTAT scored 504 points on average over the 200 rounds.

The winner

To everybody's surprise, the shortest program won. It had been submitted by the psychologist and philosopher Anatol Rapoport. It looks quite simple:

- On the first move, cooperate;
- then do whatever the other player did last.

If over 200 games both strategies always cooperate, they can each get 600 points. If both strategies always defect they each get 200 points.

Note that TITFORTAT will never outscore a single one of its opponents in their bout of 200 games:

The winner

To everybody's surprise, the shortest program won. It had been submitted by the psychologist and philosopher Anatol Rapoport. It looks quite simple:

- On the first move, cooperate;
- then do whatever the other player did last.

If over 200 games both strategies always cooperate, they can each get 600 points. If both strategies always defect they each get 200 points.

Note that TITFORTAT will never outscore a single one of its opponents in their bout of 200 games:

- If both strategies keep cooperating, they will both score 600 points.

The winner

To everybody's surprise, the shortest program won. It had been submitted by the psychologist and philosopher Anatol Rapoport. It looks quite simple:

- On the first move, cooperate;
- then do whatever the other player did last.

If over 200 games both strategies always cooperate, they can each get 600 points. If both strategies always defect they each get 200 points.

Note that TITFORTAT will never outscore a single one of its opponents in their bout of 200 games:

- If both strategies keep cooperating, they will both score 600 points.
- If the other strategy defects first then the best TITFORTAT can hope for is eventually to make up the disadvantage of five points caused thus.

Joss

Joss (named after its creator) behaves just the same as TitForTat but every so often it throws in a random defection and afterwards it goes back to cooperating.

		TitForTat	Joss

Joss

Joss (named after its creator) behaves just the same as TitForTat but every so often it throws in a random defection and afterwards it goes back to cooperating.

TitForTat	Joss	TitForTat	Joss
C	C	3	3
$\vdots$	$\vdots$	$\vdots$	$\vdots$
C	C	3	3
C	D	0	5

When TitForTat and Joss play against each other the following effect occurs: They both cooperate, racking up points, until Joss throws in one of its random defections.

Joss

Joss (named after its creator) behaves just the same as TitForTat but every so often it throws in a random defection and afterwards it goes back to cooperating.

TitForTat	Joss	TitForTat	Joss
C	C	3	3
$\vdots$	$\vdots$	$\vdots$	$\vdots$
C	C	3	3
C	D	0	5
D	C	5	0
C	D	0	5
D	C	5	0
$\vdots$	$\vdots$	$\vdots$	$\vdots$
C	D	0	5
D	D	1	1

When TitForTat and Joss play against each other the following effect occurs: They both cooperate, racking up points, until Joss throws in one of its random defections.

They are locked in a defection/cooperation loop until eventually, Joss will sneak in another defection.

Joss

TitForTat	Joss	TitForTat	Joss
C	C	3	3
$\vdots$	$\vdots$	$\vdots$	$\vdots$
C	C	3	3
C	D	0	5
D	C	5	0
C	D	0	5
D	C	5	0
$\vdots$	$\vdots$	$\vdots$	$\vdots$
C	D	0	5
D	D	1	1
D	D	1	1
D	D	1	1
$\ldots$	$\ldots$	$\ldots$	$\ldots$

When TitForTat and Joss play against each other the following effect occurs: They both cooperate, racking up points, until Joss throws in one of its random defections.

They are locked in a defection/cooperation loop until eventually, Joss will sneak in another defection.

At this point, both programs will defect forever, and their pay-off goes down to just one point per game.

Joss

TitForTat and Joss only scored 225/230 points against each other, and Joss had similar results against many other strategies. Ultimately it was hurt by its deviousness!

Downing

The psychologists Robert Downing submitted a strategy known as DOWNING. It attempts to find out how its playing partner reacts to defections, and then defects as often as it thinks it can get away with.

DOWNING

The psychologists Robert Downing submitted a strategy known as DOWNING. It attempts to find out how its playing partner reacts to defections, and then defects as often as it thinks it can get away with.

It adjusts two variables which denote

- the probability that the other side will cooperate after itself defected in the previous round;
- the probability that the other side will cooperate after itself has cooperated.

Downing

The psychologists Robert Downing submitted a strategy known as DOWNING. It attempts to find out how its playing partner reacts to defections, and then defects as often as it thinks it can get away with.

It adjusts two variables which denote

- the probability that the other side will cooperate after itself defected in the previous round;
- the probability that the other side will cooperate after itself has cooperated.

These two probabilities are then adjusted in the course of play, starting with $1/2$.

Downing

The psychologists Robert Downing submitted a strategy known as Downing. It attempts to find out how its playing partner reacts to defections, and then defects as often as it thinks it can get away with.

It adjusts two variables which denote

- the probability that the other side will cooperate after itself defected in the previous round;
- the probability that the other side will cooperate after itself has cooperated.

These two probabilities are then adjusted in the course of play, starting with $1/2$. If the two probabilities are similar then Downing will tend to defect, having decided that it makes no difference to the other side's behaviour.

DOWNING

The psychologists Robert Downing submitted a strategy known as DOWNING. It attempts to find out how its playing partner reacts to defections, and then defects as often as it thinks it can get away with.

It adjusts two variables which denote

- the probability that the other side will cooperate after itself defected in the previous round;
- the probability that the other side will cooperate after itself has cooperated.

These two probabilities are then adjusted in the course of play, starting with $1/2$. If the two probabilities are similar then DOWNING will tend to defect, having decided that it makes no difference to the other side's behaviour. If the other player instead tends to retaliate after a defection, but cooperates when DOWNING cooperates, it will decide that it is more prudent to cooperate.

Downing

The psychologists Robert Downing submitted a strategy known as Downing. It attempts to find out how its playing partner reacts to defections, and then defects as often as it thinks it can get away with.

It adjusts two variables which denote

- the probability that the other side will cooperate after itself defected in the previous round;
- the probability that the other side will cooperate after itself has cooperated.

Due to the way the two probabilities are set at the start (when Downing has no real information about its playing partner) it typically starts with defecting twice. In the tournament that led to other programs punishing it accordingly. As a consequence the program only came 10th in the tournament, gaining only 390 points on average against other entries.

Other results

Unsurprisingly RANDOM came last.

Other results

Unsurprisingly RANDOM came last. At least all the programs, into which much thought must have gone, were better than making random choices (but not all outperformed it by much!).

Other results

Unsurprisingly RANDOM came last. At least all the programs, into which much thought must have gone, were better than making random choices (but not all outperformed it by much!).

A strategy called GRUDGE, which cooperates until the other side defects and then defects forever, came 7th.

Other results

Unsurprisingly RANDOM came last. At least all the programs, into which much thought must have gone, were better than making random choices (but not all outperformed it by much!).

A strategy called GRUDGE, which cooperates until the other side defects and then defects forever, came 7th.

All the participants had been told that TITFORTAT had won a preliminary tournament, but only one of them decided to submit it.

Other results

Unsurprisingly RANDOM came last. At least all the programs, into which much thought must have gone, were better than making random choices (but not all outperformed it by much!).

A strategy called GRUDGE, which cooperates until the other side defects and then defects forever, came 7th.

All the participants had been told that TITFORTAT had won a preliminary tournament, but only one of them decided to submit it.

The others presumably thought they were clever enough to produce something better than that!

Properties of successful strategies

The point of the tournament was to find out which strategies did well in practice, and Axelrod spent some time analysing the results.

Properties of successful strategies

The point of the tournament was to find out which strategies did well in practice, and Axelrod spent some time analysing the results.

The eight top entrants were nice, that is they never defected first.

Properties of successful strategies

The point of the tournament was to find out which strategies did well in practice, and Axelrod spent some time analysing the results.

The eight top entrants were nice, that is they never defected first. All of them scored more than 470 points on average, whereas the next best strategy only managed just over 400 points.

Properties of successful strategies

The point of the tournament was to find out which strategies did well in practice, and Axelrod spent some time analysing the results.

The eight top entrants were nice, that is they never defected first. All of them scored more than 470 points on average, whereas the next best strategy only managed just over 400 points. Two nice strategies will score a ‘perfect’ 600 points against each other.

Properties of successful strategies

The point of the tournament was to find out which strategies did well in practice, and Axelrod spent some time analysing the results.

The eight top entrants were nice, that is they never defected first. All of them scored more than 470 points on average, whereas the next best strategy only managed just over 400 points. Two nice strategies will score a 'perfect' 600 points against each other.

Strategies that did well were forgiving, that is they might punish the other side for defecting but they would not hold a grudge forever.

Would-be winners

Axelrod also found that other strategies, had they been present, would have won.

Would-be winners

Axelrod also found that other strategies, had they been present, would have won.

One example of this is TitForTwoTats, a strategy that will cooperate unless the other player has defected twice in a row in which case it will defect (but just once).

Would-be winners

Axelrod also found that other strategies, had they been present, would have won.

One example of this is TitForTwoTats, a strategy that will cooperate unless the other player has defected twice in a row in which case it will defect (but just once). Ironically, this strategy was given to all contestants as a sample program which they simply could have copied!

Would-be winners

Axelrod also found that other strategies, had they been present, would have won.

One example of this is TITFORTWOTATS, a strategy that will cooperate unless the other player has defected twice in a row in which case it will defect (but just once).

The good performance of this strategy against the assembled field seems to indicate that the entries in the tournament were too keen on retaliating, where being slightly more forgiving might have been advantageous.

Would-be winners

Axelrod also found that other strategies, had they been present, would have won.

One example of this is TitForTwoTats, a strategy that will cooperate unless the other player has defected twice in a row in which case it will defect (but just once).

The good performance of this strategy against the assembled field seems to indicate that the entries in the tournament were too keen on retaliating, where being slightly more forgiving might have been advantageous.

Another example is a revised Downing which starts with cooperation rather than defection. We call the resulting strategy RDowning.

Would-be winners

Axelrod also found that other strategies, had they been present, would have won.

One example of this is TitForTwoTats, a strategy that will cooperate unless the other player has defected twice in a row in which case it will defect (but just once).

The good performance of this strategy against the assembled field seems to indicate that the entries in the tournament were too keen on retaliating, where being slightly more forgiving might have been advantageous.

Another example is a revised Downing which starts with cooperation rather than defection. We call the resulting strategy RDowning.

All in all many participants did not do very well when judging what strategies their programs might be up against.

A second computer tournament

A second tournament

Axelrod decided that the results of the first tournament where so promising that it would be worth his while to conduct a similar experiment again.

A second tournament

Axelrod decided that the results of the first tournament where so promising that it would be worth his while to conduct a similar experiment again.

This time he did not invite all the participants personally; apart from contacting all previous players he advertised in a number of computer magazines.

A second tournament

Axelrod decided that the results of the first tournament where so promising that it would be worth his while to conduct a similar experiment again.

This time he did not invite all the participants personally; apart from contacting all previous players he advertised in a number of computer magazines.

Everybody who was interested could obtain from him a description of the first tournament, together with Axelrod's analysis.

A second tournament

Axelrod decided that the results of the first tournament where so promising that it would be worth his while to conduct a similar experiment again.

This time he did not invite all the participants personally; apart from contacting all previous players he advertised in a number of computer magazines.

Everybody who was interested could obtain from him a description of the first tournament, together with Axelrod's analysis.

Question. What strategy would you have submitted this time?

A second tournament

Axelrod decided that the results of the first tournament where so promising that it would be worth his while to conduct a similar experiment again.

This time he did not invite all the participants personally; apart from contacting all previous players he advertised in a number of computer magazines.

Everybody who was interested could obtain from him a description of the first tournament, together with Axelrod's analysis.

Question. What strategy would you have submitted this time?

This time he attracted 62 entrants from six countries. Rapoport once again submitted TitForTat, while the British biologist John Maynard Smith put in TitForTwoTats. RDowning was present as well, and Axelrod entered Random to the field.

Lessons

Once again, the winner was TitForTat. Surprisingly, TitForTwoTats was only 24th, and RDowning came 38th and 40th (it was submitted by two participants). This shows the importance of the environment when it comes to determining what does well.

Lessons

Once again, the winner was TitForTat. Surprisingly, TitForTwoTats was only 24th, and RDowning came 38th and 40th (it was submitted by two participants). This shows the importance of the environment when it comes to determining what does well.

All but one of the top fifteen entries were nice. All but one of the bottom fifteen were not nice. Many people had learned this lesson, since almost two-thirds of all entries were nice.

Lessons

Once again, the winner was TitForTat. Surprisingly, TitForTwoTats was only 24th, and RDowning came 38th and 40th (it was submitted by two participants). This shows the importance of the environment when it comes to determining what does well.

All but one of the top fifteen entries were nice. All but one of the bottom fifteen were not nice.

Almost all the entries were forgiving, and therefore this criterion no longer provided a useful distinction.

Lessons

Once again, the winner was TitForTat. Surprisingly, TitForTwoTats was only 24th, and RDowning came 38th and 40th (it was submitted by two participants). This shows the importance of the environment when it comes to determining what does well.

All but one of the top fifteen entries were nice. All but one of the bottom fifteen were not nice.

Almost all the entries were forgiving, and therefore this criterion no longer provided a useful distinction.

A new criterion was being retaliatory—in order not to be exploited, a program has to hit back when the other sides defects. (However, bearing grudges is not a good idea.)

Lessons

Once again, the winner was TitForTat. Surprisingly, TitForTwoTats was only 24th, and RDowning came 38th and 40th (it was submitted by two participants). This shows the importance of the environment when it comes to determining what does well.

All but one of the top fifteen entries were nice. All but one of the bottom fifteen were not nice.

Almost all the entries were forgiving, and therefore this criterion no longer provided a useful distinction.

A new criterion was being retaliatory—in order not to be exploited, a program has to hit back when the other sides defects. (However, bearing grudges is not a good idea.)

TitForTat has all three properties identified as advantageous: It is nice (it never defects first), forgiving (it is prepared to try cooperation again if the other side shows willingness) and retaliatory (it will copy defections made by the other side in the next round).

Checking for weaknesses

Axelrod identified two strategies which worked as tests how retaliatory strategies were.

Checking for weaknesses

Axelrod identified two strategies which worked as tests how retaliatory strategies were.

He named the first one TESTER. It defects on the first move to see how the other side reacts. If the other strategy retaliates, it 'apologizes' by cooperating in the second game and playing as TITFORTAT from then on. If the other side does not respond promptly, it cooperates on the second and third move, but then defects every other round.

Checking for weaknesses

Axelrod identified two strategies which worked as tests how retaliatory strategies were.

He named the first one TESTER. It defects on the first move to see how the other side reacts. If the other strategy retaliates, it 'apologizes' by cooperating in the second game and playing as TITFORTAT from then on. If the other side does not respond promptly, it cooperates on the second and third move, but then defects every other round.

It did very well against strategies which forgave too easily, or needed more provocation to retaliate, such as TITFORTWOTATS. It also exploited strategies similar to DOWNING: Because TESTER cooperates just over half of the time, these strategies calculate that it pays off to keep cooperating—and are exploited for that. However there were not enough suckers for it to do well in the overall tournament: It only came 46th.

Checking for weaknesses

Axelrod identified two strategies which worked as tests how retaliatory strategies were.

He named the first one TESTER. It defects on the first move to see how the other side reacts. If the other strategy retaliates, it ‘apologizes’ by cooperating in the second game and playing as TITFORTAT from then on. If the other side does not respond promptly, it cooperates on the second and third move, but then defects every other round.

The other testing strategy Axelrod dubbed ‘Tranquilizer’, short TRANQ. It works by first establishing a relation of mutual cooperation (assuming the other player allows that), and then cautiously tests whether it can get away with some defections. It is careful, however, never to defect twice in a row, nor more than at a quarter of all moves. It also defects in reply to the other side defecting.

Checking for weaknesses

Axelrod identified two strategies which worked as tests how retaliatory strategies were.

He named the first one TESTER. It defects on the first move to see how the other side reacts. If the other strategy retaliates, it 'apologizes' by cooperating in the second game and playing as TITFORTAT from then on. If the other side does not respond promptly, it cooperates on the second and third move, but then defects every other round.

The other testing strategy Axelrod dubbed 'Tranquilizer', short TRANQ. It works by first establishing a relation of mutual cooperation (assuming the other player allows that), and then cautiously tests whether it can get away with some defections. It is careful, however, never to defect twice in a row, nor more than at a quarter of all moves. It also defects in reply to the other side defecting.

This fairly sophisticated program, based on reasonable sounding arguments, did poorly overall—it only came twenty-seventh (so it at least did considerably better than TESTER).

Checking for weaknesses

Axelrod identified two strategies which worked as tests how retaliatory strategies were.

He named the first one TESTER. It defects on the first move to see how the other side reacts. If the other strategy retaliates, it 'apologizes' by cooperating in the second game and playing as TITFORTAT from then on. If the other side does not respond promptly, it cooperates on the second and third move, but then defects every other round.

The other testing strategy Axelrod dubbed 'Tranquilizer', short TRANQ. It works by first establishing a relation of mutual cooperation (assuming the other player allows that), and then cautiously tests whether it can get away with some defections. It is careful, however, never to defect twice in a row, nor more than at a quarter of all moves. It also defects in reply to the other side defecting.

In order to do well against programs like TESTER and TRANQ, a program must retaliate as soon as the other side defects without provocation and so make it clear that will not tolerate this sort of behaviour.

Representative entries

In analysing the results of the second tournament further, Axelrod discovered that he could divide the strategies present into one of five categories.

Representative entries

In analysing the results of the second tournament further, Axelrod discovered that he could divide the strategies present into one of five categories.

There were 'typical' representatives of these categories present in the following sense. Let the representatives of the five categories be S_1, $S_2, \ldots, S_5$ and let $P(T, S_i)$ be the pay-off of the strategy T playing against the strategy S_i. Now consider the formula

$$120 + .202P(T, S_1) + .198P(T, S_2) + .11P(T, S_3) + .072P(T, S_4) + .086P(T, S_5).$$

Representative entries

In analysing the results of the second tournament further, Axelrod discovered that he could divide the strategies present into one of five categories.

There were 'typical' representatives of these categories present in the following sense. Let the representatives of the five categories be S_1, $S_2, \ldots, S_5$ and let $P(T, S_i)$ be the pay-off of the strategy T playing against the strategy S_i. Now consider the formula

$$120 + .202P(T, S_1) + .198P(T, S_2) + .11P(T, S_3) + .072P(T, S_4) + .086P(T, S_5).$$

This can be used to predict the score of a strategy T in the tournament to an amazing degree of accuracy (typically a difference of fewer than 10 points out of about three to four hundred). The strategy S_4 is TESTER while S_5 is TRANQ. The strategies S_1, S_2 and S_3 are nice.

Representative entries

There were ‘typical’ representatives of these categories present in the following sense. Let the representatives of the five categories be S_1, $S_2, \ldots, S_5$ and let $P(T, S_i)$ be the pay-off of the strategy T playing against the strategy S_i. Now consider the formula

$$120+.202P(T,S_1)+.198P(T,S_2)+.11P(T,S_3)+.072P(T,S_4)+.086P(T,S_5).$$

This can be used to predict the score of a strategy T in the tournament to an amazing degree of accuracy (typically a difference of fewer than 10 points out of about three to four hundred). The strategy S_4 is TESTER while S_5 is TRANQ. The strategies S_1, S_2 and S_3 are nice.

Among these, S_2, called REVISEDSTATETRANSITION by Axelrod, keeps a model of the other player as a one-step Markov process. Given the history of the game so far, it estimates the probability that the other player will defect based on the result of the previous round.

Hypothetical tournaments

This discovery gave Axelrod the idea to use the actually run tournament as a basis for hypothetical ones. By using weights, he could pretend that there was a certain percentage of various strategies present.

Hypothetical tournaments

This discovery gave Axelrod the idea to use the actually run tournament as a basis for hypothetical ones. By using weights, he could pretend that there was a certain percentage of various strategies present.

He found that TITFORTAT is robust, that is it does well in a large variety of environments. It is friendly enough to elicit cooperation, but cannot be exploited since it retaliates upon defection. It pays to cooperate with TITFORTAT.

Hypothetical tournaments

This discovery gave Axelrod the idea to use the actually run tournament as a basis for hypothetical ones. By using weights, he could pretend that there was a certain percentage of various strategies present.

He found that TITFORTAT is robust, that is it does well in a large variety of environments. It is friendly enough to elicit cooperation, but cannot be exploited since it retaliates upon defection. It pays to cooperate with TITFORTAT.

It never outscores a playing partner, so in order to do well in non zero-sum games one does not have to be better than the others on an individual basis!

Hypothetical tournaments

This discovery gave Axelrod the idea to use the actually run tournament as a basis for hypothetical ones. By using weights, he could pretend that there was a certain percentage of various strategies present.

He found that TITFORTAT is robust, that is it does well in a large variety of environments. It is friendly enough to elicit cooperation, but cannot be exploited since it retaliates upon defection. It pays to cooperate with TITFORTAT.

It never outscores a playing partner, so in order to do well in non zero-sum games one does not have to be better than the others on an individual basis!

In any of Axelrod's tournaments, ALWAYSD would have done very poorly, with a pay-off of 200 against most other strategies. Some of the strategies that finished towards the end of the field were not clearly enough distinguishable from RANDOM, so their partners gave up on trying to cooperate with them.

Further developments

A tournament organized by D.R Hofstadter and Marek Lugowski had AlwaysD come almost last, but variations of TitForTat did better than the original. They switched to AlwaysD if the other player was too uncooperative.

Further developments

A tournament organized by D.R Hofstadter and Marek Lugowski had ALWAYSD come almost last, but variations of TITFORTAT did better than the original. They switched to ALWAYSD if the other player was too uncooperative.

What is successful depends on the environment. But surely one would expect such an environment to change, as unsuccessful strategies vanish and the better ones multiply. This is the idea behind evolutionary game theory.

Further developments

In 2004 a competition with 223 entries celebrated the 20th anniversary of Axelrod's first tournament. A team from the University of Southampton submitted 60 programs which were designed to recognize each other. Then one of them would turn itself into point fodder for the other. If meeting some other program these would turn into ALWAYSD.

Further developments

In 2004 a competition with 223 entries celebrated the 20th anniversary of Axelrod's first tournament. A team from the University of Southampton submitted 60 programs which were designed to recognize each other. Then one of them would turn itself into point fodder for the other. If meeting some other program these would turn into ALWAYSD.

The top three programs were theirs, but so were a lot of the bottom ones. The Southampton team did some further research and decided that 20 of their programs would have led to a similar result.

Infinitely and indefinitely repeated versions

Variations on iterated games–infinite

In order to make an iterated PD game more realistic it makes sense to not play a fixed number of rounds. If we do not assume this then the ‘backwards induction argument’ we talked about earlier does not apply.

Variations on iterated games–infinite

In order to make an iterated PD game more realistic it makes sense to not play a fixed number of rounds. If we do not assume this then the ‘backwards induction argument’ we talked about earlier does not apply.

One alternative is to try to go for infinitely many repetitions. There are two main problems with that.

Variations on iterated games–infinite

In order to make an iterated PD game more realistic it makes sense to not play a fixed number of rounds. If we do not assume this then the ‘backwards induction argument’ we talked about earlier does not apply.

One alternative is to try to go for infinitely many repetitions. There are two main problems with that.

- It is not particularly realistic, since nobody has an infinite amount of time.

Variations on iterated games–infinite

In order to make an iterated PD game more realistic it makes sense to not play a fixed number of rounds. If we do not assume this then the 'backwards induction argument' we talked about earlier does not apply.

One alternative is to try to go for infinitely many repetitions. There are two main problems with that.

- It is not particularly realistic, since nobody has an infinite amount of time.
- Calculating expected outcomes is hard—the only sensible way of doing so is to
  - insist that all strategies have a finite definition;
  - calculate the cycle that two strategies playing against each other will eventually run into;
  - calculate the pay-off for the strategies over the cycle and thus an average pay-off per round.

Variations on iterated games–indefi nit

A more popular model is one where there is a random decision as to whether another round of the game will be played.

Variations on iterated games–indefi nit

A more popular model is one where there is a random decision as to whether another round of the game will be played.

Typically the assumption is that there is a random parameter w, and that after each round a chance experiment is conducted so that with probability w there will be another round. This is known as an indefinitely repeated game.

Variations on iterated games–indefinit

A more popular model is one where there is a random decision as to whether another round of the game will be played.

Typically the assumption is that there is a random parameter w, and that after each round a chance experiment is conducted so that with probability w there will be another round. This is known as an indefinitely repeated game.

Then the AlwaysC strategy playing against itself has an expected pay-off of

$$R + wR + w^2R + w^3R + \cdots = \frac{R}{1-w}.$$

Variations on iterated games–indefi nit

Typically the assumption is that there is a random parameter w, and that after each round a chance experiment is conducted so that with probability w there will be another round. This is known as an indefinitely repeated game.

Then the AlwaysC strategy playing against itself has an expected pay-off of

$$R + wR + w^2R + w^3R + \cdots = \frac{R}{1-w}.$$

The closer w is to 1 the higher the probability that there will be another game, and the higher the difference between the expected pay-off from mutual defection *versus* that from mutual cooperation. People sometimes call this 'the shadow of the future' to capture the idea that current behaviour is likely to influence future gains.

Variations on iterated games–indefinit

Typically the assumption is that there is a random parameter w, and that after each round a chance experiment is conducted so that with probability w there will be another round. This is known as an indefinitely repeated game.

Then the AlwaysC strategy playing against itself has an expected pay-off of

$$R + wR + w^2R + w^3R + \cdots = \frac{R}{1-w}.$$

The closer w is to 1 the higher the probability that there will be another game, and the higher the difference between the expected pay-off from mutual defection *versus* that from mutual cooperation.

There is another way of reading this, which seems to be preferred by economists: Rather than read w as a probability that there will be another game they tend to view it as a discount factor, the idea being that a pay-off in the future is worth less than one now.

The indefinitely repeated game

The indefinitely repeated game is quite complex!

The indefinitely repeated game

The indefinitely repeated game is quite complex!

Proposition 5.2 *For each of the following properties there is at least one sub-game equilibrium point for the indefinitely repeated Prisoner's Dilemma game satisfying it. In the play resulting from employing these strategies against each other it is the case that*

- *they both cooperate throughout;*
- *they both defect throughout;*
- *they both defect in the first round and cooperate ever thereafter;*
- *they both cooperate in every odd round and defect in every even one;*
- *Player 1 cooperates in even rounds and defects in odd rounds while Player 2 defects in even rounds and cooperates in odd ones.*

The indefinitely repeated game

Proposition 5.2 *For each of the following properties there is at least one sub-game equilibrium point for the indefinitely repeated Prisoner's Dilemma game satisfying it. In the play resulting from employing these strategies against each other it is the case that*

- *they both cooperate throughout;*
- *they both defect throughout;*
- *they both defect in the first round and cooperate ever thereafter;*
- *they both cooperate in every odd round and defect in every even one;*
- *Player 1 cooperates in even rounds and defects in odd rounds while Player 2 defects in even rounds and cooperates in odd ones.*

This shows that indefinitely repeated Prisoner's Dilemma has a huge number of sub-game equilibrium points, which means that it is very hard to predict just what will happen in this game. Also note that TITFORTAT is not a sub-game equilibrium point, but it is an equilibrium point in this game!

Summary of Section 5

- The generalized Prisoner's Dilemma has four pay-offs, $T > R > P > S$ with $R > (S + T)/2$. Its only equilibrium point is given by the strategy which always defects (for both players), and this is also the dominant strategy.

Summary of Section 5

- The generalized Prisoner's Dilemma has four pay-offs, $T > R > P > S$ with $R > (S+T)/2$. Its only equilibrium point is given by the strategy which always defects (for both players).
- If the game is repeated a fixed number of times, then the strategy which always defects still gives an equilibrium point, but no longer the only one. It is, however, the only sub-game equilibrium point.

Summary of Section 5

- The generalized Prisoner's Dilemma has four pay-offs, $T > R > P > S$ with $R > (S+T)/2$. Its only equilibrium point is given by the strategy which always defects (for both players).
- If the game is repeated a fixed number of times, then the strategy which always defects still gives an equilibrium point, but no longer the only one. It is, however, the only sub-game equilibrium point, that is it gives an equilibrium point in every sub-game.
- In two computer tournaments playing this game, the strategy known as TitForTat won twice. Successful strategies were nice, forgiving and retaliatory.

Summary of Section 5

- The generalized Prisoner's Dilemma has four pay-offs, $T > R > P > S$ with $R > (S+T)/2$. Its only equilibrium point is given by the strategy which always defects (for both players).
- If the game is repeated a fixed number of times, then the strategy which always defects still gives an equilibrium point, but no longer the only one. It is, however, the only sub-game equilibrium point.
- In two computer tournaments playing this game, the strategy known as TitForTat won twice. Successful strategies were nice, forgiving and retaliatory.
- More realistic versions of this game are indefinitely repeated, where after each round a chance experiments determines whether a further round will be played. This will occur with a fixed probability.

Summary of Section 5

- The generalized Prisoner's Dilemma has four pay-offs, $T > R > P > S$ with $R > (S + T)/2$. Its only equilibrium point is given by the strategy which always defects (for both players).
- If the game is repeated a fixed number of times, then the strategy which always defects still gives an equilibrium point, but no longer the only one. It is, however, the only sub-game equilibrium point.
- In two computer tournaments playing this game, the strategy known as TitForTat won twice. Successful strategies were nice, forgiving and retaliatory.
- More realistic versions of this game are indefinitely repeated, where after each round a chance experiments determines whether a further round will be played. This will occur with a fixed probability.
- Prisoner's Dilemma situation can be readily identified in the real world.