CS3191 Section 6

Games and evolution

Andrea Schalk

School of Computer Science, University of Manchester

Games and evolution

Section 6 investigates the use of games in modelling evolution.

Games and evolution

Section 6 investigates the use of games in modelling evolution.

The game we will concentrate on is still the Prisoner's Dilemma game.

Games and evolution

Section 6 investigates the use of games in modelling evolution.

The game we will concentrate on is still the Prisoner's Dilemma game.

We look at tournaments which are played over generations, and where the population of a new generation depends on what was successful in the previous round.

Games and evolution

Section 6 investigates the use of games in modelling evolution.

The game we will concentrate on is still the Prisoner's Dilemma game.

We look at tournaments which are played over generations, and where the population of a new generation depends on what was successful in the previous round.

We study which strategies can invade others and which are safe from invasion in different kinds of systems.

Games and evolution

Section 6 investigates the use of games in modelling evolution.

The game we will concentrate on is still the Prisoner's Dilemma game.

We look at tournaments which are played over generations, and where the population of a new generation depends on what was successful in the previous round.

We study which strategies can invade others and which are safe from invasion in different kinds of systems.

We consider different ways of presenting the strategies present in the system.

Games and evolution

Section 6 investigates the use of games in modelling evolution.

The game we will concentrate on is still the Prisoner's Dilemma game.

We look at tournaments which are played over generations, and where the population of a new generation depends on what was successful in the previous round.

We study which strategies can invade others and which are safe from invasion in different kinds of systems.

We consider different ways of presenting the strategies present in the system.

Lastly we look at further biological games.

An ecological tournament

Indefinitely repeated PD

The space of strategies for the indefinitely repeated Prisoner's Dilemma game is huge, and so far it has not been fully explored.

Indefinitely repeated PD

The space of strategies for the indefinitely repeated Prisoner's Dilemma game is huge, and so far it has not been fully explored.

Even if one were to limit oneself merely to strategies which are nice, forgiving and retaliatory there would be too many of them to take into account. Just how successful are these strategies?

Indefinitely repeated PD

The space of strategies for the indefinitely repeated Prisoner's Dilemma game is huge, and so far it has not been fully explored.

Even if one were to limit oneself merely to strategies which are nice, forgiving and retaliatory there would be too many of them to take into account. Just how successful are these strategies?

What one can show is that if w, the probability of playing another round, is big enough then it pays to be cooperative. Under these circumstances, TitForTat paired with itself is an equilibrium point, but not a sub-game equilibrium point.

Indefinitely repeated PD

The space of strategies for the indefinitely repeated Prisoner's Dilemma game is huge, and so far it has not been fully explored.

Even if one were to limit oneself merely to strategies which are nice, forgiving and retaliatory there would be too many of them to take into account. Just how successful are these strategies?

What one can show is that if w, the probability of playing another round, is big enough then it pays to be cooperative. Under these circumstances, TITFORTAT paired with itself is an equilibrium point, but not a sub-game equilibrium point.

On the downside, ALWAYSD paired with itself also is an equilibrium point.

Indefinitely repeated PD

The space of strategies for the indefinitely repeated Prisoner's Dilemma game is huge, and so far it has not been fully explored.

Even if one were to limit oneself merely to strategies which are nice, forgiving and retaliatory there would be too many of them to take into account. Just how successful are these strategies?

What one can show is that if w, the probability of playing another round, is big enough then it pays to be cooperative. Under these circumstances, TitForTat paired with itself is an equilibrium point, but not a sub-game equilibrium point.

On the downside, AlwaysD paired with itself also is an equilibrium point.

We will see how much we can say about this situation.

Evolution

The idea of simulating evolution using games, more specifically, the indefinitely repeated PD game, is simple:

Evolution

The idea of simulating evolution using games, more specifically, the indefinitely repeated PD game, is simple:

Assume that there are a number of individuals playing the indefinitely repeated PD game against each other, each employing some strategy. After a predefined number of encounters, they have accumulated a number of points each.

Evolution

The idea of simulating evolution using games, more specifically, the indefinitely repeated PD game, is simple:

Assume that there are a number of individuals playing the indefinitely repeated PD game against each other, each employing some strategy. After a predefined number of encounters, they have accumulated a number of points each.

In the next generation, the proportion of each strategy ought to depend on its success in the previous one. In other words, there should be more individuals employing successful strategies and fewer individuals employing unsuccessful ones.

Evolution

The idea of simulating evolution using games, more specifically, the indefinitely repeated PD game, is simple:

Assume that there are a number of individuals playing the indefinitely repeated PD game against each other, each employing some strategy. After a predefined number of encounters, they have accumulated a number of points each.

In the next generation, the proportion of each strategy ought to depend on its success in the previous one. In other words, there should be more individuals employing successful strategies and fewer individuals employing unsuccessful ones.

It does not matter much for our model how this works.

Evolution

The idea of simulating evolution using games, more specifically, the indefinitely repeated PD game, is simple:

Assume that there are a number of individuals playing the indefinitely repeated PD game against each other, each employing some strategy. After a predefined number of encounters, they have accumulated a number of points each.

In the next generation, the proportion of each strategy ought to depend on its success in the previous one. In other words, there should be more individuals employing successful strategies and fewer individuals employing unsuccessful ones.

It does not matter much for our model how this works.

- Individuals might learn from mistakes.

Evolution

The idea of simulating evolution using games, more specifically, the indefinitely repeated PD game, is simple:

Assume that there are a number of individuals playing the indefinitely repeated PD game against each other, each employing some strategy. After a predefined number of encounters, they have accumulated a number of points each.

In the next generation, the proportion of each strategy ought to depend on its success in the previous one. In other words, there should be more individuals employing successful strategies and fewer individuals employing unsuccessful ones.

It does not matter much for our model how this works.

- Individuals might learn from mistakes.
- Individuals might imitate more successful others.

Evolution

The idea of simulating evolution using games, more specifically, the indefinitely repeated PD game, is simple:

Assume that there are a number of individuals playing the indefinitely repeated PD game against each other, each employing some strategy. After a predefined number of encounters, they have accumulated a number of points each.

In the next generation, the proportion of each strategy ought to depend on its success in the previous one. In other words, there should be more individuals employing successful strategies and fewer individuals employing unsuccessful ones.

It does not matter much for our model how this works.

- Individuals might learn from mistakes.
- Individuals might imitate more successful others.
- Successful individuals might have more offspring than those who are struggling.

Local or global?

There are two different models which are commonly studied when it comes to modelling evolution.

Local or global?

There are two different models which are commonly studied when it comes to modelling evolution.

We can look at a global system, where every individual plays against every other individual, and every individual can learn from every other individual.

Local or global?

There are two different models which are commonly studied when it comes to modelling evolution.

We can look at a global system, where every individual plays against every other individual, and every individual can learn from every other individual.

Alternatively we can look at a local system, where we have a notion of individuals being neighbours. In such a system individuals play against their neighbours and learn from their neighbours. These systems are also known as territorial because of their sense of location.

Local or global?

There are two different models which are commonly studied when it comes to modelling evolution.

We can look at a global system, where every individual plays against every other individual, and every individual can learn from every other individual.

Alternatively we can look at a local system, where we have a notion of individuals being neighbours. In such a system individuals play against their neighbours and learn from their neighbours. These systems are also known as territorial because of their sense of location.

There are other variations but they are not usually studied.

Local or global?

There are two different models which are commonly studied when it comes to modelling evolution.

We can look at a global system, where every individual plays against every other individual, and every individual can learn from every other individual.

Alternatively we can look at a local system, where we have a notion of individuals being neighbours. In such a system individuals play against their neighbours and learn from their neighbours. These systems are also known as territorial because of their sense of location.

There are other variations but they are not usually studied.

We first look at global systems before transferring our insights to local ones.

An evolutionary tournament

Armed with the result from his second tournament Axelrod ran an evolutionary tournament along the lines just described.

An evolutionary tournament

Armed with the result from his second tournament Axelrod ran an evolutionary tournament along the lines just described.

He observed that by generation 5 all the lowest ranking 11 entries were down to half their initial number; mid-ranking ones roughly held their owns and the top-ranking ones started to grow in number.

An evolutionary tournament

Armed with the result from his second tournament Axelrod ran an evolutionary tournament along the lines just described.

He observed that by generation 5 all the lowest ranking 11 entries were down to half their initial number; mid-ranking ones roughly held their owns and the top-ranking ones started to grow in number.

By generation 50 the bottom third of strategies had all but disappeared, most of the middle third had started to shrink in number, and those in the top third continued to grow.

An evolutionary tournament

Armed with the result from his second tournament Axelrod ran an evolutionary tournament along the lines just described.

He observed that by generation 5 all the lowest ranking 11 entries were down to half their initial number; mid-ranking ones roughly held their owns and the top-ranking ones started to grow in number.

By generation 50 the bottom third of strategies had all but disappeared, most of the middle third had started to shrink in number, and those in the top third continued to grow.

Note that to survive in the long run, a strategy has to do well against other successful strategies, since the others are going to die out.

An evolutionary tournament

Armed with the result from his second tournament Axelrod ran an evolutionary tournament along the lines just described.

He observed that by generation 5 all the lowest ranking 11 entries were down to half their initial number; mid-ranking ones roughly held their owns and the top-ranking ones started to grow in number.

By generation 50 the bottom third of strategies had all but disappeared, most of the middle third had started to shrink in number, and those in the top third continued to grow.

Note that to survive in the long run, a strategy has to do well against other successful strategies, since the others are going to die out.

So strategies that live on exploiting non-retaliatory others will do well initially but run out of victims eventually and then vanish. They are the victims of their own success!

An evolutionary tournament

Armed with the result from his second tournament Axelrod ran an evolutionary tournament along the lines just described.

He observed that by generation 5 all the lowest ranking 11 entries were down to half their initial number; mid-ranking ones roughly held their owns and the top-ranking ones started to grow in number.

By generation 50 the bottom third of strategies had all but disappeared, most of the middle third had started to shrink in number, and those in the top third continued to grow.

Note that to survive in the long run, a strategy has to do well against other successful strategies, since the others are going to die out.

So strategies that live on exploiting non-retaliatory others will do well initially but run out of victims eventually and then vanish. They are the victims of their own success!

This tournament was won by TitForTat, which ended up with a 15% share of the entire population, fifteen times the number it started with.

Invaders and collective stability

No simple answers

Axelrod became intrigued with this situation. How could cooperation ever develop in a world where it seems advantageous to only have one's own well-being in mind?

No simple answers

Axelrod became intrigued with this situation. How could cooperation ever develop in a world where it seems advantageous to only have one's own well-being in mind?

Like he, we will from now on consider the indefinitely repeated PD with probability w.

No simple answers

Axelrod became intrigued with this situation. How could cooperation ever develop in a world where it seems advantageous to only have one's own well-being in mind?

Like he, we will from now on consider the indefinitely repeated PD with probability w.

First of all we note that there are no simple answers in this game.

No simple answers

Axelrod became intrigued with this situation. How could cooperation ever develop in a world where it seems advantageous to only have one's own well-being in mind?

Like he, we will from now on consider the indefinitely repeated PD with probability w.

First of all we note that there are no simple answers in this game.

Proposition 6.1 *In an indefinitely repeated game of Prisoner's Dilemma with two players there is no one best strategy if* w *is large enough.*

No simple answers

Axelrod became intrigued with this situation. How could cooperation ever develop in a world where it seems advantageous to only have one's own well-being in mind?

Like he, we will from now on consider the indefinitely repeated PD with probability w.

First of all we note that there are no simple answers in this game.

Proposition 6.1 *In an indefinitely repeated game of Prisoner's Dilemma with two players there is no one best strategy if w is large enough.*

In other words whether a strategy does well depends on which other strategies are present!

Proof

Assume there was a best strategy.

Proof

Assume there was a best strategy.

We first assume that that strategy cooperates on the first move.

Proof

Assume there was a best strategy.

We first assume that that strategy cooperates on the first move.

If this strategy is played against the AlwaysD strategy, then the best it can do after cooperating on the first move is to defect forever, choosing the greater pay-off P over S.

Proof

Assume there was a best strategy.

We first assume that that strategy cooperates on the first move.

If this strategy is played against the AlwaysD strategy, then the best it can do after cooperating on the first move is to defect forever, choosing the greater pay-off P over S.

Its pay-off when playing against AlwaysD is
$S + wP + w^2P + \cdots = S + \frac{wP}{1-w}$.

Proof

Assume there was a best strategy.

We first assume that that strategy cooperates on the first move.

If this strategy is played against the AlwaysD strategy, then the best it can do after cooperating on the first move is to defect forever, choosing the greater pay-off P over S.

Its pay-off when playing against AlwaysD is $S + wP + w^2P + \cdots = S + \frac{wP}{1-w}$.

AlwaysD playing against AlwaysD gets a pay-off of

$$P + wP + w^2P + \cdots = P + \frac{wP}{1-w},$$

which is higher. Since we have found a better strategy for this situation, the one we started with can't have been the best overall.

Proof

Assume there was a best strategy.

Now assume that our best strategy defects on the first move.

Proof

Assume there was a best strategy.

Now assume that our best strategy defects on the first move. If it plays the GRUDGE strategy it will get a pay-off of T for the first move, but forever thereafter, GRUDGE will defect.

Proof

Assume there was a best strategy.

Now assume that our best strategy defects on the first move. If it plays the GRUDGE strategy it will get a pay-off of T for the first move, but forever thereafter, GRUDGE will defect. Hence our strategy will get pay-off at most $T + wP + w^2P + \cdots = T + \frac{wP}{1-w}$.

Proof

Assume there was a best strategy.

Now assume that our best strategy defects on the first move. If it plays the GRUDGE strategy it will get a pay-off of T for the first move, but forever thereafter, GRUDGE will defect. Hence our strategy will get pay-off at most $T + wP + w^2P + \cdots = T + \frac{wP}{1-w}$. The ALWAYSC strategy, on the other hand, when playing against GRUDGE can expect a pay-off of

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

Proof

Assume there was a best strategy.

Now assume that our best strategy defects on the first move. If it plays the GRUDGE strategy it will get a pay-off of T for the first move, but forever thereafter, GRUDGE will defect. Hence our strategy will get pay-off at most $T + wP + w^2P + \cdots = T + \frac{wP}{1-w}$. The ALWAYSC strategy, on the other hand, when playing against GRUDGE can expect a pay-off of

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

ALWAYSC is better off than our strategy if

$$R + \frac{wR}{1-w} > T + \frac{wP}{1-w},$$ that is when $$R > T(1-w) + wP$$

which is the case if and only if

$$w > \frac{T-R}{T-P}.$$

Proof

Assume there was a best strategy.

Now assume that our best strategy defects on the first move. ALWAYSC is better off than our strategy if

$$R > T(1-w) + wP.$$

Now the latter is bigger than the former provided that $w > \frac{T-R}{T-P}$.

Hence if w is larger than this threshold value then we have once again found a strategy (ALWAYSC) which performs better (against GRUDGE) than our best strategy. Therefore such a best strategy cannot exist.

Collective stability

A sensible question is whether any one strategy is vulnerable to being invaded by another strategy.

Collective stability

A sensible question is whether any one strategy is vulnerable to being invaded by another strategy.

In a situation where A is the resident strategy it is vulnerable against an invader B if B does better against A than A does against itself.

Collective stability

A sensible question is whether any one strategy is vulnerable to being invaded by another strategy.

In a situation where A is the resident strategy it is vulnerable against an invader B if B does better against A than A does against itself.

Because if it does so then in the next generation there will be more copies of B, and they will keep multiplying.

Collective stability

A sensible question is whether any one strategy is vulnerable to being invaded by another strategy.

In a situation where A is the resident strategy it is vulnerable against an invader B if B does better against A than A does against itself.

Because if it does so then in the next generation there will be more copies of B, and they will keep multiplying.

Definition 12 *Let $P(A, B)$ be the pay-off that a strategy A receives when playing indefinitely repeated Prisoner's Dilemma against a strategy B. We say that a strategy B can* **invade** *a native strategy A if it gets a higher pay-off when playing against A then A gets when playing against itself, that is if*

$$P(B, A) > P(A, A).$$

We say that a strategy is **collectively stable** *if it cannot be invaded by any other strategy.*

Collective stability

A sensible question is whether any one strategy is vulnerable to being invaded by another strategy.

Definition 12 *Let* $P(A, B)$ *be the pay-off that a strategy* A *receives when playing indefinitely repeated Prisoner's Dilemma against a strategy* B*. We say that a strategy* B *can* **invade** *a native strategy* A *if it gets a higher pay-off when playing against* A *then* A *gets when playing against itself, that is if*

$$P(B, A) > P(A, A).$$

We say that a strategy is **collectively stable** *if it cannot be invaded by any other strategy.*

Note that in reality, an invader should be able to survive if it does just as well against the resident strategy as that strategy does against itself. This is known as drift.

Collective stability

A sensible question is whether any one strategy is vulnerable to being invaded by another strategy.

Definition 12 *Let $P(A, B)$ be the pay-off that a strategy A receives when playing indefinitely repeated Prisoner's Dilemma against a strategy B. We say that a strategy B can* **invade** *a native strategy A if it gets a higher pay-off when playing against A then A gets when playing against itself, that is if*

$$P(B, A) > P(A, A).$$

We say that a strategy is **collectively stable** *if it cannot be invaded by any other strategy.*

Collectively stable strategies can maintain themselves as a population since the can ward off all invaders.

A collectively stable strategy

Not surprisingly, the strategy which always defects is invulnerable to invasions.

A collectively stable strategy

Not surprisingly, the strategy which always defects is invulnerable to invasions.

Proposition 6.2 *The* ALWAYSD *strategy is collectively stable for all* w.

A collectively stable strategy

Not surprisingly, the strategy which always defects is invulnerable to invasions.

Proposition 6.2 *The* ALWAYSD *strategy is collectively stable for all* w*.*

The proof of this statement is an exercise.

What about TitForTat?

So what about TitForTat? How vulnerable is it against invasions?

What about TitForTat?

So what about TitForTat? How vulnerable is it against invasions?

In order to outperform the TitForTat strategy playing against itself, an invader has to defect at some stage.

What about TitForTat?

So what about TitForTat? How vulnerable is it against invasions?

In order to outperform the TitForTat strategy playing against itself, an invader has to defect at some stage.

For that defection, the invader will get pay-off T while TitForTat has to make do with pay-off S.

What about TITFORTAT?

So what about TITFORTAT? How vulnerable is it against invasions?

In order to outperform the TITFORTAT strategy playing against itself, an invader has to defect at some stage.

For that defection, the invader will get pay-off T while TITFORTAT has to make do with pay-off S. On the next move, TITFORTAT will defect itself, so the invader may get pay-off P or S in that round.

What about TITFORTAT?

So what about TITFORTAT? How vulnerable is it against invasions?

In order to outperform the TITFORTAT strategy playing against itself, an invader has to defect at some stage.

For that defection, the invader will get pay-off T while TITFORTAT has to make do with pay-off S. On the next move, TITFORTAT will defect itself, so the invader may get pay-off P or S in that round. In the latter case, they are level again, although this next round only occurs with probability w.

What about TitForTat?

So what about TitForTat? How vulnerable is it against invasions?

In order to outperform the TitForTat strategy playing against itself, an invader has to defect at some stage.

For that defection, the invader will get pay-off T while TitForTat has to make do with pay-off S. On the next move, TitForTat will defect itself, so the invader may get pay-off P or S in that round. In the latter case, they are level again, although this next round only occurs with probability w.

It will therefore have to depend on the probability w whether TitForTat is safe from invasions.

TitForTat *versus* AlwaysD

How TitForTat does fare against an invasion by AlwaysD.

TitForTat *versus* AlwaysD

How TitForTat does fare against an invasion by AlwaysD.

The play of these two against each other will consist of

- on the first move:

TitForTat *versus* AlwaysD

How TitForTat does fare against an invasion by AlwaysD.

The play of these two against each other will consist of

- on the first move: cooperation by TitForTat and defection by AlwaysD;
- on each subsequent move:

TitForTat *versus* AlwaysD

How TitForTat does fare against an invasion by AlwaysD.

The play of these two against each other will consist of

- on the first move: cooperation by TitForTat and defection by AlwaysD;
- on each subsequent move: defection by both sides.

TitForTat *versus* AlwaysD

How TitForTat does fare against an invasion by AlwaysD.

The play of these two against each other will consist of

- on the first move: cooperation by TitForTat and defection by AlwaysD;
- on each subsequent move: defection by both sides.

The pay-off for AlwaysD against TitForTat:

$$T + \frac{w}{1-w}P.$$

TitForTat *versus* AlwaysD

How TitForTat does fare against an invasion by AlwaysD.

The play of these two against each other will consist of

- on the first move: cooperation by TitForTat and defection by AlwaysD;
- on each subsequent move: defection by both sides.

The pay-off for AlwaysD against TitForTat:

$$T + \frac{w}{1-w}P.$$

The pay-off for TitForTat against TitForTat:

$$\frac{R}{1-w}.$$

TitForTat *versus* AlwaysD

How TitForTat does fare against an invasion by AlwaysD.

The play of these two against each other will consist of

- on the first move: cooperation by TitForTat and defection by AlwaysD;
- on each subsequent move: defection by both sides.

The pay-off for AlwaysD against TitForTat:

$$T + \frac{w}{1-w}P.$$

The pay-off for TitForTat against TitForTat:

$$\frac{R}{1-w}.$$

TitForTat can fight of the invasion iff

$$\frac{R}{1-w} \geq T + \frac{w}{1-w}P.$$

TitForTat *versus* AlwaysD

How TitForTat does fare against an invasion by AlwaysD.

The play of these two against each other will consist of

- on the first move: cooperation by TitForTat and defection by AlwaysD;
- on each subsequent move: defection by both sides.

TitForTat can fight of the invasion iff

$$\frac{R}{1-w} \geq T + \frac{w}{1-w}P.$$

This is true if and only if $R \geq T - wT + wP$, which is equivalent to

$$w \geq \frac{T-R}{T-P}.$$

TitForTat *versus* an arbitrary invader

If the probability w is large enough then TitForTat can fight off all invasions.

TitForTat *versus* an arbitrary invader

If the probability w is large enough then TitForTat can fight off all invasions.

Proposition 6.3 *The strategy* TitForTat *is collectively stable provided that the parameter* w *is greater than or equal to the maximum of* $(T-R)/(T-P)$ *and* $(T-R)/(R-S)$.

TitForTat *versus* an arbitrary invader

If the probability w is large enough then TitForTat can fight off all invasions.

Proposition 6.3 *The strategy* TitForTat *is collectively stable provided that the parameter* w *is greater than or equal to the maximum of* $(T - R)/(T - P)$ *and* $(T - R)/(R - S)$.

So at least a population of TitForTat strategies can be stable! The only requirement of this is that the probability that two individuals 'meet again' is sufficiently high. Niceness therefore doesn't have to be an impediment!

Proof

Any strategy that successfully invades a population of TitForTat has to defect at some point, say on move n. After that

Proof

Any strategy that successfully invades a population of TitForTat has to defect at some point, say on move n. After that

- The strategy may defect forever.
- The strategy may defect $k \geq 0$ times thereafter, and then cooperate.

Proof

Any strategy that successfully invades a population of TitForTat has to defect at some point, say on move n. After that

- The strategy may defect forever. But in that case we may discard the first $(n-1)$ rounds in which it gets the same pay-off as TitForTat and treat it as the AlwaysD strategy which we have shown cannot invade provided that $w \geq (T-R)/(T-P)$.
- The strategy may defect $k \geq 0$ times thereafter, and then cooperate.

Proof

Any strategy that successfully invades a population of TitForTat has to defect at some point, say on move n. After that

- The strategy may defect forever. ✓
- The strategy may defect $k \geq 0$ times thereafter, and then cooperate.

  Given the definition of TitForTat that means that from round n to round $n+k+1$, the strategy will accumulate a pay-off of

$$w^n(T + wP + w^2P + \cdots + w^kP + w^{k+1}S),$$

  and thereafter it is in the same situation as before (that is, TitForTat will cooperate on the next move, and the cycle repeats).

Proof

Any strategy that successfully invades a population of TitForTat has to defect at some point, say on move n. After that

- The strategy may defect forever. ✓
- The strategy may defect $k \geq 0$ times thereafter, and then cooperate.

  Given the definition of TitForTat that means that from round n to round $n+k+1$, the strategy will accumulate a pay-off of

  $$w^n(T + wP + w^2P + \cdots + w^kP + w^{k+1}S),$$

  and thereafter it is in the same situation as before . Now TitForTat's pay-off when playing against itself over these $n+k+1$ rounds is

  $$w^n(R + wR + w^2R + \cdots w^kR + w^{k+1}R).$$

Proof

Any strategy that successfully invades a population of TitForTat has to defect at some point, say on move n. After that

- The strategy may defect forever. ✓
- The strategy may defect $k \geq 0$ times thereafter, and then cooperate.

We want to see when TitForTat gets the higher pay-off, that is when

$$T + wP + w^2P + \cdots + w^kP + w^{k+1}S$$
$$\leq \quad R + wR + w^2R + \cdots w^kR + w^{k+1}R.$$

Proof

Any strategy that successfully invades a population of TitForTat has to defect at some point, say on move n. After that

- The strategy may defect forever. ✓
- The strategy may defect $k \geq 0$ times thereafter, and then cooperate.

We want to see when TitForTat gets the higher pay-off, that is when

$$\begin{aligned} & T + wP + w^2P + \cdots + w^kP + w^{k+1}S \\ \leq \quad & R + wR + w^2R + \cdots w^kR + w^{k+1}R. \end{aligned}$$

This is equivalent to

$$T - R \leq w(R-P) + \cdots + w^k(R-P) + w^{k+1}(R-S).$$

Proof ctd.

We want to show

$$T - R \leq w(R-P) + \cdots + w^k(R-P) + w^{k+1}(R-S).$$

Proof ctd.

We want to show

$$T - R \leq w(R - P) + \cdots + w^k(R - P) + w^{k+1}(R - S).$$

We proceed by induction over k.

Proof ctd.

We want to show

$$T - R \leq w(R - P) + \cdots + w^k(R - P) + w^{k+1}(R - S).$$

We proceed by induction over k.

- Base case: $k = 0$.

Proof ctd.

We want to show

$$T - R \le w(R-P) + \cdots + w^k(R-P) + w^{k+1}(R-S).$$

We proceed by induction over k.

- Base case: $k = 0$. But in that case the above inequality is

$$T - R \le w(R-S)$$

which is true by the condition on w.

Proof ctd.

We want to show

$$T - R \leq w(R - P) + \cdots + w^k(R - P) + w^{k+1}(R - S).$$

We proceed by induction over k.

- Base case: $k = 0$. ✓
- Induction step:

Proof ctd.

We want to show

$$T - R \leq w(R - P) + \cdots + w^k(R - P) + w^{k+1}(R - S).$$

We proceed by induction over k.

- Base case: $k = 0$. ✓
- Induction step: By the condition on w we know that $T - R \leq w(T - P) = w(T - R + R - P) = w(R - P) + w(T - R)$.

Proof ctd.

We want to show

$$T - R \leq w(R-P) + \cdots + w^k(R-P) + w^{k+1}(R-S).$$

We proceed by induction over k.

- Base case: $k = 0$. ✓
- Induction step: By the condition on w we know that $T - R \leq w(T-P) = w(T-R+R-P) = w(R-P) + w(T-R)$. But if the inequality holds for k then we have

$$\begin{aligned} & w(R-P) + w(T-R) \\ \leq \quad & w(R-P) + w(w(R-P) + \cdots + w^k(R-P) + w^{k+1}(R-S)) \\ = \quad & w(R-P) + w^2(R-P) + \cdots w^{k+1}(R-P) + w^{k+2}(R-S) \end{aligned}$$

which is precisely our inequality for $k+1$.

Proof ctd.

We want to show

$$T - R \leq w(R - P) + \cdots + w^k(R - P) + w^{k+1}(R - S).$$

We proceed by induction over k.

- Base case: $k = 0$. ✓
- Induction step: By the condition on w we know that $T - R \leq w(T - P) = w(T - R + R - P) = w(R - P) + w(T - R)$. But if the inequality holds for k then we have

$$\begin{aligned} & w(R - P) + w(T - R) \\ \leq \quad & w(R - P) + w(w(R - P) + \cdots + w^k(R - P) + w^{k+1}(R - S)) \\ = \quad & w(R - P) + w^2(R - P) + \cdots w^{k+1}(R - P) + w^{k+2}(R - S) \end{aligned}$$

which is precisely our inequality for $k + 1$. ✓

Proof

Any strategy that successfully invades a population of TitForTat has to defect at some point, say on move n. After that

- The strategy may defect forever. ✓
- The strategy may defect $k \geq 0$ times thereafter, and then cooperate. We want to show that

$$T + wP + w^2P + \cdots + w^kP + w^{k+1}S \leq R + wR + w^2R + \cdots w^kR + w^{k+1}R.$$

This is equivalent to

$$T - R \leq w(R-P) + \cdots + w^k(R-P) + w^{k+1}(R-S).$$

Proof

Any strategy that successfully invades a population of TitForTat has to defect at some point, say on move n. After that

- The strategy may defect forever. ✓
- The strategy may defect $k \geq 0$ times thereafter, and then cooperate. We want to show that

$$\begin{aligned} & T + wP + w^2P + \cdots + w^kP + w^{k+1}S \\ \leq \quad & R + wR + w^2R + \cdots w^kR + w^{k+1}R. \end{aligned}$$

This is equivalent to

$$T - R \leq w(R-P) + \cdots + w^k(R-P) + w^{k+1}(R-S).$$ ✓

Which strategies are stable?

In order for a strategy to be collectively stable it has to react to defections from the other side. This can be made precise.

Which strategies are stable?

In order for a strategy to be collectively stable it has to react to defections from the other side. This can be made precise.

Theorem 6.4 *A strategy A is collectively stable if and only if it defects when the B's cumulative score so far is too great. Specifically, it defects on move $n+1$ if the score of B up to game n exceeds that of A playing n games against itself minus*

$$w^{n-1}(T + \frac{w}{1-w}P).$$

Which strategies are stable?

In order for a strategy to be collectively stable it has to react to defections from the other side. This can be made precise.

Theorem 6.4 *A strategy A is collectively stable if and only if it defects when the B's cumulative score so far is too great. Specifically, it defects on move $n+1$ if the score of B up to game n exceeds that of A playing n games against itself minus*

$$w^{n-1}(T + \frac{w}{1-w}P).$$

Proof. We do not provide a proof on this course. □

Which strategies are stable?

In order for a strategy to be collectively stable it has to react to defections from the other side. This can be made precise.

Theorem 6.4 *A strategy A is collectively stable if and only if it defects when the B's cumulative score so far is too great. Specifically, it defects on move $n+1$ if the score of B up to game n exceeds that of A playing n games against itself minus*

$$w^{n-1}(T + \frac{w}{1-w}P).$$

There is some flexibility in the concept of being collectively stable. A strategy typically has a number of points at which it can react to a defection by the other side. This theorem tells us, however, that it must react when it falls too far behind.

Hitting back

However, a strategy cannot afford not to react to a defection by the other side.

Hitting back

However, a strategy cannot afford not to react to a defection by the other side.

Proposition 6.5 *For a nice strategy to be collectively stable it must react in some way to the very first defection of its playing partner.*

Hitting back

However, a strategy cannot afford not to react to a defection by the other side.

Proposition 6.5 *For a nice strategy to be collectively stable it must react in some way to the very first defection of its playing partner.*

Proof. If a nice strategy does not react at all to a defection on, say, move n then the strategy which defects on move n and cooperates in every remaining round will exceed its pay-off by $w^{n-1}(T - S)$, and thus can invade. □

Hitting back

However, a strategy cannot afford not to react to a defection by the other side.

Proposition 6.5 *For a nice strategy to be collectively stable it must react in some way to the very first defection of its playing partner.*

Proof. If a nice strategy does not react at all to a defection on, say, move n then the strategy which defects on move n and cooperates in every remaining round will exceed its pay-off by $w^{n-1}(T - S)$, and thus can invade. $\square$

Note that this proposition does not say that the nice strategy must hit back immediately.

Hitting back

However, a strategy cannot afford not to react to a defection by the other side.

Proposition 6.5 *For a nice strategy to be collectively stable it must react in some way to the very first defection of its playing partner.*

Proof. If a nice strategy does not react at all to a defection on, say, move n then the strategy which defects on move n and cooperates in every remaining round will exceed its pay-off by $w^{n-1}(T - S)$, and thus can invade. □

Note that this proposition does not say that the nice strategy must hit back immediately. TITFORTWOTATS is a strategy which does not satisfy this criterion and therefore is not collectively stable.

Hitting back

However, a strategy cannot afford not to react to a defection by the other side.

Proposition 6.5 *For a nice strategy to be collectively stable it must react in some way to the very first defection of its playing partner.*

Proof. If a nice strategy does not react at all to a defection on, say, move n then the strategy which defects on move n and cooperates in every remaining round will exceed its pay-off by $w^{n-1}(T - S)$, and thus can invade. □

Note that this proposition does not say that the nice strategy must hit back immediately. TitForTwoTats is a strategy which does not satisfy this criterion and therefore is not collectively stable.

We can see here that Axelrod's notion of being retaliatory is required for strategies to be successful.

Invasions of ALWAYSD?

Our starting point for this was the question whether a nice strategy like TITFORTAT could be successful in a climate that rewards selfishness. If we assume that at the start, we only have ALWAYSD strategies, which are collectively stable, how can nice strategies ever develop?

Invasions of AlwaysD?

Our starting point for this was the question whether a nice strategy like TitForTat could be successful in a climate that rewards selfishness. If we assume that at the start, we only have AlwaysD strategies, which are collectively stable, how can nice strategies ever develop?

What if our assumption that invaders always arrive alone is wrong?

Invasions of AlwaysD?

Our starting point for this was the question whether a nice strategy like TitForTat could be successful in a climate that rewards selfishness. If we assume that at the start, we only have AlwaysD strategies, which are collectively stable, how can nice strategies ever develop?

What if our assumption that invaders always arrive alone is wrong?

The whole point of nice strategies is that they do better with each other than AlwaysD does with itself!

Invasions of AlwaysD?

Our starting point for this was the question whether a nice strategy like TitForTat could be successful in a climate that rewards selfishness. If we assume that at the start, we only have AlwaysD strategies, which are collectively stable, how can nice strategies ever develop?

What if our assumption that invaders always arrive alone is wrong?

The whole point of nice strategies is that they do better with each other than AlwaysD does with itself!

We will now examine the situation that occurs when we assume that invaders come in clusters.

Invasion by clusters

Invasion by clusters

Assume we have a PD game with $w = .9$ and pay-off matrix

	defect	cooperate
defect	1	5
cooperate	0	3

Invasion by clusters

Assume we have a PD game with $w = .9$ and pay-off matrix

	defect	cooperate
defect	1	5
cooperate	0	3

The expected number of rounds to be played is 10, so two AlwaysD strategies playing against each other can expect 10 points.

Invasion by clusters

Assume we have a PD game with $w = .9$ and pay-off matrix

	defect	cooperate
defect	1	5
cooperate	0	3

The expected number of rounds to be played is 10, so two AlwaysD strategies playing against each other can expect 10 points.

A single TitForTat strategy introduced in the population can expect 9 points per encounter, and the AlwaysD strategy it plays against will get 14 points.

Invasion by clusters

Assume we have a PD game with $w = .9$ and pay-off matrix

	defect	cooperate
defect	1	5
cooperate	0	3

The expected number of rounds to be played is 10, so two ALWAYSD strategies playing against each other can expect 10 points.

A single TITFORTAT strategy introduced in the population can expect 9 points per encounter, and the ALWAYSD strategy it plays against will get 14 points.

But two TITFORTAT strategies playing against each other can expect 30 points each!

Invasion by clusters

Assume we have a PD game with $w = .9$ and pay-off matrix

	defect	cooperate
defect	1	5
cooperate	0	3

The expected number of rounds to be played is 10, so two AlwaysD strategies playing against each other can expect 10 points.

A single TitForTat strategy introduced in the population can expect 9 points per encounter, and the AlwaysD strategy it plays against will get 14 points.

But two TitForTat strategies playing against each other can expect 30 points each!

Let us assume that there is a proportion of p TitForTat strategies in a population of AlwaysDs. The pay-off for TitForTat is $30p + 9(1 - p)$, whereas for AlwaysD it is $14p + 10(1 - p)$.

Invasion by clusters

Assume we have a PD game with $w = .9$ and pay-off matrix

	defect	cooperate
defect	1	5
cooperate	0	3

The expected number of rounds to be played is 10, so two AlwaysD strategies playing against each other can expect 10 points.

Let us assume that there is a proportion of p TitForTat strategies in a population of AlwaysDs. The pay-off for TitForTat is $30p + 9(1 - p)$, whereas for AlwaysD it is $14p + 10(1 - p)$.

The former is bigger than the latter if and only if $16p - (1 - p) > 0$ which is the case if and only if $p > 1/17 \approx .0588$.

Invasion by clusters

Assume we have a PD game with $w = .9$ and pay-off matrix

	defect	cooperate
defect	1	5
cooperate	0	3

The expected number of rounds to be played is 10, so two ALWAYSD strategies playing against each other can expect 10 points.

Let us assume that there is a proportion of p TITFORTAT strategies in a population of ALWAYSDs. The pay-off for TITFORTAT is $30p + 9(1 - p)$, whereas for ALWAYSD it is $14p + 10(1 - p)$.

The former is bigger than the latter if and only if $16p - (1 - p) > 0$ which is the case if and only if $p > 1/17 \approx .0588$.

In other words, as soon as just 6% of all members of the population are TITFORTAT it pays to be nice!

Invasion in territorial systems

So far we have made the assumption that every individual plays with every other individual.

Invasion in territorial systems

So far we have made the assumption that every individual plays with every other individual.

Alternatively we can assume that every individual only plays with its neighbours (as long as we introduce a notion of locality—we call these systems territorial).

Invasion in territorial systems

So far we have made the assumption that every individual plays with every other individual.

Alternatively we can assume that every individual only plays with its neighbours (as long as we introduce a notion of locality—we call these systems territorial).

In a territorial system, where one strategy might typically have 4 or 8 neighbours, a cluster of TitForTat strategies can invade a population of AlwaysDs provided that they all have at least one TitForTat neighbour.

Invasion in territorial systems

So far we have made the assumption that every individual plays with every other individual.

Alternatively we can assume that every individual only plays with its neighbours (as long as we introduce a notion of locality—we call these systems territorial).

In a territorial system, where one strategy might typically have 4 or 8 neighbours, a cluster of TITFORTAT strategies can invade a population of ALWAYSDs provided that they all have at least one TITFORTAT neighbour.

Hence in a territorial system a cluster of two TITFORTAT strategies which are neighbours can invade a system of ALWAYSD strategies!

Invasion in territorial systems

So far we have made the assumption that every individual plays with every other individual.

Alternatively we can assume that every individual only plays with its neighbours (as long as we introduce a notion of locality—we call these systems territorial).

In a territorial system, where one strategy might typically have 4 or 8 neighbours, a cluster of TitForTat strategies can invade a population of AlwaysDs provided that they all have at least one TitForTat neighbour.

Hence in a territorial system a cluster of two TitForTat strategies which are neighbours can invade a system of AlwaysD strategies!

We will find out more about territorial systems later.

Nice strategies withstand clusters

We have just seen that AlwaysD is vulnerable against invasions by clusters. What about TitForTat?

Nice strategies withstand clusters

We have just seen that ALWAYSD is vulnerable against invasions by clusters. What about TITFORTAT?

Proposition 6.7 *Assume that in our model, any cluster of invaders is small enough that the majority of the native population will interact only with itself. Then if a nice strategy is collectively stable it cannot be invaded by a cluster of individuals.*

Nice strategies withstand clusters

We have just seen that ALWAYSD is vulnerable against invasions by clusters. What about TITFORTAT?

Proposition 6.7 *Assume that in our model, any cluster of invaders is small enough that the majority of the native population will interact only with itself. Then if a nice strategy is collectively stable it cannot be invaded by a cluster of individuals.*

In other words, cooperation can develop in a world of egoism and take over that world, as long as it is not just single individuals that appear.

Nice strategies withstand clusters

We have just seen that ALWAYSD is vulnerable against invasions by clusters. What about TITFORTAT?

Proposition 6.7 *Assume that in our model, any cluster of invaders is small enough that the majority of the native population will interact only with itself. Then if a nice strategy is collectively stable it cannot be invaded by a cluster of individuals.*

In other words, cooperation can develop in a world of egoism and take over that world, as long as it is not just single individuals that appear.

For such a 'take-over' it is necessary that the probability w for interacting with the same individual as well as that of one invader meeting another p being high enough. Once they have taken over, nice strategies are fairly resistant against counter-invasions.

Nice strategies withstand clusters

We have just seen that ALWAYSD is vulnerable against invasions by clusters. What about TITFORTAT?

Proposition 6.7 *Assume that in our model, any cluster of invaders is small enough that the majority of the native population will interact only with itself. Then if a nice strategy is collectively stable it cannot be invaded by a cluster of individuals.*

In other words, cooperation can develop in a world of egoism and take over that world, as long as it is not just single individuals that appear.

For such a 'take-over' it is necessary that the probability w for interacting with the same individual as well as that of one invader meeting another p being high enough. Once they have taken over, nice strategies are fairly resistant against counter-invasions. (Note that a strategy which is immune against invasion by any one strategy might still be invaded by a mixture of two or more strategies.)

Proof

Let A be a nice strategy which is collectively stable, and let B be a potential invader. Assume that the proportion of Bs among the As is p.

Proof

Let A be a nice strategy which is collectively stable, and let B be a potential invader. Assume that the proportion of Bs among the As is p.

Then the expected pay-off in one generation for B is $pP(B, B) + (1 - p)P(B, A)$,

Proof

Let A be a nice strategy which is collectively stable, and let B be a potential invader. Assume that the proportion of Bs among the As is p.

Then the expected pay-off in one generation for B is $pP(B,B) + (1-p)P(B,A)$, whereas that for A is $P(A,A)$.

Proof

Let A be a nice strategy which is collectively stable, and let B be a potential invader. Assume that the proportion of Bs among the As is p.

Then the expected pay-off in one generation for B is $pP(B,B) + (1-p)P(B,A)$, whereas that for A is $P(A,A)$. Hence the expected pay-off for B is larger than that for A if and only if

$$pP(B,B) + (1-p)P(B,A) - P(A,A) > 0.$$

Proof

Let A be a nice strategy which is collectively stable, and let B be a potential invader. Assume that the proportion of Bs among the As is p.

Then the expected pay-off in one generation for B is $pP(B, B) + (1 - p)P(B, A)$, whereas that for A is $P(A, A)$. Hence the expected pay-off for B is larger than that for A if and only if

$$pP(B, B) + (1 - p)P(B, A) - P(A, A) > 0.$$

We know that $P(B, B) \leq P(A, A)$.

Proof

Let A be a nice strategy which is collectively stable, and let B be a potential invader. Assume that the proportion of Bs among the As is p.

Then the expected pay-off in one generation for B is $pP(B, B) + (1 - p)P(B, A)$, whereas that for A is $P(A, A)$. Hence the expected pay-off for B is larger than that for A if and only if

$$pP(B, B) + (1 - p)P(B, A) - P(A, A) > 0.$$

We know that $P(B, B) \leq P(A, A)$. So at most $P(B, B) = P(A, A)$.

Proof

Let A be a nice strategy which is collectively stable, and let B be a potential invader. Assume that the proportion of Bs among the As is p.

Then the expected pay-off in one generation for B is $pP(B,B) + (1-p)P(B,A)$, whereas that for A is $P(A,A)$. Hence the expected pay-off for B is larger than that for A if and only if

$$pP(B,B) + (1-p)P(B,A) - P(A,A) > 0.$$

We know that $P(B,B) \leq P(A,A)$. So at most $P(B,B) = P(A,A)$.

In that case the above is equivalent to

$$(p-1)P(A,A) + (1-p)P(B,A) = (1-p)\Big(P(B,A) - P(A,A)\Big) > 0,$$

which is equivalent to

$$P(B,A) > P(A,A).$$

Proof

Let A be a nice strategy which is collectively stable, and let B be a potential invader. Assume that the proportion of Bs among the As is p.

In that case the above is equivalent to

$$(p-1)P(A,A)+(1-p)P(B,A)=(1-p)\Big(P(B,A)-P(A,A)\Big)>0,$$

which is equivalent to

$$P(B,A)>P(A,A).$$

But that is the condition that B cannot invade A.

Territorial systems

Because territorial systems work on a different basis we have to revisit the notions we created for our ‘global’ systems.

Territorial systems

Because territorial systems work on a different basis we have to revisit the notions we created for our ‘global’ systems.

We say that a strategy can territorially invade a population consisting of another strategy if, eventually, every location in the territory holds an individual employing the new strategy. We say that a strategy is territorially stable if it cannot be territorially invaded.

Territorial systems

Because territorial systems work on a different basis we have to revisit the notions we created for our ‘global’ systems.

We say that a strategy can territorially invade a population consisting of another strategy if, eventually, every location in the territory holds an individual employing the new strategy. We say that a strategy is territorially stable if it cannot be territorially invaded.

It turns out that quite a few of our conclusions for the evolutionary systems where each individual interacts with each other individual carry over to the territorial system.

Stability

It turns out that the two notions of stability are related.

Stability

It turns out that the two notions of stability are related.

Proposition 6.8 *If a strategy is collectively stable then it is territorially stable.*

Stability

It turns out that the two notions of stability are related.

Proposition 6.8 *If a strategy is collectively stable then it is territorially stable.*

Proof. A strategy can only survive in the territorial system if there is an individual in the next generation which carries it on. But that will only be the case if it is more successful against the native strategy than that strategy is against itself, which is precisely the condition for it being able to invade a population consisting entirely of that native strategy. □

Example of a territorial system

Here is a system where

$$w = .3,$$
$$T = 56,$$
$$R = 29,$$
$$P = 6 \text{ and}$$
$$S = 0.$$

Each strategy had four neighbours.

The simulation started with a single ALWAYSD invader in the centre of a system of TITFORTAT.

Example of a territorial system

Here is a system where

$$w = .3,$$
$$T = 56,$$
$$R = 29,$$
$$P = 6 \text{ and}$$
$$S = 0.$$

Each strategy had four neighbours.

The simulation started with a single AlwaysD invader in the centre of a system of TitForTat.

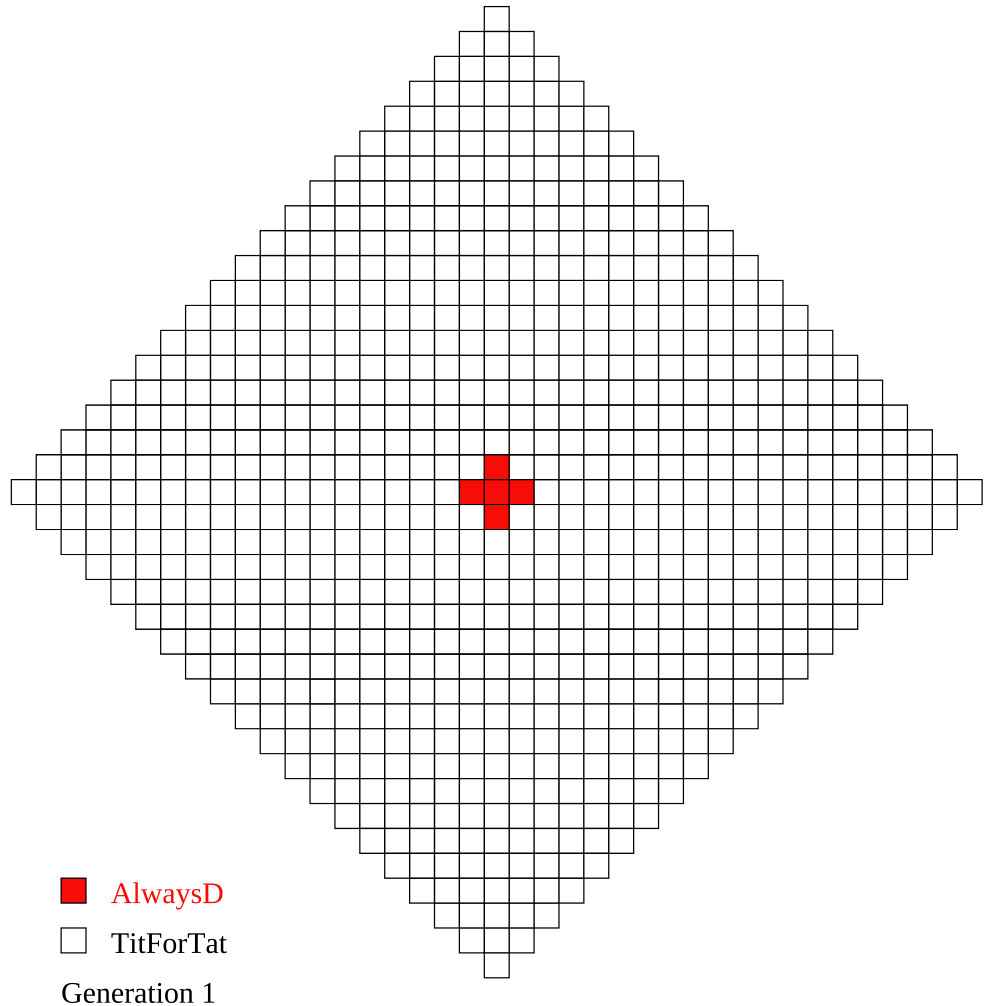

Example of a territorial system

Here is a system where

$$w = .3,$$
$$T = 56,$$
$$R = 29,$$
$$P = 6 \text{ and}$$
$$S = 0.$$

Each strategy had four neighbours.

The simulation started with a single AlwaysD invader in the centre of a system of TitForTat.

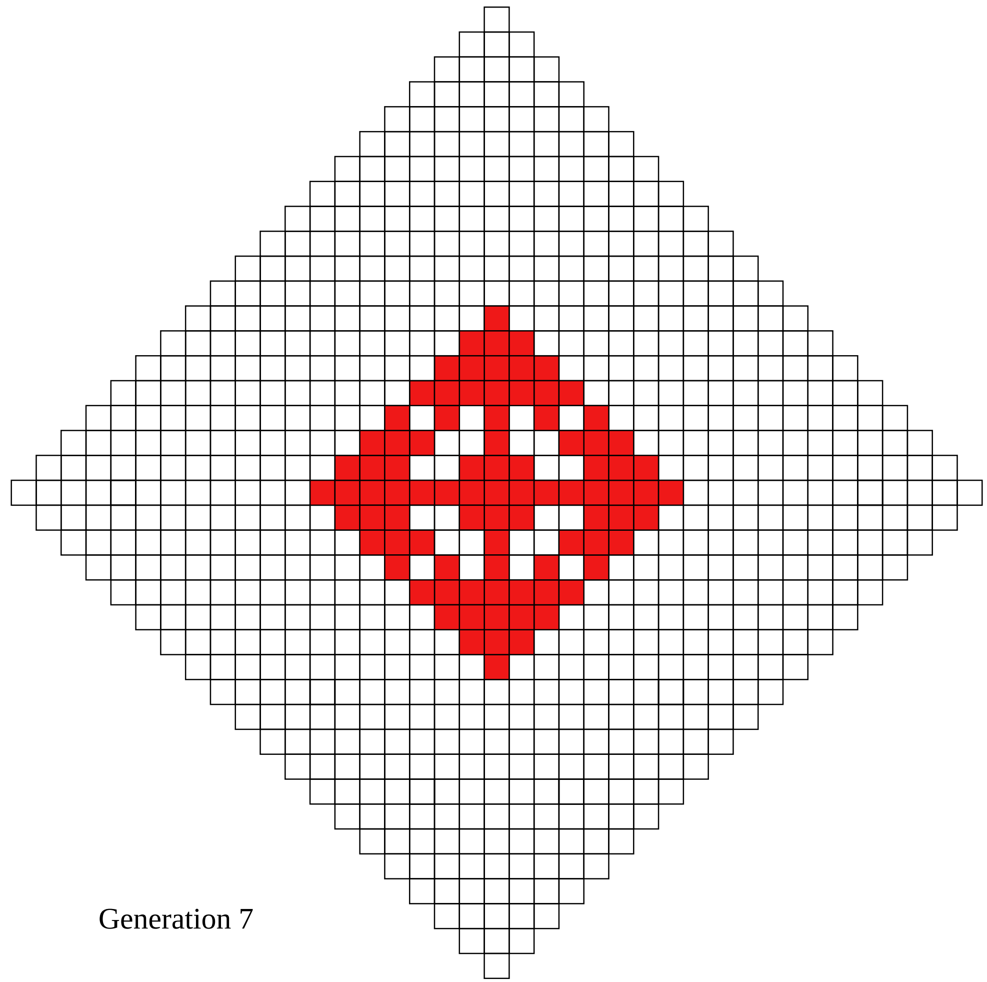

Generation 7

Example of a territorial system

Here is a system where

$$w = .3,$$
$$T = 56,$$
$$R = 29,$$
$$P = 6 \text{ and}$$
$$S = 0.$$

Each strategy had four neighbours.

The simulation started with a single ALWAYSD invader in the centre of a system of TITFORTAT.

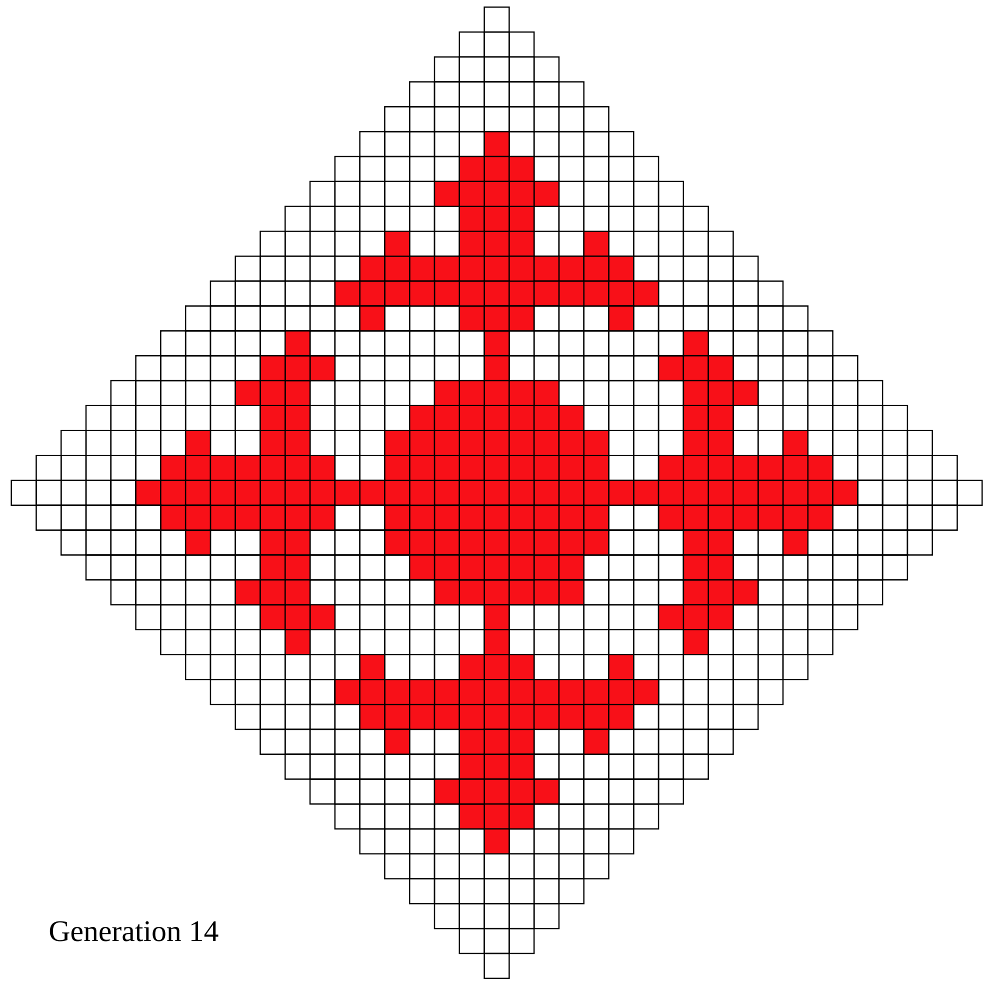

Example of a territorial system

Here is a system where

$$w = .3,$$
$$T = 56,$$
$$R = 29,$$
$$P = 6 \text{ and}$$
$$S = 0.$$

Each strategy had four neighbours.

The simulation started with a single ALWAYSD invader in the centre of a system of TITFORTAT.

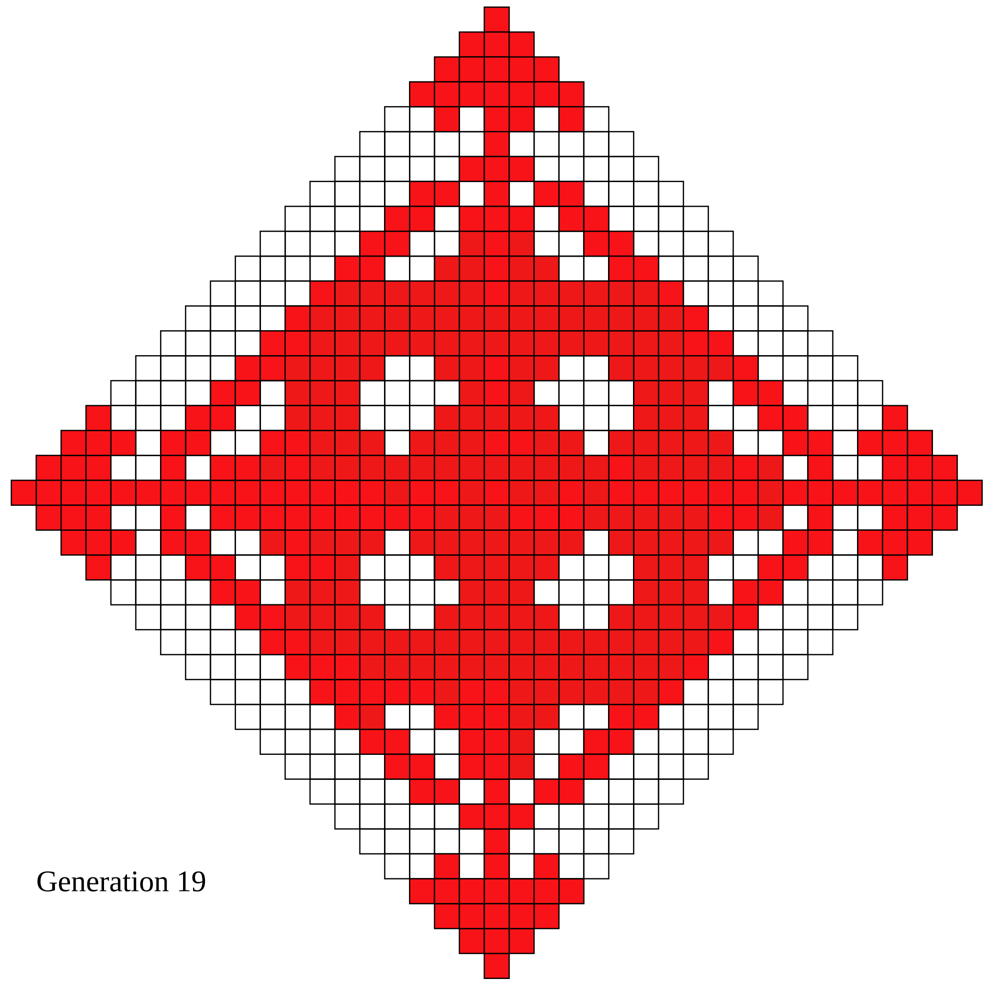

Generation 19

Territorial system with tournament str

Axelrod decided to try a territorial tournament with the strategies submitted to the second tournament. He picked an (18×14) grid with four representatives each. Every strategy had four neighbours. This is a typical final state.

Territorial system with tournament str

Axelrod decided to try a territorial tournament with the strategies submitted to the second tournament. He picked an (18×14) grid with four representatives each. Every strategy had four neighbours. This is a typical final state.

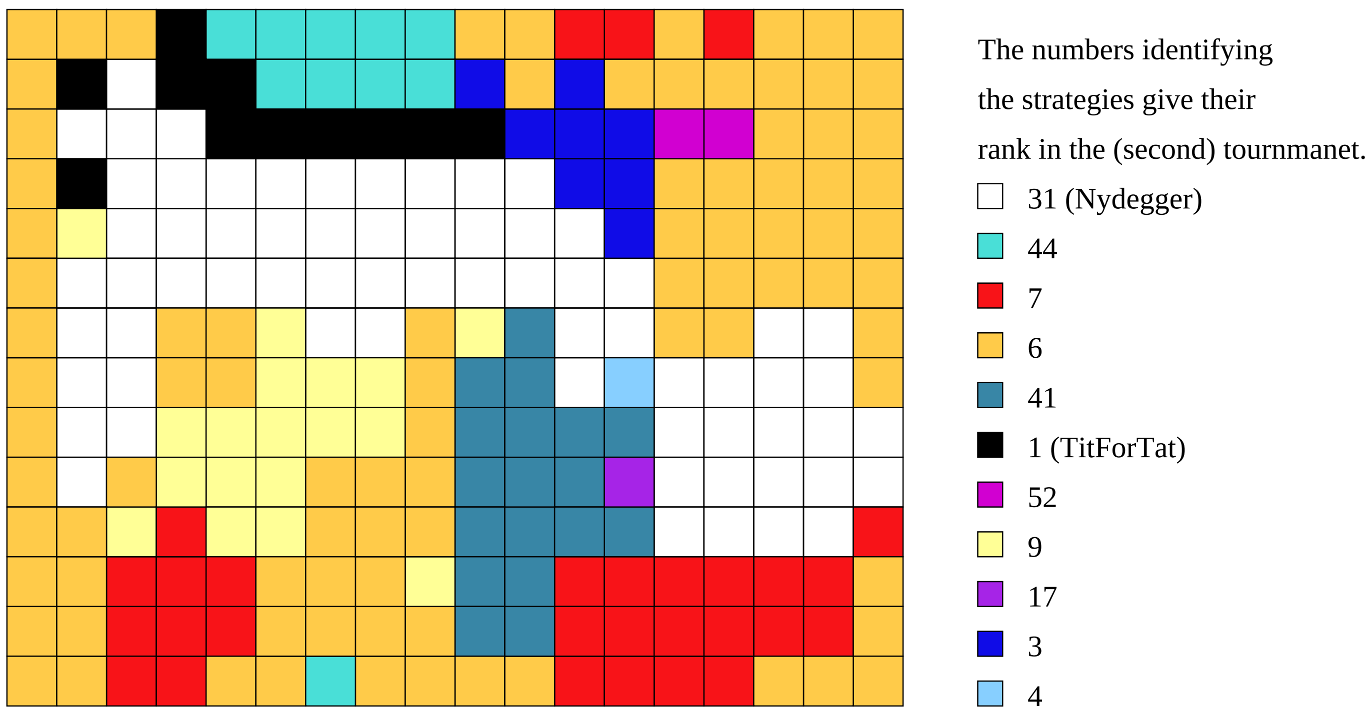

Territorial system with tournament str

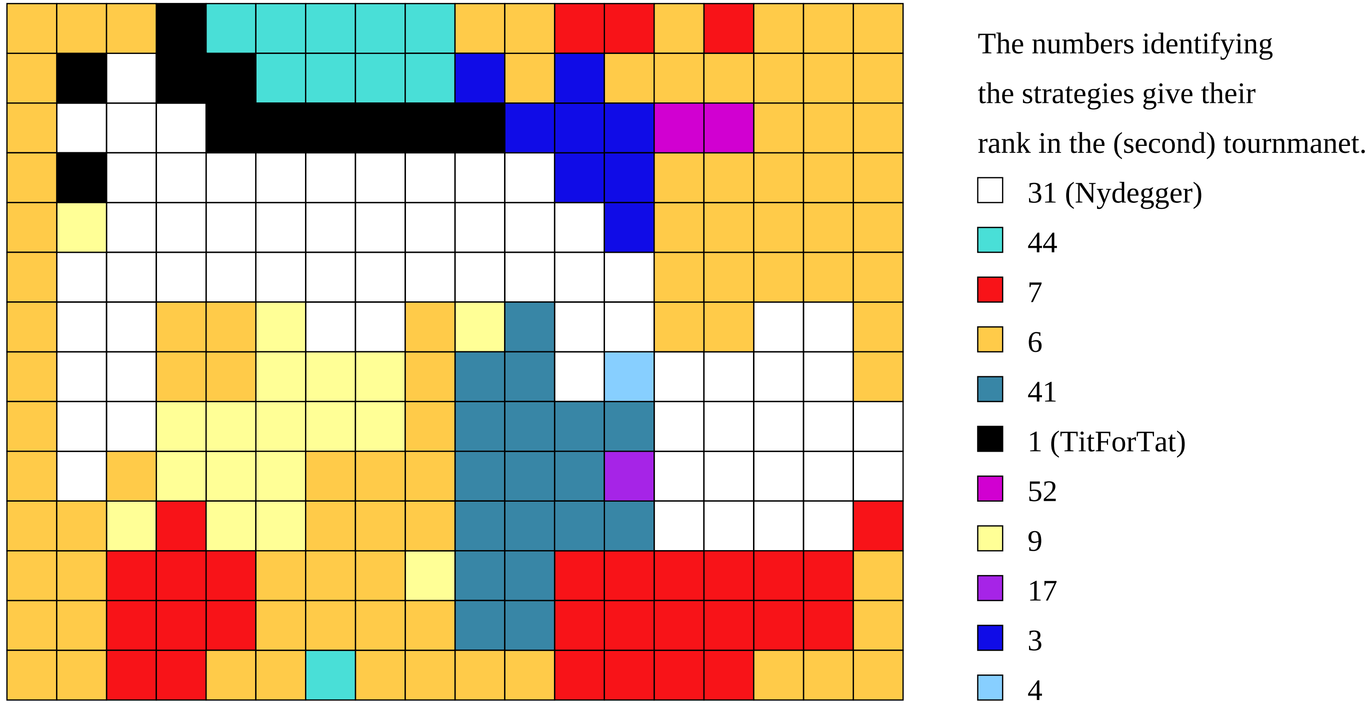

- Not all surviving strategies did well in the tournament;

Territorial system with tournament str

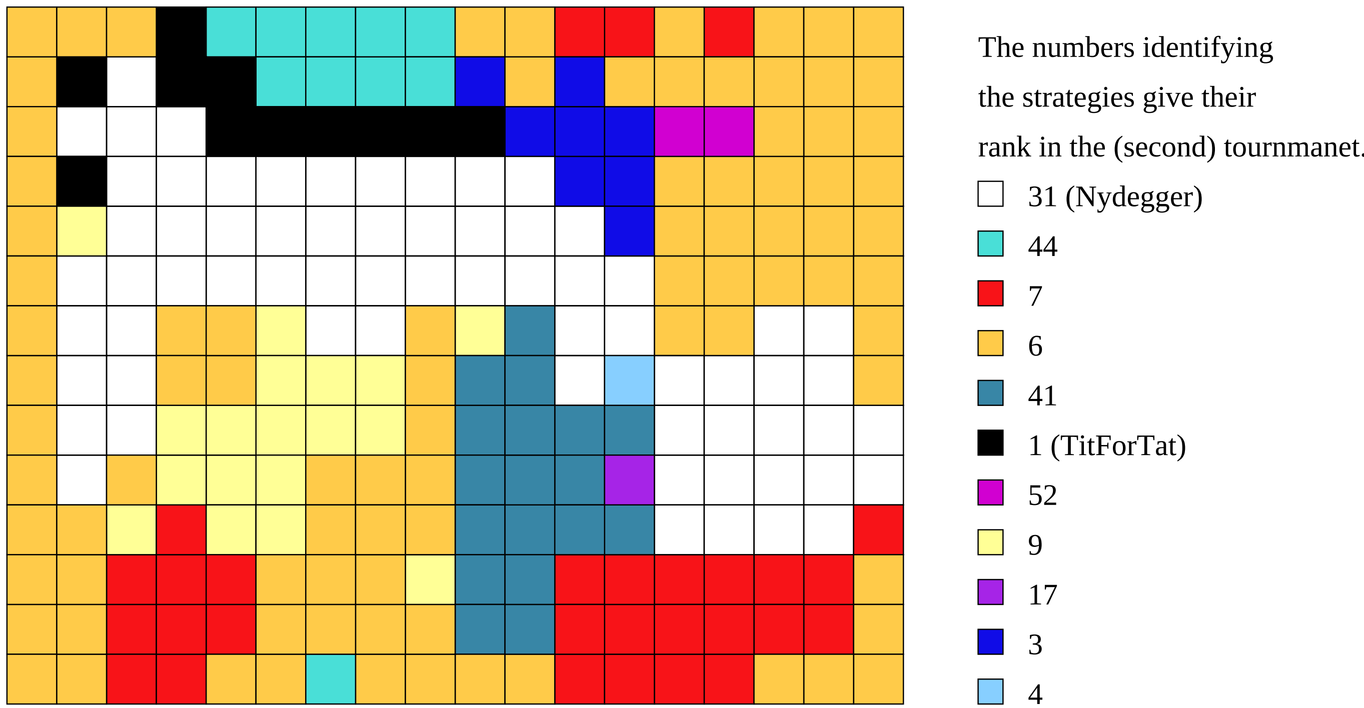

- Not all surviving strategies did well in the tournament;
- not all strategies that did well in the tournament survive;

Territorial system with tournament str

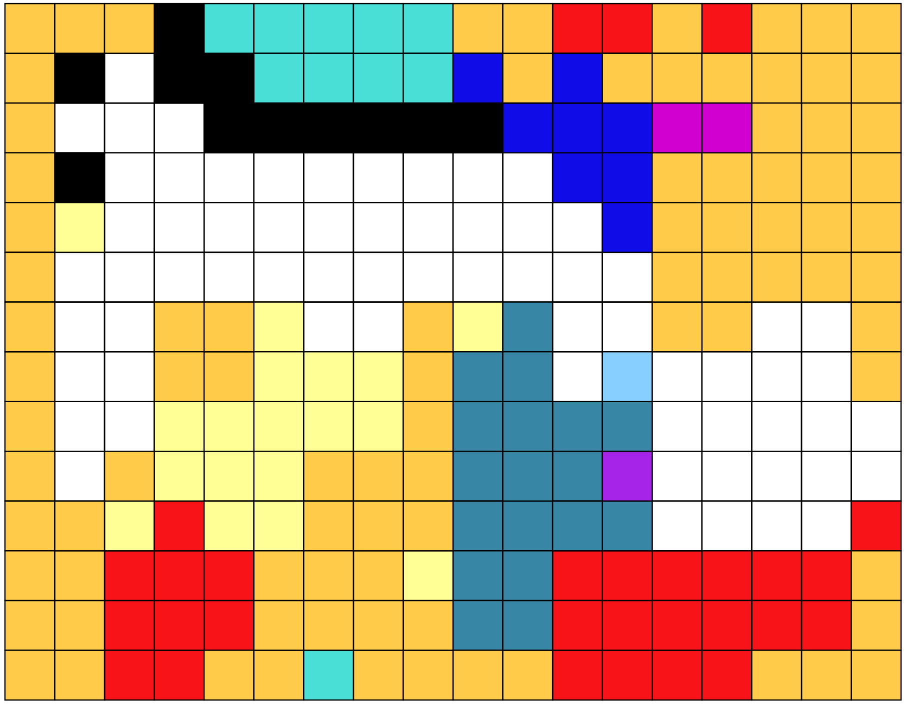

The numbers identifying the strategies give their rank in the (second) tournmanet.

- 31 (Nydegger)
- 44
- 7
- 6
- 41
- 1 (TitForTat)
- 52
- 9
- 17
- 3
- 4

- Not all surviving strategies did well in the tournament;
- not all strategies that did well in the tournament survive;
- most surviving strategies form clusters;

Territorial system with tournament str

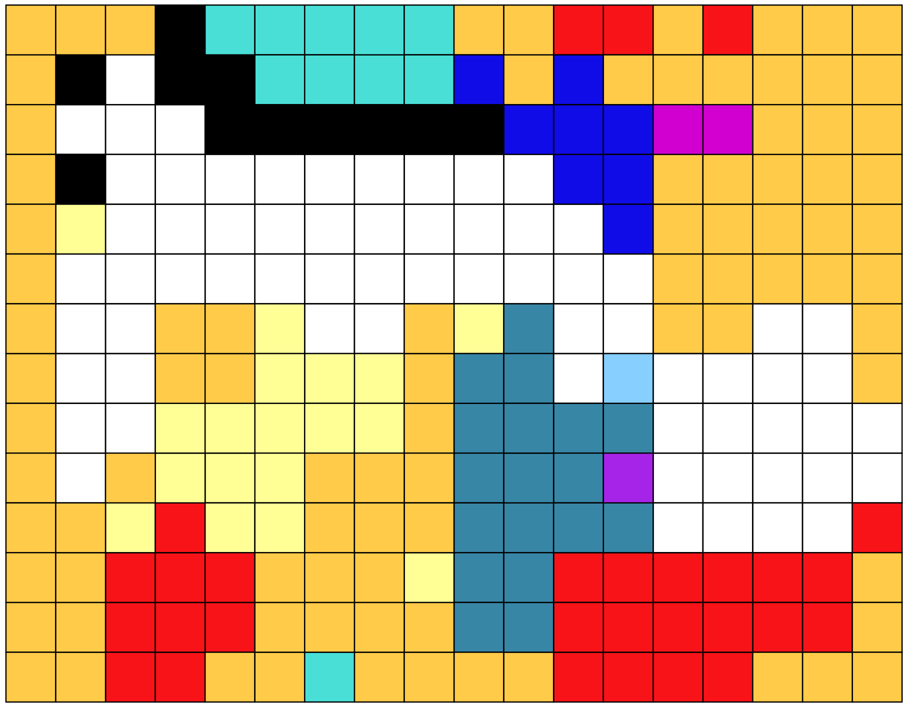

The numbers identifying the strategies give their rank in the (second) tournmanet.

- 31 (Nydegger)
- 44
- 7
- 6
- 41
- 1 (TitForTat)
- 52
- 9
- 17
- 3
- 4

- Not all surviving strategies did well in the tournament;
- not all strategies that did well in the tournament survive;
- most surviving strategies form clusters;
- all surviving rules are nice.

Territorial system with tournament str

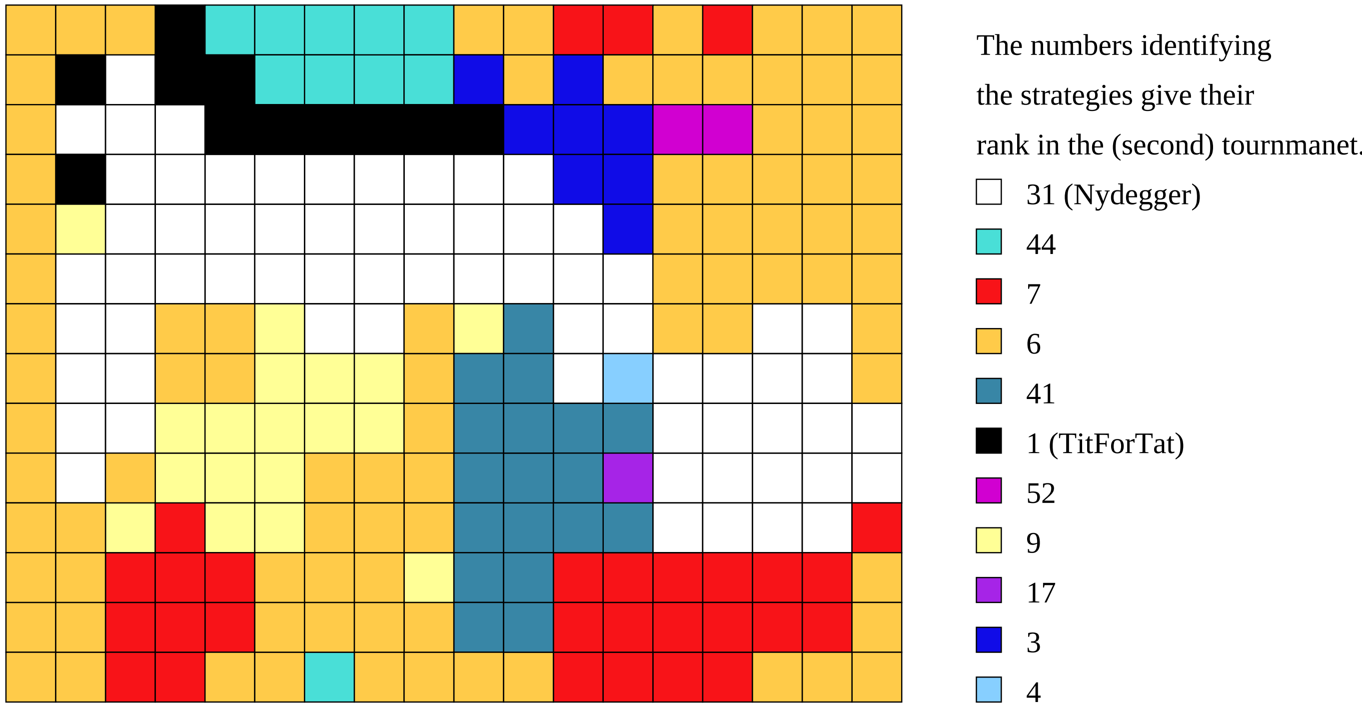

TitForTat did well in all the simulations, on average increasing its numbers from 4 to 17.

Territorial system with tournament str

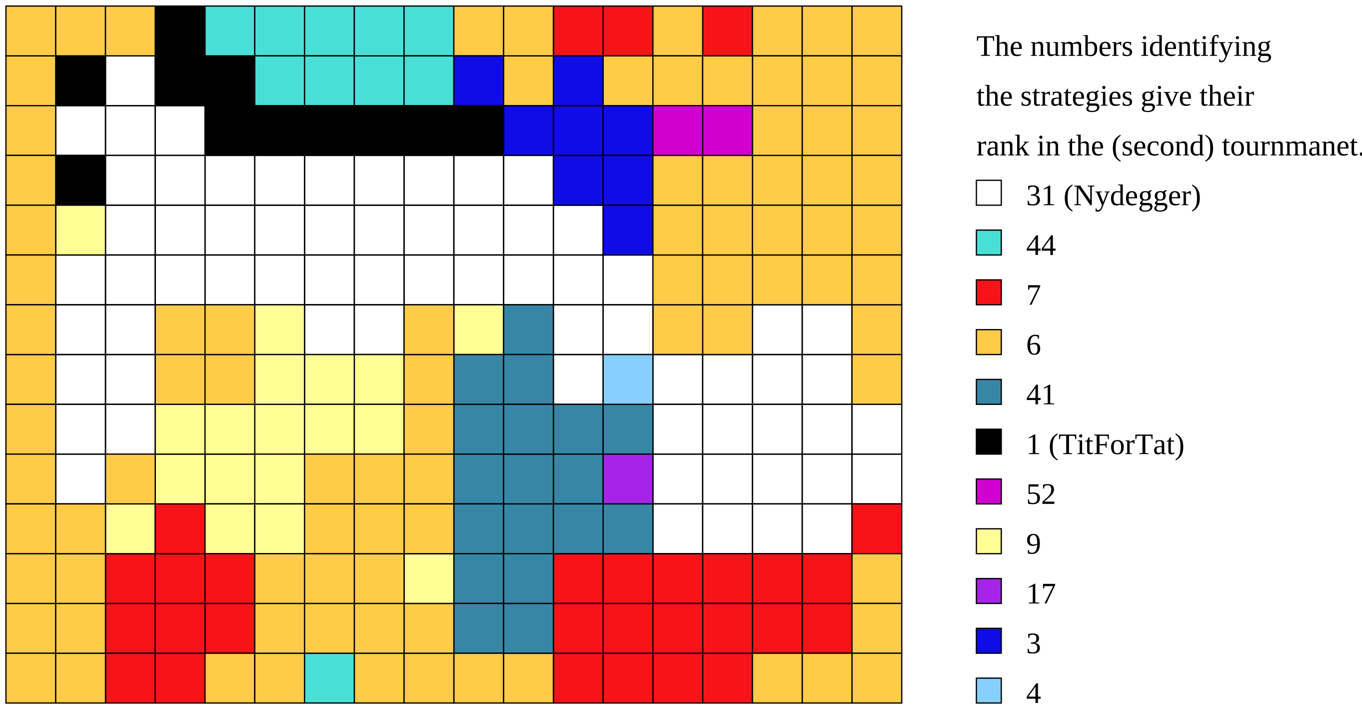

NYDEGGER did extremely well, despite the fact that it only finished 31st in the tournament proper: On average, it had 40 copies.

Territorial system with tournament str

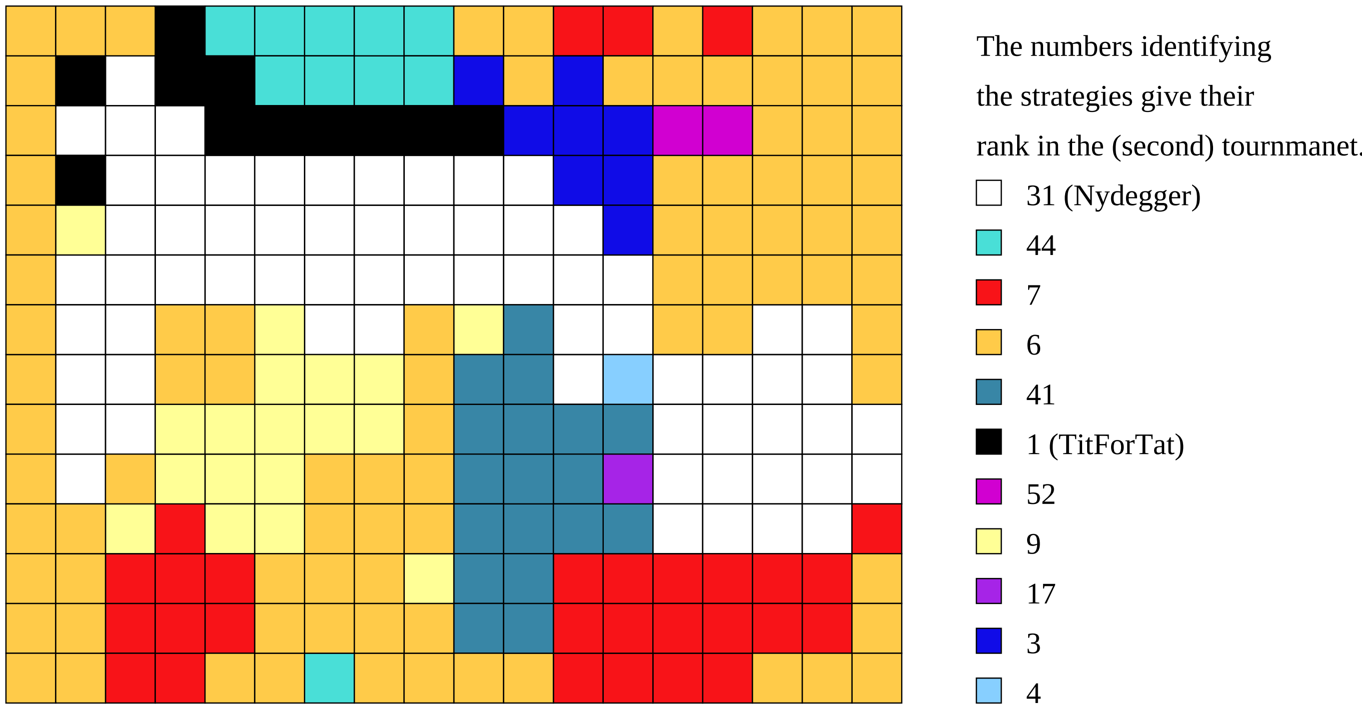

NYDEGGER did extremely well, despite the fact that it only finished 31st in the tournament proper: On average, it had 40 copies.

Which strategies do well? Once all strategies are nice, no more change will occur. So those that do well are the ones that are best at exploiting the original population!

Territorial system with tournament str

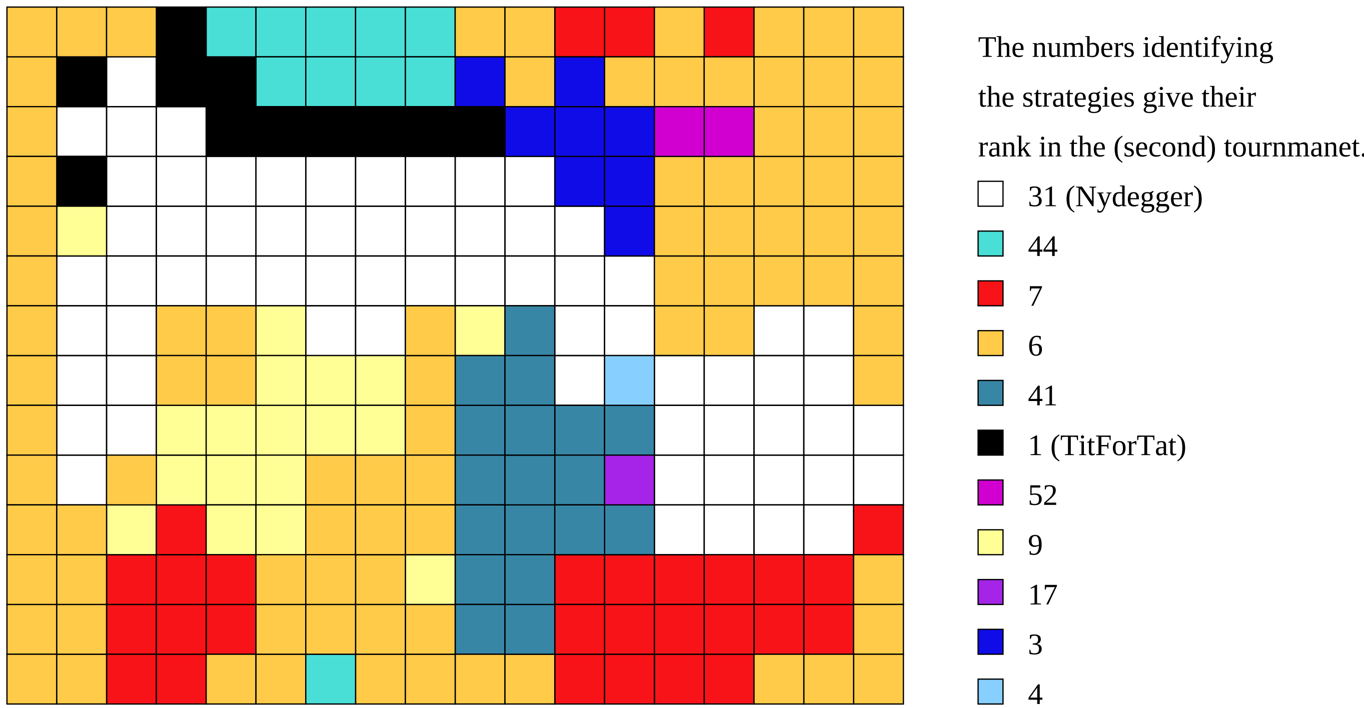

NYDEGGER did extremely well, despite the fact that it only finished 31st in the tournament proper: On average, it had 40 copies.

NYDEGGER is a complicated strategy which makes a number of case-distinctions based on the previous three rounds. It is nice. When the other side defects NYDEGGER sometimes gets it strategy to 'apologize' by cooperating while NYDEGGER defects.

Territorial system with tournament str

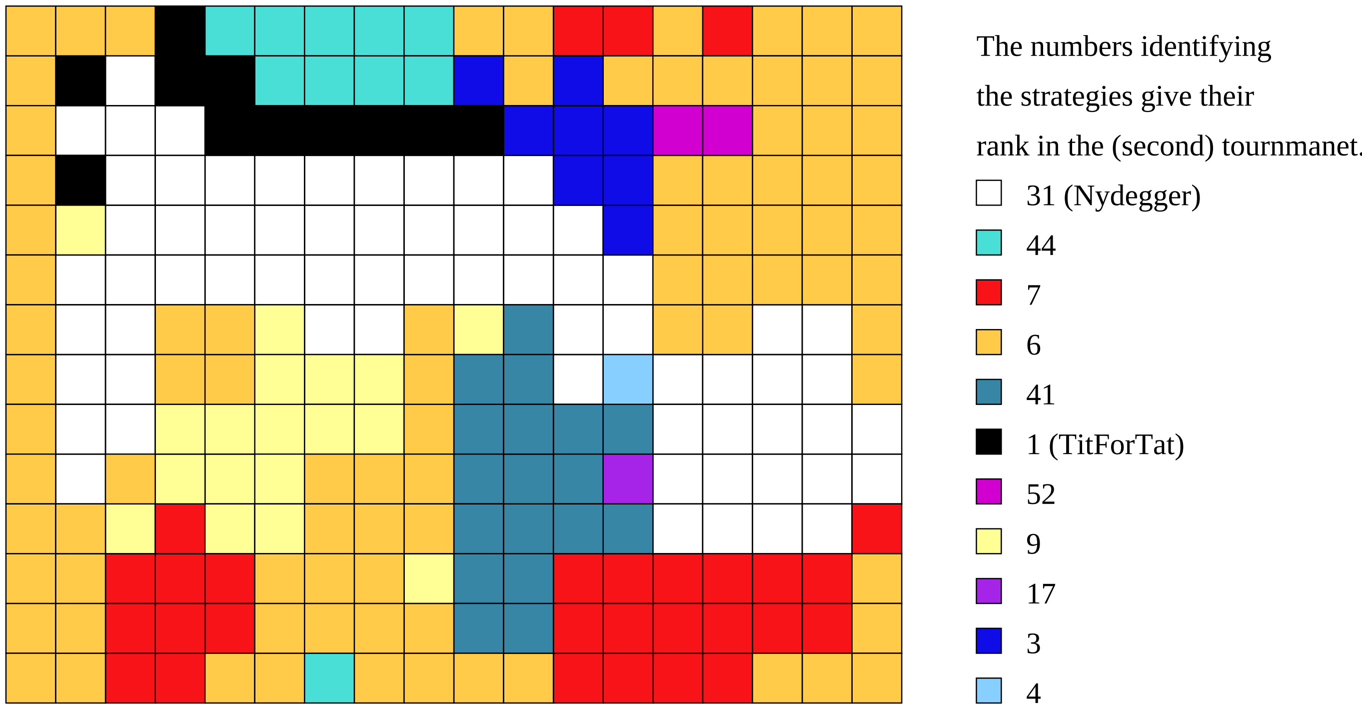

NYDEGGER did extremely well, despite the fact that it only finished 31st in the tournament proper: On average, it had 40 copies.
Whenever a copy of NYDEGGER has one ‘apologetic’ neighbour it will do considerably better than all of its other neighbours. Thus it converts a number of strategies to its own ideas.

What does it all say?

What do Axelrod’s simulations tell us about cooperation, and its evolution?

What does it all say?

What do Axelrod’s simulations tell us about cooperation, and its evolution?

Many people have argued that this set-up is not very realistic when it comes to describing real-world phenomena:

What does it all say?

What do Axelrod's simulations tell us about cooperation, and its evolution?

Many people have argued that this set-up is not very realistic when it comes to describing real-world phenomena:

Question. Can you think of reasons why this might be?

What does it all say?

What do Axelrod's simulations tell us about cooperation, and its evolution?

Many people have argued that this set-up is not very realistic when it comes to describing real-world phenomena:

- The strategies present at the start are a somewhat eclectic mix,

What does it all say?

What do Axelrod's simulations tell us about cooperation, and its evolution?

Many people have argued that this set-up is not very realistic when it comes to describing real-world phenomena:

- The strategies present at the start are a somewhat eclectic mix,
- there is no notion of 'misunderstanding' or 'error',

What does it all say?

What do Axelrod's simulations tell us about cooperation, and its evolution?

Many people have argued that this set-up is not very realistic when it comes to describing real-world phenomena:

- The strategies present at the start are a somewhat eclectic mix,
- there is no notion of 'misunderstanding' or 'error',
- no mutations are present to change the environment.

What does it all say?

What do Axelrod's simulations tell us about cooperation, and its evolution?

Many people have argued that this set-up is not very realistic when it comes to describing real-world phenomena:

- The strategies present at the start are a somewhat eclectic mix,
- there is no notion of 'misunderstanding' or 'error',
- no mutations are present to change the environment.

On the plus side, Axelrod has provided us with an interesting new approach with promising results. Many people have now taken up the idea and applied them to more realistic frameworks.

What does it all say?

What do Axelrod's simulations tell us about cooperation, and its evolution?

Many people have argued that this set-up is not very realistic when it comes to describing real-world phenomena:

- The strategies present at the start are a somewhat eclectic mix,
- there is no notion of 'misunderstanding' or 'error',
- no mutations are present to change the environment.

On the plus side, Axelrod has provided us with an interesting new approach with promising results. Many people have now taken up the idea and applied them to more realistic frameworks.

A lesson to take away: If individuals tend to imitate their successful neighbours, it really pays to be outstandingly successful under at least some circumstances (because that generates converts), even if one's average performance is below that of the average of the entire population.

Beyond Axelrod

Beyond Axelrod

Since Axelrod’s trail-breaking work in the eighties, a lot more has been done. Approaches we will consider at least briefly are

Beyond Axelrod

Since Axelrod's trail-breaking work in the eighties, a lot more has been done. Approaches we will consider at least briefly are

- strategies that can learn;

Beyond Axelrod

Since Axelrod's trail-breaking work in the eighties, a lot more has been done. Approaches we will consider at least briefly are

- strategies that can learn;
- strategies that allow error *via* making use of probabilistic events;

Beyond Axelrod

Since Axelrod's trail-breaking work in the eighties, a lot more has been done. Approaches we will consider at least briefly are

- strategies that can learn;
- strategies that allow error *via* making use of probabilistic events;
- strategies based on finite state machines.

Learning strategies

One approach to strategies that can learn is to assume they have a probability with which they cooperate, and change that probability based on what has happened previously. For the strategy n-PAVLOV, P_n, the probability that it

Learning strategies

One approach to strategies that can learn is to assume they have a probability with which they cooperate, and change that probability based on what has happened previously. For the strategy n-PAVLOV, P_n, the probability that it

repeats its current actions   cooperates (p) changes to

Learning strategies

One approach to strategies that can learn is to assume they have a probability with which they cooperate, and change that probability based on what has happened previously. For the strategy n-PAVLOV, P_n, the probability that it

repeats its current actions

- increases by $1/n$ if it received pay-off R;

cooperates (p) changes to

- $p + 1/n$ if it received pay-off R;

Learning strategies

One approach to strategies that can learn is to assume they have a probability with which they cooperate, and change that probability based on what has happened previously. For the strategy n-PAVLOV, P_n, the probability that it

repeats its current actions

- increases by $1/n$ if it received pay-off R;
- decreases by $2/n$ if it received pay-off S;

cooperates (p) changes to

- $p + 1/n$ if it received pay-off R;
- $p - 2/n$ if it received pay-off S;

Learning strategies

One approach to strategies that can learn is to assume they have a probability with which they cooperate, and change that probability based on what has happened previously. For the strategy n-PAVLOV, P_n, the probability that it

repeats its current actions cooperates (p) changes to

- increases by $1/n$ if it received pay-off R; — $p + 1/n$ if it received pay-off R;
- decreases by $2/n$ if it received pay-off S; — $p - 2/n$ if it received pay-off S;
- decreases by $1/n$ if it received pay-off P; — $p + 1/n$ if it received pay-off P;

Learning strategies

One approach to strategies that can learn is to assume they have a probability with which they cooperate, and change that probability based on what has happened previously. For the strategy n-PAVLOV, P_n, the probability that it

repeats its current actions	cooperates (p) changes to
increases by $1/n$ if it received pay-off R;	$p + 1/n$ if it received pay-off R;
decreases by $2/n$ if it received pay-off S;	$p - 2/n$ if it received pay-off S;
decreases by $1/n$ if it received pay-off P;	$p + 1/n$ if it received pay-off P;
increases by $2/n$ if it received pay-off T.	$p - 2/n$ if it received pay-off T.

Learning strategies

For the strategy n-PAVLOV, P_n, the probability that it

repeats its current actions

- increases by $1/n$ if it received pay-off R;
- decreases by $2/n$ if it received pay-off S;
- decreases by $1/n$ if it received pay-off P;
- increases by $2/n$ if it received pay-off T.

cooperates (p) changes to

- $p + 1/n$ if it received pay-off R;
- $p - 2/n$ if it received pay-off S;
- $p + 1/n$ if it received pay-off P;
- $p - 2/n$ if it received pay-off T.

Note that 0-PAVLOV is just a form of the RANDOM strategy.

Learning strategies

For the strategy n-PAVLOV, P_n, the probability that it

repeats its current actions

- increases by $1/n$ if it received pay-off R;
- decreases by $2/n$ if it received pay-off S;
- decreases by $1/n$ if it received pay-off P;
- increases by $2/n$ if it received pay-off T.

cooperates (p) changes to

- $p + 1/n$ if it received pay-off R;
- $p - 2/n$ if it received pay-off S;
- $p + 1/n$ if it received pay-off P;
- $p - 2/n$ if it received pay-off T.

1-PAVLOV will on the next move

- cooperate if both parties chose the same move in the current round;
- defect if both parties chose different moves in the current round.

The behaviour of Pavlovian strategies

The PAVLOV family of strategies can be viewed as learning from the experience they make.

The behaviour of Pavlovian strategies

The PAVLOV family of strategies can be viewed as learning from the experience they make.

Because the the result of the previous round is treated as a stimulus, this form of learning fits into Skinner's operant conditioning model for learning.

The behaviour of Pavlovian strategies

The PAVLOV family of strategies can be viewed as learning from the experience they make.

Because the the result of the previous round is treated as a stimulus, this form of learning fits into Skinner's operant conditioning model for learning.

This is also a model which is deemed realistic when it comes to describing animal learning.

The behaviour of Pavlovian strategies

The PAVLOV family of strategies can be viewed as learning from the experience they make.

Because the the result of the previous round is treated as a stimulus, this form of learning fits into Skinner's operant conditioning model for learning.

This is also a model which is deemed realistic when it comes to describing animal learning.

When paired with a responsive strategy, the various PAVLOV strategies eventually reach a state where they cooperate almost exclusively.

The behaviour of Pavlovian strategies

The PAVLOV family of strategies can be viewed as learning from the experience they make.

Because the the result of the previous round is treated as a stimulus, this form of learning fits into Skinner's operant conditioning model for learning.

This is also a model which is deemed realistic when it comes to describing animal learning.

When paired with a responsive strategy, the various PAVLOV strategies eventually reach a state where they cooperate almost exclusively.

It can take such a strategy a fairly long time to learn to cooperate when paired with another Pavlovian strategy or TITFORTAT.

The behaviour of Pavlovian strategies

The PAVLOV family of strategies can be viewed as learning from the experience they make.

Because the the result of the previous round is treated as a stimulus, this form of learning fits into Skinner's operant conditioning model for learning.

This is also a model which is deemed realistic when it comes to describing animal learning.

When paired with a responsive strategy, the various PAVLOV strategies eventually reach a state where they cooperate almost exclusively.

It can take such a strategy a fairly long time to learn to cooperate when paired with another Pavlovian strategy or TITFORTAT.

They typically outperform TITFORTAT against versions of the RANDOM strategy, provided the probability for cooperation is at least $1/2$.

Cycling populations

One criticism levelled against Axelrod's work is that he made no attempt to find a realistic initial population. He also did not allow for mutations of strategies.

Cycling populations

One criticism levelled against Axelrod's work is that he made no attempt to find a realistic initial population. He also did not allow for mutations of strategies.

In a world where almost all strategies are nice, a mutation leading to ALWAYSC would be able to survive.

Cycling populations

One criticism levelled against Axelrod’s work is that he made no attempt to find a realistic initial population. He also did not allow for mutations of strategies.

In a world where almost all strategies are nice, a mutation leading to AlwaysC would be able to survive.

But once there are AlwaysC strategies present, other mutants which exploit their generous behaviour might establish themselves (such as AlwaysD).

Cycling populations

One criticism levelled against Axelrod's work is that he made no attempt to find a realistic initial population. He also did not allow for mutations of strategies.

In a world where almost all strategies are nice, a mutation leading to AlwaysC would be able to survive.

But once there are AlwaysC strategies present, other mutants which exploit their generous behaviour might establish themselves (such as AlwaysD).

This suggests that such a population should go through cycles, whereas in Axelrod's world, all populations become stable eventually. (This is due to his lack of making true chance experiments and not allowing mutations.)

Reactive strategies

Two biologies, Nowak and Sigmund, decided to work with reactive strategies. Such a strategy $R(r, p, q)$ will

Reactive strategies

Two biologies, Nowak and Sigmund, decided to work with reactive strategies. Such a strategy $R(r, p, q)$ will

- cooperate on the first move with probability r;

Reactive strategies

Two biologies, Nowak and Sigmund, decided to work with reactive strategies. Such a strategy $R(r, p, q)$ will

- cooperate on the first move with probability r;
- cooperate with probability p if the other player cooperated in the previous round;

Reactive strategies

Two biologies, Nowak and Sigmund, decided to work with reactive strategies. Such a strategy $R(r, p, q)$ will

- cooperate on the first move with probability r;
- cooperate with probability p if the other player cooperated in the previous round;
- cooperate with probability q if the other player defected in the previous round.

Reactive strategies

Two biologies, Nowak and Sigmund, decided to work with reactive strategies. Such a strategy $R(r, p, q)$ will

- cooperate on the first move with probability r;
- cooperate with probability p if the other player cooperated in the previous round;
- cooperate with probability q if the other player defected in the previous round.

Note that value for the first move r plays no big role over time, and therefore we will ignore it in our discussion.

Reactive strategies

Two biologies, Nowak and Sigmund, decided to work with reactive strategies. Such a strategy $R(r, p, q)$ will

- cooperate on the first move with probability r;
- cooperate with probability p if the other player cooperated in the previous round;
- cooperate with probability q if the other player defected in the previous round.

Question. Can you find any of our known strategies among the reactive ones?

Reactive strategies

Two biologies, Nowak and Sigmund, decided to work with reactive strategies. Such a strategy $R(r, p, q)$ will

- cooperate on the first move with probability r;
- cooperate with probability p if the other player cooperated in the previous round;
- cooperate with probability q if the other player defected in the previous round.

Question. Can you find any of our known strategies among the reactive ones?

AlwaysD is $R(0, 0, 0)$.

Reactive strategies

Two biologies, Nowak and Sigmund, decided to work with reactive strategies. Such a strategy $R(r, p, q)$ will

- cooperate on the first move with probability r;
- cooperate with probability p if the other player cooperated in the previous round;
- cooperate with probability q if the other player defected in the previous round.

Question. Can you find any of our known strategies among the reactive ones?

ALWAYSD is $R(0, 0, 0)$. TITFORTAT is $R(1, 1, 0)$.

Reactive strategies

Two biologies, Nowak and Sigmund, decided to work with reactive strategies. Such a strategy $R(r, p, q)$ will

- cooperate on the first move with probability r;
- cooperate with probability p if the other player cooperated in the previous round;
- cooperate with probability q if the other player defected in the previous round.

Question. Can you find any of our known strategies among the reactive ones?

AlwaysD is $R(0, 0, 0)$. TitForTat is $R(1, 1, 0)$. AlwaysC is $R(1, 1, 1)$.

Reactive strategies

Two biologies, Nowak and Sigmund, decided to work with reactive strategies. Such a strategy $R(r, p, q)$ will

- cooperate on the first move with probability r;
- cooperate with probability p if the other player cooperated in the previous round;
- cooperate with probability q if the other player defected in the previous round.

Question. Can you find any of our known strategies among the reactive ones?

AlwaysD is $R(0, 0, 0)$. TitForTat is $R(1, 1, 0)$. AlwaysC is $R(1, 1, 1)$.

The generous version of TitForTat, known as GenTitForTat has $r = p = 1$, but rather than cooperating with probability 0 when the other side has defected last, it will cooperate with probability

$$\min\{1 - \frac{T-R}{R-S}, \frac{R-P}{T-P}\}.$$

Noise

In order to model error Nowak and Sigmund did not allow strategies where p and q were equal to 0 or 1.

Noise

In order to model error Nowak and Sigmund did not allow strategies where p and q were equal to 0 or 1.

Their idea is that no being is that perfect!

Noise

In order to model error Nowak and Sigmund did not allow strategies where p and q were equal to 0 or 1.

Their idea is that no being is that perfect!

The idea is that in the real world, people (or animals) might misread somebody, give out the wrong signals, or otherwise have problems communicating.

The world of reactive strategies

Nowak and Sigmund seeded their population with a random selection of these strategies.

The world of reactive strategies

Nowak and Sigmund seeded their population with a random selection of these strategies.

Most of the time the strategies that do well are those closest to AlwaysD, that is, strategies for which p and q are close to 0.

The world of reactive strategies

Nowak and Sigmund seeded their population with a random selection of these strategies.

Most of the time the strategies that do well are those closest to AlwaysD, that is, strategies for which p and q are close to 0.

However, if there is at least one TitForTat-like strategy in the initial population then everything changes:

The world of reactive strategies

Nowak and Sigmund seeded their population with a random selection of these strategies.

Most of the time the strategies that do well are those closest to ALWAYSD, that is, strategies for which p and q are close to 0.

However, if there is at least one TITFORTAT-like strategy in the initial population then everything changes:

At the start, this strategies (and its copies) struggles to survive.

The world of reactive strategies

Nowak and Sigmund seeded their population with a random selection of these strategies.

Most of the time the strategies that do well are those closest to AlwaysD, that is, strategies for which p and q are close to 0.

However, if there is at least one TitForTat-like strategy in the initial population then everything changes:

At the start, this strategies (and its copies) struggles to survive. The AlwaysD-like strategies live on the strategies for which both p and q are relatively large.

The world of reactive strategies

Nowak and Sigmund seeded their population with a random selection of these strategies.

Most of the time the strategies that do well are those closest to ALWAYSD, that is, strategies for which p and q are close to 0.

However, if there is at least one TITFORTAT-like strategy in the initial population then everything changes:

At the start, this strategies (and its copies) struggles to survive. The ALWAYSD-like strategies live on the strategies for which both p and q are relatively large. But over time, these 'victim' strategies vanish, and then TITFORTAT-like strategies start growing in number at the cost of the ALWAYSD strategies.

The world of reactive strategies

Nowak and Sigmund seeded their population with a random selection of these strategies.

Most of the time the strategies that do well are those closest to AlwaysD, that is, strategies for which p and q are close to 0.

However, if there is at least one TitForTat-like strategy in the initial population then everything changes:

At the start, this strategies (and its copies) struggles to survive. The AlwaysD-like strategies live on the strategies for which both p and q are relatively large. But over time, these 'victim' strategies vanish, and then TitForTat-like strategies start growing in number at the cost of the AlwaysD strategies.

Once the exploiters have gone, GenTitForTat takes over, and then evolution stops.

The world of reactive strategies

Nowak and Sigmund seeded their population with a random selection of these strategies.

Most of the time the strategies that do well are those closest to AlwaysD, that is, strategies for which p and q are close to 0.

However, if there is at least one TitForTat-like strategy in the initial population then everything changes:

At the start, this strategies (and its copies) struggles to survive. The AlwaysD-like strategies live on the strategies for which both p and q are relatively large. But over time, these 'victim' strategies vanish, and then TitForTat-like strategies start growing in number at the cost of the AlwaysD strategies.

Once the exploiters have gone, GenTitForTat takes over, and then evolution stops. Nowak and Sigmund concluded that while TitForTat is vital for cooperation to evolve, persistent patterns of cooperation in the real world are more likely to be due to GenTitForTat.

More complicated strategies

Nowak and Sigmund then ran a second series of simulations, with a wider class of strategies. They decided to allow four random values to describe a strategy, p_1, p_2, p_3, and p_4 so that it would be possible to take the strategy's own last move into account and not just the other player's.

More complicated strategies

Nowak and Sigmund then ran a second series of simulations, with a wider class of strategies. They decided to allow four random values to describe a strategy, p_1, p_2, p_3, and p_4 so that it would be possible to take the strategy's own last move into account and not just the other player's.

A strategy $S(p_1, p_2, p_3, p_4)$ will cooperate on the next move with

- probability p_1 if in the current round, both players cooperated;
- probability p_2 if in the current round, it cooperated while the other side defected;
- probability p_3 if in the current round, it defected while the other side cooperated;
- probability p_4 if in the current round, both sides defected.

More complicated strategies

Nowak and Sigmund then ran a second series of simulations, with a wider class of strategies. They decided to allow four random values to describe a strategy, p_1, p_2, p_3, and p_4 so that it would be possible to take the strategy's own last move into account and not just the other player's.

A strategy $S(p_1, p_2, p_3, p_4)$ will cooperate on the next move with

- probability p_1 if in the current round, both players cooperated;
- probability p_2 if in the current round, it cooperated while the other side defected;
- probability p_3 if in the current round, it defected while the other side cooperated;
- probability p_4 if in the current round, both sides defected.

TitForTat is $S(1, 0, 1, 0)$, and all reactive strategies in general are still represented: they are the ones with

$$p_1 = p_3 \qquad \text{and} \qquad p_2 = p_4.$$

Modelling mutation

They took an initial population of strategies all playing $S(.5, .5, .5, .5)$.

Modelling mutation

They took an initial population of strategies all playing $S(.5, .5, .5, .5)$.

Every 100 generations a small number of randomly chosen mutants was introduced, where the probabilities p_1 to p_4 were changed.

Modelling mutation

They took an initial population of strategies all playing $S(.5, .5, .5, .5)$.

Every 100 generations a small number of randomly chosen mutants was introduced, where the probabilities p_1 to p_4 were changed. They used the proportional evolutionary model rather than the territorial one.

Modelling mutation

They took an initial population of strategies all playing $S(.5, .5, .5, .5)$.

Every 100 generations a small number of randomly chosen mutants was introduced, where the probabilities p_1 to p_4 were changed. They used the proportional evolutionary model rather than the territorial one.

After 10 million generations, 90% of all simulations had reached a state of steady mutual cooperation.

Modelling mutation

They took an initial population of strategies all playing $S(.5, .5, .5, .5)$.

Every 100 generations a small number of randomly chosen mutants was introduced, where the probabilities p_1 to p_4 were changed. They used the proportional evolutionary model rather than the territorial one.

After 10 million generations, 90% of all simulations had reached a state of steady mutual cooperation.

But in only 8.3% of these was the dominating strategy TitForTat or GenTitForTat.

Modelling mutation

They took an initial population of strategies all playing $S(.5, .5, .5, .5)$.

Every 100 generations a small number of randomly chosen mutants was introduced, where the probabilities p_1 to p_4 were changed. They used the proportional evolutionary model rather than the territorial one.

After 10 million generations, 90% of all simulations had reached a state of steady mutual cooperation.

But in only 8.3% of these was the dominating strategy TitForTat or GenTitForTat.

In the remaining ones it was strategies close to $S(1, 0, 0, 1)$ which flourished. But this is precisely the strategy 1-Pavlov!

Modelling mutation

They took an initial population of strategies all playing $S(.5, .5, .5, .5)$.

Every 100 generations a small number of randomly chosen mutants was introduced, where the probabilities p_1 to p_4 were changed. They used the proportional evolutionary model rather than the territorial one.

After 10 million generations, 90% of all simulations had reached a state of steady mutual cooperation.

But in only 8.3% of these was the dominating strategy TitForTat or GenTitForTat.

In the remaining ones it was strategies close to $S(1, 0, 0, 1)$ which flourished. But this is precisely the strategy 1-Pavlov!

This strategy stays with its previous decision if it received the higher of the two pay-offs available (that is T (over R) and P (over S)). Otherwise it changes its mind in the next move.

Modelling mutation

They took an initial population of strategies all playing $S(.5, .5, .5, .5)$.

Every 100 generations a small number of randomly chosen mutants was introduced, where the probabilities p_1 to p_4 were changed. They used the proportional evolutionary model rather than the territorial one.

After 10 million generations, 90% of all simulations had reached a state of steady mutual cooperation.

But in only 8.3% of these was the dominating strategy TitForTat or GenTitForTat.

In the remaining ones it was strategies close to $S(1, 0, 0, 1)$ which flourished. But this is precisely the strategy 1-Pavlov!

This strategy had been disparagingly called 'simpleton' by Rapoport and others: It cooperates with AlwaysD on every other move, and against TitForTat it can be locked into a sequence where it receives repeating pay-offs of T, P, S.

Modelling mutation

They took an initial population of strategies all playing $S(.5, .5, .5, .5)$.

Every 100 generations a small number of randomly chosen mutants was introduced, where the probabilities p_1 to p_4 were changed. They used the proportional evolutionary model rather than the territorial one.

After 10 million generations, 90% of all simulations had reached a state of steady mutual cooperation.

But in only 8.3% of these was the dominating strategy TitForTat or GenTitForTat.

In the remaining ones it was strategies close to $S(1, 0, 0, 1)$ which flourished. But this is precisely the strategy 1-Pavlov!

This strategy makes it hard for strategies like AlwaysD to gain a foothold.

Finite state machines

Other researchers decided to explore simulations where all strategies are represented by finite state machines.

Finite state machines

Other researchers decided to explore simulations where all strategies are represented by finite state machines.

Here is TitForTat as such a machine

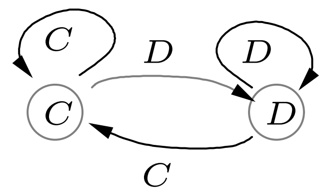

Finite state machines

Other researchers decided to explore simulations where all strategies are represented by finite state machines. Linster conducted a tournament with all strategies which can be expressed using such automata with two states.

Finite state machines

Other researchers decided to explore simulations where all strategies are represented by finite state machines. Linster conducted a tournament with all strategies which can be expressed using such automata with two states. There are several ways of encoding AlwaysC and AlwaysD, and so there are only 22 different strategies in this world.

Finite state machines

Other researchers decided to explore simulations where all strategies are represented by finite state machines. Linster conducted a tournament with all strategies which can be expressed using such automata with two states. He conducted tournaments with mutations (which were very rare), sometimes with invasion forces where as much as 1% of the population could consist of invaders.

Finite state machines

Other researchers decided to explore simulations where all strategies are represented by finite state machines. Linster conducted a tournament with all strategies which can be expressed using such automata with two states.

No single strategy dominated Linster's populations.

Finite state machines

Other researchers decided to explore simulations where all strategies are represented by finite state machines. Linster conducted a tournament with all strategies which can be expressed using such automata with two states.

No single strategy dominated Linster's populations. The GRUDGE strategy did best on average—this could not do well in Nowak and Sigmund's world because defections caused by noise would make it defect forever.

Finite state machines

Other researchers decided to explore simulations where all strategies are represented by finite state machines. Linster conducted a tournament with all strategies which can be expressed using such automata with two states.

No single strategy dominated Linster's populations. The GRUDGE strategy did best on average—this could not do well in Nowak and Sigmund's world because defections caused by noise would make it defect forever.

Other strategies which did well were TITFORTAT, 1-PAVLOV, ALWAYSC and $S(0,1,1,0)$.

Finite state machines

Other researchers decided to explore simulations where all strategies are represented by finite state machines. Linster conducted a tournament with all strategies which can be expressed using such automata with two states.

No single strategy dominated Linster's populations. The GRUDGE strategy did best on average—this could not do well in Nowak and Sigmund's world because defections caused by noise would make it defect forever.

Other strategies which did well were TITFORTAT, 1-PAVLOV, ALWAYSC and $S(0,1,1,0)$.

His results suggest that there may be stable mixes of strategies.

Use of simulations

All the results we have just looked at rely vitally on the use of simulations.

Use of simulations

All the results we have just looked at rely vitally on the use of simulations.

The underlying models are too complicated to obtain any results by purely theoretic means, so simulations are the only method that works in practice.

Use of simulations

All the results we have just looked at rely vitally on the use of simulations.

The underlying models are too complicated to obtain any results by purely theoretic means, so simulations are the only method that works in practice.

Such results have to be taken with a grain of salt in that there are a number of possible errors that may occur when implementing them.

Use of simulations

All the results we have just looked at rely vitally on the use of simulations.

The underlying models are too complicated to obtain any results by purely theoretic means, so simulations are the only method that works in practice.

Such results have to be taken with a grain of salt in that there are a number of possible errors that may occur when implementing them.

It is therefore always a good idea to test the simulation on simple situations where the theory can predict the outcome.

Use of simulations

All the results we have just looked at rely vitally on the use of simulations.

The underlying models are too complicated to obtain any results by purely theoretic means, so simulations are the only method that works in practice.

Such results have to be taken with a grain of salt in that there are a number of possible errors that may occur when implementing them.

It is therefore always a good idea to test the simulation on simple situations where the theory can predict the outcome.

This provides a connection between computer science, and areas where the simulations come from, such as sociology, biology or the political sciences.

More biological games

More biological games

There are many other games that play a role in biology. One example is the Hawk-Dove game.

More biological games

There are many other games that play a role in biology. One example is the Hawk-Dove game.

This is an attempt by biologists to explain behaviour such as ritual fighting, which very rarely damages one of the contestants.

More biological games

There are many other games that play a role in biology. One example is the Hawk-Dove game.

This is an attempt by biologists to explain behaviour such as ritual fighting, which very rarely damages one of the contestants.

Take two male stags fighting for a group of females. They start with a prolonged roaring match, followed by a parallel walk, followed by a direct contest of strength where the two interlock antlers and push against each other. At any time, one of them (usually the intruder) can turn away and break off the fight.

More biological games

There are many other games that play a role in biology. One example is the Hawk-Dove game.

This is an attempt by biologists to explain behaviour such as ritual fighting, which very rarely damages one of the contestants.

Take two male stags fighting for a group of females. They start with a prolonged roaring match, followed by a parallel walk, followed by a direct contest of strength where the two interlock antlers and push against each other. At any time, one of them (usually the intruder) can turn away and break off the fight.

Why does not one of the stags attack the other during the ‘parallel walk’ phase, where the flank of the opponent makes an enticing target? Such an aggressive stag might well have advantages if all other stags would retreat under such an assault.

More biological games

There are many other games that play a role in biology. One example is the Hawk-Dove game.

This is an attempt by biologists to explain behaviour such as ritual fighting, which very rarely damages one of the contestants.

To explain this and a number of similar phenomena, consider a game where there are two strategies, the HAWK and the DOVE strategy.

More biological games

There are many other games that play a role in biology. One example is the Hawk-Dove game.

This is an attempt by biologists to explain behaviour such as ritual fighting, which very rarely damages one of the contestants.

To explain this and a number of similar phenomena, consider a game where there are two strategies, the HAWK and the DOVE strategy.

The DOVE strategy will pretend that it is willing to fight, but when the situation gets serious it will retreat.

More biological games

There are many other games that play a role in biology. One example is the Hawk-Dove game.

This is an attempt by biologists to explain behaviour such as ritual fighting, which very rarely damages one of the contestants.

To explain this and a number of similar phenomena, consider a game where there are two strategies, the HAWK and the DOVE strategy.

The DOVE strategy will pretend that it is willing to fight, but when the situation gets serious it will retreat.

The HAWK, on the other hand, will keep fighting until either it is too severely injured to continue or until the opponent retreats.

The Hawk-Dove game

Let us assume that the contestants in this game are fighting for some 'gain in fitness' G.

The Hawk-Dove game

Let us assume that the contestants in this game are fighting for some 'gain in fitness' G.

If a HAWK meets a DOVE then the DOVE runs away and the HAWK gets a pay-off of G.

The Hawk-Dove game

Let us assume that the contestants in this game are fighting for some ‘gain in fitness’ G.

If a HAWK meets a DOVE then the DOVE runs away and the HAWK gets a pay-off of G.

If HAWK meets HAWK, then a serious fight will ensue. That will reduce the loser’s fitness by C. If either HAWK has a $.5$ chance of winning, the pay-off is $(G-C)/2$ for each of them.

The Hawk-Dove game

Let us assume that the contestants in this game are fighting for some 'gain in fitness' G.

If a HAWK meets a DOVE then the DOVE runs away and the HAWK gets a pay-off of G.

If HAWK meets HAWK, then a serious fight will ensue. That will reduce the loser's fitness by C. If either HAWK has a $.5$ chance of winning, the pay-off is $(G - C)/2$ for each of them. It is typically assumed that C is bigger than G, making $G - C$ negative.

The Hawk-Dove game

Let us assume that the contestants in this game are fighting for some 'gain in fitness' G.

If a HAWK meets a DOVE then the DOVE runs away and the HAWK gets a pay-off of G.

If HAWK meets HAWK, then a serious fight will ensue. That will reduce the loser's fitness by C. If either HAWK has a $.5$ chance of winning, the pay-off is $(G - C)/2$ for each of them.

If two DOVES meet each other they may pretend to fight for a long time, which costs L. So the winner gets $G - L$, and the loser $-L$. If again each side has a $.5$ chance of winning, the expected pay-off is $(G - 2L)/2$.

The Hawk-Dove game

Let us assume that the contestants in this game are fighting for some 'gain in fitness' G.

If a HAWK meets a DOVE then the DOVE runs away and the HAWK gets a pay-off of G.

If HAWK meets HAWK, then a serious fight will ensue. That will reduce the loser's fitness by C. If either HAWK has a $.5$ chance of winning, the pay-off is $(G - C)/2$ for each of them.

If two DOVES meet each other they may pretend to fight for a long time, which costs L. So the winner gets $G - L$, and the loser $-L$. If again each side has a $.5$ chance of winning, the expected pay-off is $(G - 2L)/2$.

This symmetric game can be described by the following matrix giving the pay-off for Player 1.

	HAWK	DOVE
HAWK	$(G - C)/2$	G
DOVE	0	$(G - 2L)/2$

The Hawk-Dove game continued

	Hawk	Dove
Hawk	$(G-C)/2$	G
Dove	0	$(G-2L)/2$

It is assumed that L is much smaller than C.

The Hawk-Dove game continued

	Hawk	Dove
Hawk	$(G-C)/2$	G
Dove	0	$(G-2L)/2$

It is assumed that L is much smaller than C.

The fewer Hawks there are the better the chance of meeting a Dove, and the better Hawks do on average.

The Hawk-Dove game continued

	Hawk	Dove
Hawk	$(G - C)/2$	G
Dove	0	$(G - 2L)/2$

It is assumed that L is much smaller than C.

The fewer Hawks there are the better the chance of meeting a Dove, and the better Hawks do on average.

Here is a specific example. Let $G = 50$, $C = 100$ and $L = 10$ (points). This is the resulting pay-off matrix.

	Hawk	Dove
Hawk	-25	50
Dove	0	15

Extreme populations

	Hawk	Dove
Hawk	-25	50
Dove	0	15

Extreme populations

	Hawk	Dove
Hawk	-25	50
Dove	0	15

In a population consisting entirely of Doves, on average the score from a contest is 15.

Extreme populations

	HAWK	DOVE
HAWK	-25	50
DOVE	0	15

In a population consisting entirely of DOVES, on average the score from a contest is 15.

If a mutant HAWK turns up he will meet a DOVE in every contest gaining 50 points.

Extreme populations

	Hawk	Dove
Hawk	-25	50
Dove	0	15

In a population consisting entirely of Doves, on average the score from a contest is 15.

If a mutant Hawk turns up he will meet a Dove in every contest gaining 50 points. This is much better than a Dove manages, and therefore the Hawk genes will spread quite rapidly, leading to an increase in the number of Hawks.

Extreme populations

	Hawk	Dove
Hawk	-25	50
Dove	0	15

In a population consisting entirely of Doves, on average the score from a contest is 15.

If a mutant Hawk turns up he will meet a Dove in every contest gaining 50 points. This is much better than a Dove manages, and therefore the Hawk genes will spread quite rapidly, leading to an increase in the number of Hawks.

In a population consisting entirely of Hawks the average pay-off from a contest is -25! A single Dove in such a population is at an advantage: While it loses all its fights, it at least gets an average pay-off of 0 as opposed to -25. This would lead to an increase of the number of Doves.

Stable populations?

	HAWK	DOVE
HAWK	-25	50
DOVE	0	15

Stable populations?

	HAWK	DOVE
HAWK	-25	50
DOVE	0	15

In a population with a proportion of p DOVES and $(1-p)$ HAWKS, the average pay-off of one contest
for a DOVE is and that for a HAWK

$$p\frac{G-2L}{2} = 15p, \qquad pG+(1-p)\frac{G-C}{2} = 50p-25(1-p) = 75p-25.$$

Stable populations?

	HAWK	DOVE
HAWK	-25	50
DOVE	0	15

In a population with a proportion of p DOVES and $(1-p)$ HAWKS, the average pay-off of one contest
for a DOVE is and that for a HAWK

$$p\frac{G-2L}{2} = 15p, \qquad pG+(1-p)\frac{G-C}{2} = 50p-25(1-p) = 75p-25.$$

In a balanced population, neither is at an advantage and these are equal. This happens precisely when

$15p = 75p - 25$, which is true if and only if $p = 5/12$.

Stable populations?

	HAWK	DOVE
HAWK	-25	50
DOVE	0	15

In a population with a proportion of p DOVES and $(1-p)$ HAWKS, the average pay-off of one contest

for a DOVE is and that for a HAWK

$$p\frac{G-2L}{2} = 15p, \qquad pG+(1-p)\frac{G-C}{2} = 50p-25(1-p) = 75p-25.$$

In a balanced population, neither is at an advantage and these are equal. This happens precisely when

$15p = 75p - 25$, which is true if and only if $p = 5/12$.

A population with $5/12$ DOVES and $7/12$ HAWKS is stable, and the average pay-off for an individual is $75/12 = 6.25$.

Stable populations?

	Hawk	Dove
Hawk	-25	50
Dove	0	15

In a population with a proportion of p Doves and $(1-p)$ Hawks, the average pay-off of one contest
for a Dove is and that for a Hawk

$$p\frac{G-2L}{2} = 15p, \qquad pG+(1-p)\frac{G-C}{2} = 50p-25(1-p) = 75p-25.$$

A population with $5/12$ Doves and $7/12$ Hawks is stable, and the average pay-off for an individual is $75/12 = 6.25$.

Note that if everybody agreed to be a Dove, there would be a much higher pay-off per contest for the individual, and thus for the entire population! But such a population wouldn't be stable.

Mixed strategy stable populations

	Hawk	Dove
Hawk	-25	50
Dove	0	15

A mixed population is not the only way of reaching a stable population.

Mixed strategy stable populations

	Hawk	Dove
Hawk	-25	50
Dove	0	15

We could interpret the game as one where the pure strategies are the Hawk and Dove strategy, but where each contestant picks a mixed strategy for himself.

Mixed strategy stable populations

	HAWK	DOVE
HAWK	-25	50
DOVE	0	15

We could interpret the game as one where the pure strategies are the HAWK and DOVE strategy, but where each contestant picks a mixed strategy for himself.

Question. Which population of mixed strategies might be stable?

Mixed strategy stable populations

	Hawk	Dove
Hawk	-25	50
Dove	0	15

We could interpret the game as one where the pure strategies are the Hawk and Dove strategy, but where each contestant picks a mixed strategy for himself.

Question. Which population of mixed strategies might be stable?

It turns out that a stable population, that is an equilibrium point for the game is given by everybody adopting the mixed strategy $(7/12, 5/12)$.

Mixed strategy stable populations

	Hawk	Dove
Hawk	-25	50
Dove	0	15

We could interpret the game as one where the pure strategies are the Hawk and Dove strategy, but where each contestant picks a mixed strategy for himself.

It turns out that a stable population, that is an equilibrium point for the game is given by everybody adopting the mixed strategy $(7/12, 5/12)$.

This correspondence between stable populations and mixed strategy equilibrium points only works in the situation where the game is symmetric and there are only two strategies for each player.

Mixed strategy stable populations

	Hawk	Dove
Hawk	-25	50
Dove	0	15

We could interpret the game as one where the pure strategies are the Hawk and Dove strategy, but where each contestant picks a mixed strategy for himself.

It turns out that a stable population, that is an equilibrium point for the game is given by everybody adopting the mixed strategy $(7/12, 5/12)$.

This correspondence between stable populations and mixed strategy equilibrium points only works in the situation where the game is symmetric and there are only two strategies for each player.

We can use Proposition 2.4 to prove that the former gives rise to the latter, but we do not prove the converse.

Equilibria and stability

Since this game is symmetric, an optimal strategy for Player 1 is also optimal for Player 2, which is why we only had to solve one equation.

Equilibria and stability

Since this game is symmetric, an optimal strategy for Player 1 is also optimal for Player 2, which is why we only had to solve one equation.

This means that $(7/12, 5/12)$ is the best response to itself.

Equilibria and stability

Since this game is symmetric, an optimal strategy for Player 1 is also optimal for Player 2, which is why we only had to solve one equation.

This means that $(7/12, 5/12)$ is the best response to itself.

Clearly in a population consisting entirely of such optimal strategies, every invader will do worse against these than they do against themselves, and therefore such a population cannot be invaded.

Equilibria and stability

Since this game is symmetric, an optimal strategy for Player 1 is also optimal for Player 2, which is why we only had to solve one equation.

This means that $(7/12, 5/12)$ is the best response to itself.

Clearly in a population consisting entirely of such optimal strategies, every invader will do worse against these than they do against themselves, and therefore such a population cannot be invaded.

However, if there are more than two strategies around (and contests are on a one-on-one basis) then this changes.

Equilibria and stability

Since this game is symmetric, an optimal strategy for Player 1 is also optimal for Player 2, which is why we only had to solve one equation.

This means that $(7/12, 5/12)$ is the best response to itself.

Clearly in a population consisting entirely of such optimal strategies, every invader will do worse against these than they do against themselves, and therefore such a population cannot be invaded.

However, if there are more than two strategies around (and contests are on a one-on-one basis) then this changes.

Also among biologists the idea that an invader would have to outperform the resident strategy to succeed is not geenerally accepted. They do not consider the equilibrium point as a truly stable situation: Strategies which perform as well against the resident strategy as that strategy does against itself might still spread.

More strategies: Retaliator

The Retaliator starts by behaving similar to a Dove, but when attacked (by a Hawk, for example), it retaliates.

More strategies: RETALIATOR

The RETALIATOR starts by behaving similar to a DOVE, but when attacked (by a HAWK, for example), it retaliates.

Hence it behaves like a HAWK when paired with a HAWK, and like a DOVE when paired with a DOVE.

More strategies: RETALIATOR

The RETALIATOR starts by behaving similar to a DOVE, but when attacked (by a HAWK, for example), it retaliates.

Hence it behaves like a HAWK when paired with a HAWK, and like a DOVE when paired with a DOVE.

	HAWK	DOVE	RETALIATOR
HAWK	$(G-C)/2$	G	$(G-C)/2$
DOVE	0	$(G-2L)/2$	$(G-2L)/2$
RETALIATOR	$(G-C)/2$	$(G-2L)/2$	$(G-2L)/2$

More strategies: RETALIATOR

The RETALIATOR starts by behaving similar to a DOVE, but when attacked (by a HAWK, for example), it retaliates.

Hence it behaves like a HAWK when paired with a HAWK, and like a DOVE when paired with a DOVE.

	HAWK	DOVE	RETALIATOR
HAWK	$(G-C)/2$	G	$(G-C)/2$
DOVE	0	$(G-2L)/2$	$(G-2L)/2$
RETALIATOR	$(G-C)/2$	$(G-2L)/2$	$(G-2L)/2$

If $L=0$ in this game, then the only stable population is a mixture of HAWKS and DOVES, without any RETALIATORS.

More strategies: BULLY

We add a fourth strategy, BULLY, which behaves like a HAWK until it is seriously attacked (by a HAWK, for example) in which case it turns into a DOVE.

More strategies: BULLY

We add a fourth strategy, BULLY, which behaves like a HAWK until it is seriously attacked (by a HAWK, for example) in which case it turns into a DOVE.

There is no stable population at all for this game, so the system oscillates.

More strategies: BULLY

We add a fourth strategy, BULLY, which behaves like a HAWK until it is seriously attacked (by a HAWK, for example) in which case it turns into a DOVE.

There is no stable population at all for this game, so the system oscillates.

So there is nothing in the mathematical theory which says that such a game has to have a stable form!

Change of RETALIATOR

	HAWK	DOVE	RETALIATOR
HAWK	$(G-C)/2$	G	$(G-C)/2$
DOVE	0	$(G-2L)/2$	$(G-2L)/2$
RETALIATOR	$(G-C)/2$	$(G-2L)/2$	$(G-2L)/2$

Change of RETALIATOR

	HAWK	DOVE	RETALIATOR
HAWK	$(G-C)/2$	G	$(G-C)/2$
DOVE	0	$(G-2L)/2$	$(G-2L)/2$
RETALIATOR	$(G-C)/2$	$(G-2L)/2$	$(G-2L)/2$

Question. What happens if we remove all HAWKS from this system?

Change of Retaliator

	Hawk	Dove	Retaliator
Hawk	$(G-C)/2$	G	$(G-C)/2$
Dove	0	$(G-2L)/2$	$(G-2L)/2$
Retaliator	$(G-C)/2$	$(G-2L)/2$	$(G-2L)/2$

Question. What happens if we remove all Hawks from this system?

In the absence of a Hawk, Retaliator and Dove are indistinguishable.

Change of RETALIATOR

	HAWK	DOVE	RETALIATOR
HAWK	$(G-C)/2$	G	$(G-C)/2$
DOVE	0	$(G-2L)/2$	$(G-2L)/2$
RETALIATOR	$(G-C)/2$	$(G-2L)/2$	$(G-2L)/2$

We can fix this by assuming that when paired with a DOVE, there is a slight chance that RETALIATOR may find out that escalating the fight will win it.

Change of RETALIATOR

	HAWK	DOVE	RETALIATOR
HAWK	$(G - C)/2$	G	$(G - C)/2$
DOVE	0	$(G - 2L)/2$	$(G - 2L)/2$
RETALIATOR	$(G - C)/2$	$(G - 2L)/2$	$(G - 2L)/2$

We can fix this by assuming that when paired with a DOVE, there is a slight chance that RETALIATOR may find out that escalating the fight will win it. It then seems only fair to assume that a HAWK has an advantage when paired with a RETALIATOR since it will escalate first.

Change of RETALIATOR

	HAWK	DOVE	RETALIATOR
HAWK	$(G-C)/2$	G	$(G-C)/2$
DOVE	0	$(G-2L)/2$	$(G-2L)/2$
RETALIATOR	$(G-C)/2$	$(G-2L)/2$	$(G-2L)/2$

We can fix this by assuming that when paired with a DOVE, there is a slight chance that RETALIATOR may find out that escalating the fight will win it. It then seems only fair to assume that a HAWK has an advantage when paired with a RETALIATOR since it will escalate first.

	HAWK	DOVE	RETALIATOR
HAWK	$(G-C)/2$	G	$(G-C+E)/2$
DOVE	0	$G/2$	$(G-E)/2$
RETALIATOR	$(G-C-E)/2$	$(G+E)/2$	$G/2$

Stable populations for three strategies

	Hawk	Dove	Retaliator
Hawk	$(G-C)/2$	G	$(G-C+E)/2$
Dove	0	$G/2$	$(G-E)/2$
Retaliator	$(G-C-E)/2$	$(G+E)/2$	$G/2$

Stable populations for three strategies

	HAWK	DOVE	RETALIATOR
HAWK	$(G-C)/2$	G	$(G-C+E)/2$
DOVE	0	$G/2$	$(G-E)/2$
RETALIATOR	$(G-C-E)/2$	$(G+E)/2$	$G/2$

This game has two stable populations, one consisting entirely of RETALIATORS and one consisting of a mixture of HAWKS and DOVES.

Stable populations for three strategies

	HAWK	DOVE	RETALIATOR
HAWK	$(G-C)/2$	G	$(G-C+E)/2$
DOVE	0	$G/2$	$(G-E)/2$
RETALIATOR	$(G-C-E)/2$	$(G+E)/2$	$G/2$

This game has two stable populations, one consisting entirely of RETALIATORS and one consisting of a mixture of HAWKS and DOVES.

We will not work any of these out in detail; they are just meant to give an idea of the variety of situations that are possible with this setup.

Yet more biological games

There are other strategies one might add to this game, and there are different games that describe slightly different situations.

Yet more biological games

There are other strategies one might add to this game, and there are different games that describe slightly different situations.

In particular when the potential gain G is small, contests often become asymmetric: The two contestants do not fight on equal grounds, for example because one is an intruder and the other on home territory.

Yet more biological games

There are other strategies one might add to this game, and there are different games that describe slightly different situations.

In particular when the potential gain G is small, contests often become asymmetric: The two contestants do not fight on equal grounds, for example because one is an intruder and the other on home territory.

A certain kind of butterfly, for example, seeks out sunny spots in the hope of being joined by a female. If the spot is already occupied, the intruder gives up very quickly.

Yet more biological games

There are other strategies one might add to this game, and there are different games that describe slightly different situations.

In particular when the potential gain G is small, contests often become asymmetric: The two contestants do not fight on equal grounds, for example because one is an intruder and the other on home territory.

In such fights there typically is a considerable advantage for the home side. This seems sensible, because the home side knows the territory in question, and there are good reasons for striving to be a resident. This makes fights a lot shorter, and thus less costly, and gives a 'natural' solution, namely a stable population.

Yet more biological games

There are other strategies one might add to this game, and there are different games that describe slightly different situations.

In particular when the potential gain G is small, contests often become asymmetric: The two contestants do not fight on equal grounds, for example because one is an intruder and the other on home territory.

In such fights there typically is a considerable advantage for the home side. This seems sensible, because the home side knows the territory in question, and there are good reasons for striving to be a resident. This makes fights a lot shorter, and thus less costly, and gives a 'natural' solution, namely a stable population.

However there's a type of Mexican social spider which, when disturbed tries to find a new hiding place. If it darts into a crevice occupied by another spider the occupant will leave and seek a new place for itself.

Summary of Section 6

- The indefinitely repeated PD can be used to model evolution. There is no single best strategy if w is large enough.

Summary of Section 6

- The indefinitely repeated PD can be used to model evolution. There is no single best strategy if w is large enough.
- The core concept is collective stability, that is, being safe from invasions. Examples of such strategies are AlwaysD (always), and TitForTat (if w is large enough).

Summary of Section 6

- The indefinitely repeated PD can be used to model evolution. There is no single best strategy if w is large enough.
- The core concept is collective stability, that is, being safe from invasions.
- Nice strategies have to react to the first defection of a playing partner to be collectively stable, and one can define a rule of when a collectively stable strategy will have to defect.

Summary of Section 6

- The indefinitely repeated PD can be used to model evolution. There is no single best strategy if w is large enough.
- The core concept is collective stability, that is, being safe from invasions.
- Nice strategies have to react to the first defection of a playing partner to be collectively stable.
- Invasion becomes easier for nice strategies if they invade in clusters, but nice collectively stable strategies are safe from invasions. In many ways, TitForTat is as successful a strategy as it can be in such a world.

Summary of Section 6

- The indefinitely repeated PD can be used to model evolution. There is no single best strategy if w is large enough.
- The core concept is collective stability, that is, being safe from invasions.
- Nice strategies have to react to the first defection of a playing partner to be collectively stable.
- Invasion becomes easier for nice strategies if they invade in clusters, but nice collectively stable strategies are safe from invasions.
- We can model localized interaction in territorial system.

Summary of Section 6

- The indefinitely repeated PD can be used to model evolution. There is no single best strategy if w is large enough.
- The core concept is collective stability, that is, being safe from invasions.
- Nice strategies have to react to the first defection of a playing partner to be collectively stable.
- Invasion becomes easier for nice strategies if they invade in clusters, but nice collectively stable strategies are safe from invasions.
- We can model localized interaction in territorial system.
- Beyond Axelrod, people have introduced noise and simple learning (with probabilistic strategies or finite state machines).

Summary of Section 6

- The indefinitely repeated PD can be used to model evolution. There is no single best strategy if w is large enough.
- The core concept is collective stability, that is, being safe from invasions.
- Nice strategies have to react to the first defection of a playing partner to be collectively stable.
- Invasion becomes easier for nice strategies if they invade in clusters, but nice collectively stable strategies are safe from invasions.
- We can model localized interaction in territorial system.
- Beyond Axelrod, people have introduced noise and simple learning.
- There are other games such as the Hawk-Dove game that are used in biology to explain the point of balance of stable populations.