PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY B BIOLOGICAL SCIENCES

The Spatial Struggle of Tit-For-Tat and Defect

V. C. L. Hutson and G. T. Vickers

Phil. Trans. R. Soc. Lond. B 1995 **348**, 393-404
doi: 10.1098/rstb.1995.0077

The spatial struggle of tit-for-tat and defect

V. C. L. HUTSON AND G. T. VICKERS

School of Mathematics and Statistics, Applied Mathematics Section, The University of Sheffield, Sheffield S3 7RH, U.K.

SUMMARY

The pioneering work by Trivers (1971), Axelrod (1984) and Axelrod & Hamilton (1981) has stimulated continuing interest in explaining the evolution of cooperation by game theory, in particular, the iterated prisoner's dilemma and the strategy of tit-for-tat. However these models suffer from a lack of biological reality, most seriously because it is assumed that players meet opponents at random from the population and, unless the population is very small, this excludes the repeated encounters necessary for tit-for-tat to prosper. To meet some of the objections, we consider a model with two types of players, defectors (D) and tit-for-tat players (T), in a spatially homogeneous environment with player densities varying continuously in space and time. Players only encounter neighbours but move at random in space. The analysis demonstrates major new conclusions, the three most important being as follows. First, stable coexistence with constant densities of both players is possible. Second, stable coexistence in a pattern (a spatially inhomogeneous stationary state) may be possible when it is impossible for constant distributions (even unstable ones) to exist. Third, invasion by a very small number of T-players is sometimes possible (in contrast with the usual predictions) and so a mutation to tit-for-tat may lead to a population of defectors being displaced by the T-players.

1. INTRODUCTION

One of the most interesting problems in evolutionary biology is to find a convincing explanation for the evolution of cooperative (or social) behaviour among unrelated individuals in a species. There are a number of mechanisms that have been suggested (see for example the review by Mesterton-Gibbons & Dugatkin 1992) and it is likely that each should be valid under certain circumstances and even that all could sometimes operate together. However, one particular class of model, that based on repeated games – in particular the iterated prisoner's dilemma (IPD) – appears particularly attractive and has been much studied over the past 25 years. One of the most interesting strategies is tit-for-tat (TFT), where a player cooperates until the other defects after which he does the same as his opponent did on the previous play. Some early references are Trivers (1971), Brown *et al.* (1982), Axelrod & Hamilton (1981), Maynard Smith (1982), Axelrod (1984) and elaborations of the original model are introduced in for example Feldman & Thomas (1987), Nowak & Sigmund (1990, 1994) and Nowak *et al.* (1994). Much of the literature on IPD is concerned with theoretical aspects and there is rather little published, quantitative work on the application of these ideas to real populations. Nonetheless cooperation is widely observed even in situations where kin selection does not operate, and a great deal of attention has focused on the model recently. To quote just one authority (Dawkins 1989, p. 229): 'These conditions [for the operation of the IPD and TFT] are certainly met all around the living kingdoms'.

Except in certain special circumstances, the above models appear to ignore certain features which seem central to the problem. For example, a basic assumption is that players repeatedly play either a single opponent or a range of opponents chosen at random from the population. Both of these assumptions are unrealistic. For unless the total population is very small, the assumption of random interactions with the whole population will imply that one is unlikely to meet an opponent again. This is inconsistent with the repeated interactions which are essential if the IPD is to predict the evolution of cooperation (see Dugatkin & Wilson, p. 689, 1991). Also, players will only meet their near neighbours (in space) and mobility will affect the outcome (as whether or not players have met before is clearly crucial). For the same reason, the death rate will also be vital. It is inevitable then that the spatial relation of players is a crucial factor. A number of other shortcomings can also be mentioned. One of the most important is that the classical models do not explain how a small number of players who always play TFT (T-players) can invade a population of players who always defect (D-players). Nor do they allow a mixed population of T- and D-players to coexist stably. There are a number of other criticisms which have been made. In some cases suggestions (sometimes perhaps a little artificial) have been given as to more effective mechanisms, some representative references with interesting comments being Eshel & Cavalli-Sforza (1982), Dugatkin & Wilson (1991), Enquist & Leimar (1991), Mesterton-Gibbons (1992), Mesterton-Gibbons & Dugatkin (1992) and Houston (1993). The importance of a spatial structure for the

Phil. Trans. R. Soc. Lond. B (1995) **348**, 393–404
Printed in Great Britain

IPD has long been realized (see, for example, Axelrod 1981, 1984). However, only recently has a start been made on the analysis of detailed quantitative models (see the discussions in Dugatkin & Wilson 1991; Ferriere & Michod 1994).

Our objective in this investigation is to introduce a model which, although as simple as possible, contains the key factors necessary to make the IPD game with T- and D-players more realistic. In the model players can only play their neighbours in space and 'remember' their previous opponents only until they move out of the neighbourhood. Thus if the opponent subsequently returns it is treated as a stranger. This assumption is adopted partly because it simplifies the model somewhat and is at the same time a good approximation in many circumstances. The model could be modified to take account of players with more powerful memories, but this is very unlikely to have a crucial effect on the results. The environment is assumed spatially homogeneous, and players move in space in a random manner at a rate controlled by the parameters μ_T (for T-players) and μ_D (for D-players). We suppose that the spatial density of players is large enough for a continuous model to be a good approximation, and represent the random spatial motion by a standard diffusion approximation leading to a pair of reaction-diffusion equations. To model the fact that players remember individuals whom they have played before, it is necessary to introduce a 'getting-to-know' function g, the proportion of pairs of T- and D-players in a neighbourhood who have already met; g thus varies in both space and time. The relation of our model with classical IPD models can be seen by comparing the roles of g and w the (constant) probability of players meeting again. The value of w should clearly be derived from the dynamics of the model rather than imposed on it, as emphasized in Houston (1993) and Dugatkin & Wilson (1991); in the current model g plays exactly this role. We believe that the present model represents a significant improvement over previous mathematical attempts to describe the IPD and explain cooperation. Our results show that spatial structure with diffusion sometimes leads to significantly new predictions. Thus in any quantitative study of the IPD in real populations effort should be made to measure the dispersion rates of the various behaviour-types as well as the pay-offs.

The model is derived and the underlying assumptions carefully described in §2, where the relation to previous models (including those with spatial effects) is also discussed. We then carry out an extensive investigation, partly theoretical and partly computational, of various key points of importance for the evolution of cooperation. In §3 it is assumed that the distribution of players is spatially homogeneous. It is worth noting that even without diffusion new effects are to be expected. This is because the death rate on its own has the result of reducing the number of players which have previously had encounters. Diffusion of course increases this effect. We find that there are various possibilities depending on the parameter values including the stable coexistence of T- and D-players.

Next we consider effects due to spatial inhomogeneities in the population. First, in §4 the question of stable coexistence of time-independent spatially inhomogeneous states, or 'patterns' is studied. We show that through a Turing instability such patterns are produced for certain ranges of parameters. Next, it is shown that under certain circumstances a stable coexistence in a pattern is possible when coexistence is impossible for a spatially constant distribution of players.

Finally, in §5 we consider the question of invasion. It is probably most realistic to assume that in many circumstances a mutation will lead to a spatially limited distribution of new strategists, say TFT. What are the conditions under which this can invade? How do the mobilities of the two species affect the situation? For example it is commonly assumed (see Dugatkin & Wilson, p. 690 (1991) and the critical remarks in Ferriere & Michod 1994), that a low mobility of T-players is good for themselves. However, we find that no such simple conclusion is possible: sometimes it is good for T-players to have an intermediate mobility. However, our results do show that a high mobility of D-players is good for them and bad for T-players. The investigation is partly based on the analysis of travelling waves, but we also present simulations showing that sometimes a very small spatially limited 'blip' of T-players may invade a population of D-players. Indeed the total number of T-players needed to invade may tend to zero as μ_T tends to zero, so an arbitrarily small initial distribution of almost sedentary T-players may invade a population of D-players. The paper concludes with a discussion of results in §6.

2. THE MODEL

Let us first review the classical IPD. This is played between two individuals of similar appearance, each of whom has to choose between the strategies C (cooperate) and D (defect). The pay-offs to the contestants after a single round of the game are shown in table 1. Note that it is always assumed that

$$\gamma > \alpha > \delta > \beta \text{ and } 2\alpha > \beta + \gamma.$$

If the players only meet once then the only rational strategy is for each to defect and so both players receive the pay-off δ. However, the situation changes if the same two players compete repeatedly leading to the IPD. There are now more options available to the players because each may decide to play C or D on a round by any rule (deterministic or probabilistic) of their choice. The rules that are adopted here are D (which now means defect on every round) and T (the tit-for-tat strategy). A T-player always starts by cooperating and then repeats its opponent's last play. Note that a player cannot refuse to play the game even

Table 1. *The matrix of pay-offs for the prisoner's dilemma with the strategies C and D.*

	C	D
C	α	β
D	γ	δ

Table 2. *The pay-off matrix for the* IPD *with the strategies* TFT and D.

	TFT	D
TFT	$\alpha/(1-w)$	$\beta+[w\delta/(1-w)]$
D	$\gamma+[w\delta/(1-w)]$	$\delta/(1-w)$

if it recognizes an uncooperative opponent; this is a standard assumption in the IPD. If the probability of a further round is w then the expected pay-offs to the T-and D-players at the conclusion of their meeting are as given in table 2.

Now consider a population of individuals, each of whom is either a T-player or a D-player. These individuals are able to move around in space and so meet different opponents. At each meeting just one round of the game is played. The rate of encounters will depend upon the population density and so it is necessary to modify the above pay-offs. It is not reasonable to require that each encounter lasts for the same length of time because this would lead to difficulties when the population density was high. It is now assumed that the pay-offs α, β, γ and δ are measured in units of *per capita* growth rate. For example, if D-players meet opponents half of whom play D and half play C, then their *per capita* growth rate will be $(\gamma+\delta)/2$. Note that this pay-off is independent of the number of encounters, for because pay-off depends upon encounters per unit time, it does not matter whether there are few encounters (each lasting a long time) or many encounters (each lasting a short time).

For the strategy T to make sense, T-players must be able to recognize a past adversary. On first encounter, i.e. with an unrecognized opponent, a T-player always cooperates. The pay-offs are not received until after the encounter ends and only then does each player realize what its opponent has been playing. This is exactly the same as for the classical situation of just two players. If a T-player meets someone on a second (or later) occasion then he instantly recognizes that opponent, remembers what they played last time, and plays the same himself (recall that he cannot refuse to play an opponent). Thus the model equations not only depend upon the number densities of T- and D-players (which may vary in space and time), they must also keep track of the types of encounters that have occurred, so that when a T-player meets a D-player the probability of this being a first encounter can be calculated.

In the model developed here, the probability of a further round is not fixed *a priori* (which is the case for the classical two-person IPD) but rather it is part of the solution. This is possible because diffusion, births and deaths are explicitly included in the model equations. The spatial environment in which the population resides is assumed for simplicity to be one-dimensional. When the model equations have been derived there will be no difficulty in extending them to a higher dimensional space and indeed some numerical results for a two-dimensional region are presented. The ultimate equations will take the form of partial differential equations but for the purposes of their derivation space is discretized. Suppose that the line, taken as the x-axis, is divided into contiguous cells or sites of size l. These sites may be thought of as containing social groups, or they might correspond to the home-ranges. They must be sufficiently large that relatedness is not important. Many authors have commented upon the fact that cooperation most easily flourishes when the game is played within a family (or at least a related) group essentially because of kin selection. However we are specifically looking for an explanation for cooperation when individuals are unrelated and kin selection is excluded in our model. Also the cells must contain a sufficiently large number of individuals for probabilistic arguments to be appropriate and for the resulting differential (as opposed to difference) equations to be valid. It will also be necessary to introduce Δt, a time interval in which a fraction m of the individuals in a cell leave that cell. It must be borne in mind that Δt and l are biologically significant parameters, corresponding to generation time and habitat size, and are not mathematical constructs which are later made infinitesimally small. However, Δt must be small compared with the timescale of the problem and l must be small compared with the overall size of the region. In the course of time there will be contests played between many of the pairs of individuals in a cell. Players may also leave the cell by dying or by moving to an adjacent cell and also new players may enter by being born or by migrating from an adjacent cell. As the games are played, the contestants accrue a pay-off which influences their reproductive success. Simple haploid genetics (like begets like) is assumed. A T-player will recognize an opponent on subsequent occasions provided that neither has left the cell where the encounter occurred. Individuals who leave and who later return are treated as first encounters.

It should be emphasized that our model is continuous and the approach will not be valid if the densities are extremely low. There may be additional effects due to very small group size, but these will not be covered in our model. Also there will be a population limitation term in the final equations (see next paragraph) which will forbid very high densities. This automatically removes the possibility of encounter rates being unrealistically high (when there may be too little time for recognition of the opponent).

Suppose that a typical cell contains lu T-players and lv D-players so that u and v are the line densities of the two types. The pay-offs in the game (i.e. α, β, γ, δ) determine the *per capita* growth rate. Thus if the pay-off were ξ to each individual of a group of size z then their numbers would increase at a rate ξz in the absence of all other effects. Let the number of D-players that a typical T-player has already met (and so will recognize next time they meet) be gvl. This T-player can meet any of three types of individual and the fraction of time which will be occupied by each type of meeting is as follows:

T-player	known D-player	unknown D-player
$u/(u+v)$	$gv/(u+v)$	$(1-g)\,v/(u+v)$

The corresponding pay-offs are α, δ and β which imply that the rate of increase of the T-players is:

$$u\{[u\alpha/(u+v)]+[gv\delta/(u+v)]+[(1-g)v\beta/(u+v)]\}$$
$$=\{[u/(u+v)](\alpha u+\beta v)\}+[G(\delta-\beta)/(u+v)]$$

where $G = guv$. Similarly the rate of increase of D-players due to interactions is:

$$v\{[v\delta/(u+v)]+[gu\delta/(u+v)]+[(1-g)\,u\gamma/(u+v)]\}$$
$$=\{[v/(u+v)](\gamma u+\delta v)\}-[G(\gamma-\delta)/(u+v)]$$

because a typical D-player will have already met *gul* of the T-players in its cell and so only receives δ when it plays them again. These growth rates have to be modified by population limitation terms which we take to have the forms $uF(u+v)$ and $vF(u+v)$ for the T- and D-players respectively. In the following discussion it is assumed that logistic growth is appropriate, but other limitation terms would be equally possible. Thus we have:

$$F(U) = b-(d+\sigma U)$$

where b and d are the intrinsic birth and death rates and σ measures the importance of overcrowding. It will be assumed that the pay-offs are always beneficial and contribute positively to the birth rate (rather than negatively to the death rate), and so α, β, γ, δ will all be considered non-negative. Consequently the *per capita* death rate of either type of player is $[d+\sigma(u+v)]$ and is independent of the game. The final effect that has to be considered is migration. As is usual this is modelled by a random walk. Specifically, if m_T is the probability that a T-player moves out of a cell into an adjacent one in time Δt, then the diffusion rate of the T-players is $\mu_T = m_T l^2/(2\Delta t)$. Similarly the diffusion rate of the D-players can be written as $\mu_D = m_D l^2/(2\Delta t)$. The final form of the equations for the growth of the populations is:

$$(\partial u/\partial t) = \{[u/(u+v)](\alpha u+\beta v)\}+\{[G/(u+v)](\delta-\beta)\}$$
$$+[uF(u+v)]+[\mu_T(\partial^2 u/\partial x^2)] \quad (1)$$

$$(\partial v/\partial t) = \{[v/(u+v)](\gamma u+\delta v)\}-\{[G/(u+v)](\gamma-\delta)\}$$
$$+[vF(u+v)]+\mu_D(\partial^2 v/\partial x^2) \quad (2)$$

where

$$F(U) = b-d-\sigma U. \quad (3)$$

The possible variation of G in time and space now has to be considered. Recall that $G = guv$ is the number density of T-D pairs within a cell that have already met. This will be affected by: (i) deaths; (ii) migration; and (iii) encounters. The first two of these will change G but not g and so if Δu, Δv are the changes in u,v due to the combined effects of death and migration in the time interval Δt, then the corresponding change in G is

$$g(u\Delta v+v\Delta u) = G[(\Delta u/u)+(\Delta v/v)]$$
$$= G[2d+2\sigma(u+v)+(m_T/\Delta t)+(m_D/\Delta t)]\,\Delta t.$$

The number of encounters per unit time in a cell is taken to be proportional to the number of pairings within that cell. The encounter rate k will be used for the constant of proportionality. It follows that the number of encounters that each individual has in unit time is proportional to the local value of $u+v$ and also that the change in G due to meetings in the time interval Δt is $(uv-G)\,k\Delta t$ where k is the encounter rate. We thus have

$$(\partial G/\partial t) = kuv-G[k+2d+2\sigma(u+v)+\theta] \quad (4)$$

where

$$\theta = (2/l^2)(\mu_T+\mu_D). \quad (5)$$

The model equations (1–5) have to be supplemented by boundary conditions. We shall always use Neumann conditions for the spatial boundary conditions (which corresponds to the biologically reasonable condition that there is no net flux of players through the boundaries) together with appropriate conditions for u, v and G at $t = 0$. It is clear that if the spatial region is two dimensional only minor changes have to be made to the model equations (apart from interpreting u and v as numbers per unit area) and the general forms of the boundary and initial conditions are unchanged.

If we ignore equation (4) and take g to be a constant then the equations above provide a simple dynamic for the model of Maynard Smith (1982) or Axelrod (1984) for the IPD with g replacing w, the probability that two players have another round of their contest. Thus we have a model for the IPD in which each contest has the same expected number of rounds. For the IPD it is well-known that when the population is composed of T- and D-players the equilibrium point at which the whole population is composed of defectors, all-D, is always stable (i.e. always an evolutionarily stable strategy or ESS) and that the other boundary equilibrium solution, all-T, is also stable provided that

$$w > (\gamma-\alpha)/(\gamma-\delta). \quad (6)$$

For the model presented above, the solution all-D is also always stable (provided that it is viable, i.e. provided that $b+\delta > d$). The conditions for all-T to be stable are the viability condition $b+\alpha > d$ and

$$k(\alpha-\delta)/(\gamma-\alpha) > 2\alpha+2b+\theta. \quad (7)$$

Also, from the equation for G, it is readily found that the value of g at this boundary equilibrium point is

$$\{(2b+k+\alpha+\gamma+\theta)-[(2b+k+\alpha+\gamma+\theta)^2$$
$$-4k(\gamma-\delta)]^{\frac{1}{2}}\}/[2(\gamma-\delta)] \quad (8)$$

and from (8) it can be shown that (7) is equivalent to (6). The critical case of equality is when an interior equilibrium is crossing the u-axis and this causes a change in the stability of the all-T equilibrium point.

In terms of g the contribution of the game to the growth rate of the T-players is

$$[u/(u+v)]\{\alpha u+[\beta(1-g)+g\delta]\,v\} \quad (9)$$

and for the D-players it is

$$[v/(u+v)]\{[\gamma(1-g)+\delta g]\,u+\delta v\}. \quad (10)$$

These expressions show that g plays the same role as w in the IPD. In particular the expected number of rounds in a contest between D- and T-players (conditional upon the players having met) is $1/(1-g)$ at an equilibrium. Although g is the biologically more important variable, it is mathematically more convenient to use G in the final equations.

The population limitation term assumed (equation 3) corresponds to logistic growth and is in the spirit of Cressman & Dash (1987) and Cressman (1992) who propose a modification of the classical Taylor-Jonker replicator equations. The model actually used here appears superior to the classical IPD with dynamics because the mobility, population density and mortality determine the expected length of game. Thus the model meets a criticism of the usual IPD made by Houston (1993). We discuss finally comparable models which involve a spatial factor together with the restriction that players only have encounters with their neighbours.

Cellular automata models which are discrete in space or time have appeared fairly frequently in the literature. One of the most recent is that due to Nowak & May (1992,1993), see also comments in Sigmund (1992) and Huberman & Glance (1993). This model does not incorporate any memory for the players and is thus not an iterated game. It is interesting in that it nonetheless predicts coexistence of cooperators and defectors. The model is based on two assumptions crucially different from ours and a direct comparison is impossible: (i) it is a 'one-round' game; and (ii) it assumes discrete spatial patches with one individual per cell and discrete time (non-overlapping generations). Concerning the second assumption, it may be argued that all models which are discrete in space suffer from being sensitive to the updating procedure assumed and of course discrete generations also encourage curious effects. Here we have overlapping generations and continuous space so that possibly artificial effects due to discretization are completely avoided.

A model for cooperation which is broadly in the same spirit as ours is that described in an interesting paper by Dugatkin & Wilson (1991). This model is stochastic and only defectors are assumed to migrate. Thus the technicalities differ quite considerably from those in the present paper.

An approach which is fairly similar to that presented here is that due to Ferriere & Michod (1994). Their model is simpler and they are able to utilize known results (Hutson & Vickers 1992) concerning the solution of the governing partial differential equations to predict the direction of travelling waves. However, they assume that their equivalent to g is constant in space and time. Thus their model cannot predict the occurrence of patterns. It does incorporate the important feature that players will contest with neighbours, and its relative simplicity is also a very attractive feature.

3. SPATIALLY HOMOGENEOUS PLAYER DENSITIES

The principal aim in deriving the model in §2 is to consider the effect of spatially inhomogeneous distributions of players. However, we shall see that even if the distributions are homogeneous, there are significant differences in the results from those for standard models. Initially this may appear surprising, but on re-examination there are good reasons for the differences. In the case of our model: first, players only play with their near neighbours and secondly they may leave the neighbourhood or die. Both of these factors cause the memory of previous encounters to be lost and have a crucial effect on the outcome.

When the densities u and v of the T- and D-players and the getting-to-know function g depend only on t (but not on x), the system shown by equations 1–5 reduces to the following system of ordinary differential equations.

$$\mathrm{d}u/\mathrm{d}t = \{[u/(u+v)](\alpha u+\beta v)\} + \{[G/(u+v)](\delta-\beta)\} + u[b-d-\sigma(u+v)] \quad (11)$$

$$\mathrm{d}v/\mathrm{d}t = \{[v/(u+v)](\gamma u+\delta v)\} - \{[G/(u+v)](\gamma-\delta)\} + v[b-d-\sigma(u+v)] \quad (12)$$

$$\mathrm{d}G/\mathrm{d}t = kuv - G[k+2d+2\sigma(u+v)+\theta], \quad (13)$$

where

$$\theta = (2/l^2)(\mu_\mathrm{T}+\mu_\mathrm{D}). \quad (14)$$

Of course, although spatial homogeneity holds, diffusion still has a significant role, entering the above system through the parameter θ. Furthermore, because of the existence of a death rate, even when $\theta = 0$ the model does not reduce to a standard IPD model because of the variation in g.

Let us start by noting that with a standard model, that is with $g = G/(uv)$ constant (of either replicator or Cressman type), the game will have one of two simple structures. In one of these there are stable equilibria with only T-players and only D-players (and an unstable coexistence equilibrium), and in the other there is a stable all-D state and the only other equilibrium, that with just T-players, is unstable. Consider first cases with $\mu_\mathrm{T} = \mu_\mathrm{D} = 0$. Then, in addition to the possibilities just described, there is a range of parameters leading to a stable coexistence of cooperators and defectors with two coexistence states (one stable and one unstable) as well as a stable state with only D-players and an unstable state with only T-players. The phenomenon of having stable coexistence of the two types of players also occurs in the model of Dugatkin & Wilson (1991). Of course exactly the same structure is possible with μ_T and μ_D positive. Typical cases are illustrated in figure 1.

A question of particular interest is how a mutation of TFT in a population of D-players can lead to a stable state where some or all players are TFT's. The difficulty is that a mutation is likely to cause only a small number

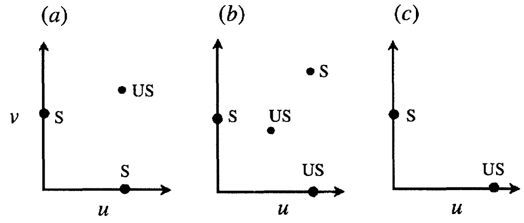

Figure 1. S and US denote stable and unstable equilibria respectively. The sequence (*a*), (*b*), (*c*) corresponds to decreasing α which causes T-players to be less 'fit'. In (*a*) all-T is stable, in (*b*) all-T is not stable, but stable coexistence is still possible and in (*c*) there is no stable state with T-players.

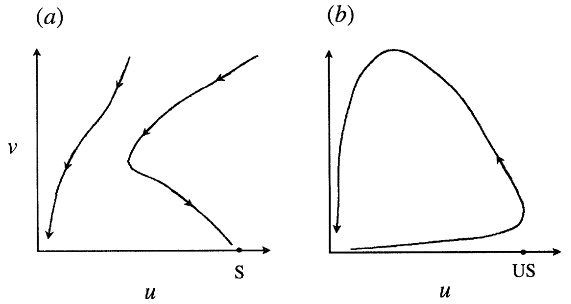

Figure 2. Where defectors cannot live on their own: (*a*) the all-T equilibrium is stable; (*b*) the all-T equilibrium is unstable and T-players are vulnerable to invasion by D-players with the eventual result that both populations are wiped out.

of new types to appear and it is difficult to see how these may increase. However, in the present model for certain parameter ranges the basin of attraction of the all-D equilibrium is extremely small, so a small number of T-players may invade. We shall return to this question in §5 when considering spatially inhomogeneous effects as it is probably more realistic to suppose that a mutation is restricted to a limited region of space and ask how it evolves.

A final possibility is of some interest. For certain parameter ranges D-players cannot exist on their own but T-players can. If this point is stable the final state may be either extinction of both types, or TFT only, depending on the initial populations (see figure 2*a*). On the other hand, the all-T equilibrium may be unstable. In this case it may happen that if a small number of D-players is introduced then their number will increase, but these will have a deleterious effect on the T-players and both populations will become extinct (see figure 2*b*). The well-being of TFT is fragile in this case. Another possibility is that there may be a stable interior equilibrium; biologically this corresponds to D-players existing as parasites on T-players, whilst being unable to live on their own.

In conclusion, it is important to note that a consideration of spatially homogeneous populations alone may lead to misleading results. The dangers are illustrated by two specific examples considered in detail in §4. In the first, a stable equilibrium becomes unstable when spatial inhomogeneity is introduced, a pattern being produced. In the second there is a pattern when coexistence is impossible for homogeneous densities. It is therefore dangerous to ignore spatial inhomogeneity. It appears that the IPD should be tackled via a model in which spatial inhomogeneity is taken into account. We next turn to these issues.

4. PATTERNS

It is shown in this section that the model equations 1–5 can have stable, stationary, spatially non-homogeneous solutions (i.e. patterns), and this will be illustrated later when numerical solutions are shown for a two-dimensional region. In the first instance a standard linearization argument is presented to show that a stable solution of the spatially homogeneous equations 11–14 may be unstable for the full equations provided that the diffusion rates are suitably chosen. In other words these equilibria are sometimes susceptible to Turing instabilities.

The behaviour of these patterns has been investigated numerically when some parameter of the problem is varied. The most remarkable finding is that it is possible for the equilibrium point (which originally gave rise to the pattern via a Turing instability) to disappear and yet the pattern remain: this is reminiscent of the smile of the Cheshire cat. When the equilibrium point has disappeared there is a spatial pattern which has no counterpart in the equations 11–14. Indeed the only equilibrium solutions that these latter equations then have correspond to the population being composed entirely of either T-players or D-players. There is not even an unstable equilibrium solution implying co-existence and yet the equations with spatial terms have a stable solution in which both types of player are present. This phenomenon is further commented upon in the final section.

The standard technique to find a Turing bifurcation, see for example Murray (1989), for a system of equations of the form:

$$(\partial u_i/\partial t) = [f_i(\boldsymbol{u}) + (\mu_i \Delta u_i)]$$
$$(1 \leqslant i \leqslant n,\ \boldsymbol{u}\in\mathbb{R}^n,\ \boldsymbol{x}\in\Omega\subset\mathbb{R}^2) \quad (15)$$

which have an equilibrium solution $\boldsymbol{u} = \bar{u}$, is to assume a solution of the form

$$u_i = [\bar{u}_i + (c_i \phi e^{\omega t})] \quad (1 \leqslant i \leqslant n). \quad (16)$$

Here ϕ is an eigenfunction for the Laplacian on Ω, i.e.

$$\Delta\phi = -\rho\phi\ (\boldsymbol{x}\in\Omega),\ (\partial\phi/\partial n) = 0\ (\boldsymbol{x}\in\partial\Omega) \quad (17)$$

for some eigenvalue ρ. For example, if $\Omega = [0, 1]$ then $\rho = m^2\pi^2$ for some integer m. If the solution 16 is substituted into the equations 15 and a linearization in the c_i's carried out, it is readily found that ω must satisfy an eigenvalue equation of degree n. For the model equations under consideration, these equations are algebraically complicated and have to be solved together with the equilibrium conditions which are also non-trivial. No striking simplifications appear possible, but the equations are quite easy to handle numerically. For any given set of parameters (α, β, γ, δ, σ, b, d, k, μ_T, μ_D, θ) it is necessary to compute the interior equilibrium point(s) (of course, there may not be any), determine their stability (by checking the real parts of ω when μ_T and μ_D are both zero) and then, if there is a stable equilibrium point, determine whether or not it can be made unstable by a suitable choice of μ_T and μ_D.

Consider first the behaviour of the system of ordinary differential equations 11–14. Figure 3 indicates one of the possible modes of behaviour of the equilibrium points in the u-v plane as α is varied. For small values of α there are only the two boundary equilibria, all-T (which moves along the u-axis) and all-D (the point P, which is independent of α). The first is unstable and the second one stable. As α is increased, a pair of equilibria appear first at A. One of these is unstable and moves along the branch AB as α increases. It must be remembered that α is not an entirely free parameter.

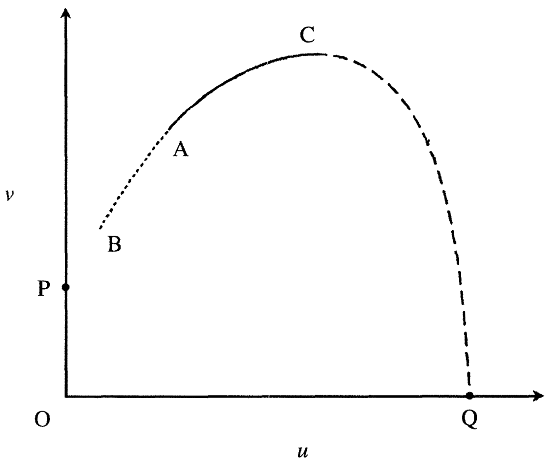

Figure 3. Typically movement of equilibrium points in the u-v plane as α is increased. The dotted portion of the curve denotes unstable points, the solid portion denotes those susceptible to a Turing instability, and the dashed portion denotes unconditionally stable points.

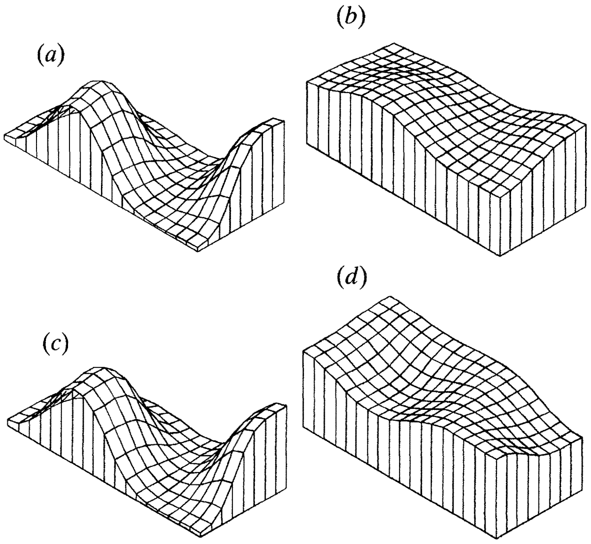

Figure 4. Pattern solutions when there is no internal equilibrium solutions: $\alpha = 12$, $\beta = 0$, $\gamma = 16$, $\delta = 1$, $\sigma = 1$, $\theta = 0.5$, $b = 1$, $d = 1$, $\mu_T = 0.02$, $\mu_D = 0.3$. (*a*), (*b*), (*c*) and (*d*) show the distribution of u, v, G and g in space, their maximum values being 2.1, 2.9, 1.5 and 0.39, respectively.

It is the constraint that α has to be less than γ which here prevents the unstable equilibrium point from reaching O. The other equilibrium point, which is stable, moves along the curve AC. As α is increased further, this point moves towards the u-axis and reaches it at Q. When it does so the boundary equilibrium point all-T undergoes a bifurcation and is now stable for all larger values of α.

Consider now the same parameter set but this time with the system of partial differential equations 1–5. When the equilibrium point is on the section AC of the curve in figure 3, it is capable of undergoing a Turing instability if μ_T and μ_D are suitably chosen. At C, which almost corresponds to a maximum in v, the equilibrium point becomes unconditionally stable, i.e. Turing instabilities are no longer possible. Let α be reduced and track a solution of the model equations rather than just the homogeneous equilibrium points in the u–v plane. For large α there is just a stable, spatially constant solution corresponding to a point on the u-axis (together with the ever-present stable equilibrium point P). As α is reduced, a stable (but still spatially uniform) solution with both T- and D-players develops. This is the branch QC. But, providing the diffusion rates have been suitably chosen, a pattern is formed somewhere along the branch C to A. This non-homogeneous, stable solution can be followed as α decreases towards its value corresponding to the point A. But by the time that α has this critical value, the pattern solution is sufficiently separated from the equilibrium point that its existence continues on after this point has disappeared (by merging with the unstable equilibrium solution that has been moving along BA). The existence of a pattern in a situation where there is not even an interior equilibrium point is unusual. Figure 4 illustrates, for a two-dimensional rectangular region with aspect ratio 2:1, such a pattern.

The above gives an overview of the essential features. We now present some further analysis which describes some of the more complex behaviour of the model. Figure 5 illustrates the nature of the equilibrium points in the (α,δ)-plane, all other parameters having been fixed. The pair (α,δ) were chosen because of their particular importance; α being the pay-off to two cooperators, δ being the pay-off to two defectors. The permitted region, which is defined by

$$\gamma > \alpha > \max[\delta, (\beta+\gamma)/2],$$

is shown in figure 5 bounded by solid lines. For the parameter set shown in the figure the following lists the three events which occur as α and δ are varied.

1. An interior equilibrium point may appear out of the all-T equilibrium point $(\bar{u}, 0)$. When this happens the stability of the all-T point is changed. This is shown as a dot-dash line in the figure.
2. Two interior roots may appear (originally as a double root). This is shown as a line of short dashes. One of the roots is always stable and the other unstable.
3. An interior, stable equilibrium point may cease to be unconditionally stable and instead become susceptible to a Turing instability. This event is bounded by the line of long dashes.

For other parameter sets (i.e. values of α, β, γ, δ, σ, b, d, k, μ_T, μ_D, θ) there are two further possibilities.

4. An all-T or all-D equilibrium point may appear from the origin.
5. An interior equilibrium point may appear out of the origin.

The lines corresponding to the events 1,2 and 3 above, divide the permitted region of the α-δ plane into four subregions which have been coded according to the types of equilibrium points present on the u-axis, on the v-axis and in the interior. We use s and u to denote stability and instability, respectively. Thus the code $u/s/us$ implies that there is an unstable point on the u-axis (i.e. all-T), a stable point on the v-axis (all-D) and two interior points, one of which is stable and the other unstable. The letter t indicates the presence of a stable point (for the system 11–13) which may undergo a

Turing instability if the diffusion rates are suitably chosen. Also a dash in the code indicates that there is no equilibrium point in that region. It is possible to generate very complicated pictures of sub-regions. Certainly there may be as many as 11 in one diagram (without any parameter value exceeding 2). The dotted line in figure 5 indicates how, for the situation described in figure 3, the behaviour changes with varying α. The points labelled Q, C, A and B correspond to the respectively labelled points in the earlier figure. However, the scenario may be more involved than that shown in figure 3. For example it may happen that as α is increased the representative point passes through sub-regions which are coded as follows:

u/-/-
internal double root appears
u/-/*ut*
one root enters the origin
u/-/*t*
interior point ceases to be of Turing type
u/-/*s*
interior point crosses the *u*-axis
s/-/-.

Although this is an extreme case, subsequences of the above serve to show how the picture may change as the pay-offs are varied. Although the effect of increasing α (the cooperators pay-off) is very complicated, the extreme values of α give easily interpreted results. Small α makes it difficult for the T-players to exist (typically an unstable point on the *u*-axis) whereas large α gives a stable population of T-players. What is intriguing is the catalogue of intermediate possibilities as internal equilibrium points appear (and disappear) and whose stability change. They may be regarded as transitional situations in which T- and D-players are truly struggling for domination. However, some of the regions of parameter space occupied by these situations are large and it may not be possible, because of the nature of the game, to evolve out of them (or at least not by any permissible change in the pay-offs). Thus it is not possible to rule out any of them.

In interpreting the biological implications of these results, an important consideration is the strength of the fitness effects of playing the IPD relative to the background population pressures. A reasonable measure of the relative size of these two effects is provided by the ratio of the equilibrium population sizes of all-T and all-D individuals, i.e. by $\bar{u}/\bar{v}$. If this ratio is large the game is important, whereas if it is near unity it is the underlying population dynamics which are dominant and then most of the interesting effects do not occur. It could be argued that defectors effectively 'fight' one another, whereas cooperators assist each other (as in several of the situations discussed by Axelrod (1984)). A quite large value of $\bar{u}/\bar{v}$ may be appropriate in such circumstances. On the other hand, a mutation (perhaps of a D-player to a T-player) might lead to a more modest ratio, say less than 3.

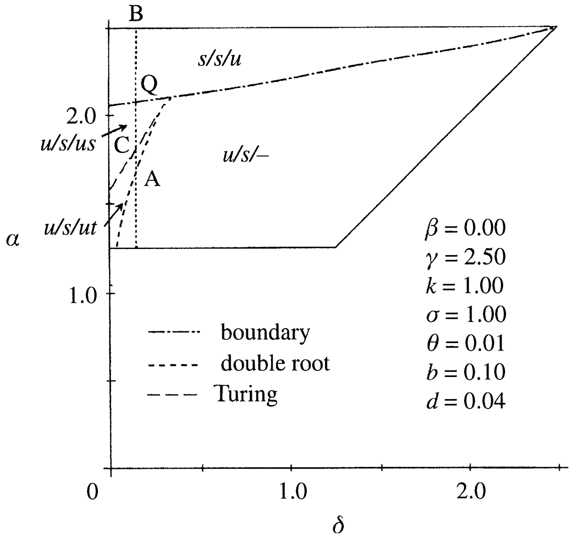

Figure 5. The behaviour of equilibrium points in the (α, δ) plane showing where double roots and Turing instabilities occur and also where an interior point crosses the boundary. The dotted line shows the effect of changing α; the points B, Q, C and A correspond with those in figure 3.

As mentioned earlier, a pattern formed at a Turing instability may be tracked beyond the existence of the equilibrium point. For example, using the parameter set of figure 5 and $\alpha = 1.5$, $\mu_T = 0.005$, $\mu_D = 0.03$ a pattern was traced from $\delta = 0.075$ to $\delta = 0.25$. The equilibrium point disappears at $\delta \simeq 0.13$. For these parameters the game dominates over the birth and death terms. The ratio $\bar{u}/\bar{v}$ is around 6 for the situation just described. However it is possible to make this ratio arbitrarily close to 2 and still get a Turing instability (although rather extreme values of the diffusion rates may be needed for it to manifest itself). It is certainly easier to get a Turing bifurcation when the fitness effects of the game dominate the population terms, but such an eventuality is not prohibited by modest values of $\bar{u}/\bar{v}$. Generally speaking, increasing δ (the defectors pay-off) has rather simpler consequences than increasing α. Interior equilibrium points tend to lose their stability, via a Turing bifurcation, and then coalesce (with an unstable point) so that the final result is all-D as the only stable equilibrium point. Of course the limiting value of δ is α at which point the defectors are getting just the same pay-offs as the cooperators (when each plays with one of its own type) and so the cooperators are never doing better.

The most important point is that the inclusion of space has produced a phenomenon which would not have been guessed if the spatially homogeneous equations 11–13 alone had been studied.

5. INVASION BY T-PLAYERS

From the point of view of the evolution of cooperation, one of the most difficult problems is to understand how, in a population of D-players a small proportion of T-players (produced by mutation in the same species or by invasion of a new type) can increase and lead to a situation where only T-players exist, or perhaps where there is a stable coexistence of T and D-players. We examine this question here and point out the crucial role that spatial inhomogeneities in distribution play. A second key point is to discover the role that the

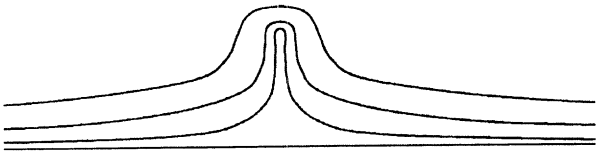

Figure 6. Diagrammatic representation of spread of T-players with an initially small spatially restricted distribution.

mobilities (for the T and D-strategists) play in determining which strategy is the more robust, and this is the second main area of investigation in this section.

We shall, for simplicity, restrict ourselves in this section to the case where there are exactly two stable equilibria, one with D-players alone and one with T-players alone, and an unstable coexistence equilibrium. (Closely analogous arguments hold when there is a stable coexistence equilibrium instead of the all-T equilibrium.) When only spatially homogeneous distributions are considered, the situation is relatively straightforward. The phase space is split into two regions, the 'basins of attraction' of the equilibria. However, the situation is a great deal more complicated for spatially inhomogeneous distributions. For the basins are now sets of functions, the initial densities, rather than sets of points, and the 'shape' of the functions will be an important factor. In addition the diffusion coefficients μ_T and μ_D will play an equally important role. These considerations have led to the idea of the 'dominance' of equilibria (see Fife 1979). Roughly, the all-T equilibrium for example is said to be dominant if a localized distribution of T-players in a population of D-players grows and wipes out the D's. A definition which is probably equivalent (although no strict proof of this exists at the moment) is that the travelling wave (see below for the definition) of invasion will go towards the all-T equilibrium. A theoretical treatment presents rather formidable difficulties, and in this investigation we shall restrict ourselves to a computational study, with the intention of returning to the problem from a more theoretical point of view later.

We consider a spatial domain $[-1,1]$ and suppose that in a homogeneous population of D-players of density $\bar{v}$ a small, spatially restricted group of T-players appears. We assume that the initial values u_0, v_0 satisfy $u_0(x)+v_0(x) = \bar{v}$ and

$$u_0(x) = 0.75 \exp(-\rho x^2).$$

Define

$$F = (1/2\bar{v})\int_{-1}^{1} u_0(x)\mathrm{d}x.$$

F is thus the total initial population of invading T-players expressed as a fraction of the total population, $2\bar{v}$, of D-players. Suppose first that $\mu_T = \mu_D$. If ρ is not too large (i.e. F is not too small), the T-players will invade and take over completely as schematically shown in figure 6: that is, the all-T equilibrium is dominant.

Table 3. *The fraction of T-players required to successfully invade a population.*

(Parameters are and $\alpha = 7$, $\beta = 0$, $\gamma = 9$, $\delta = 1$, $k = 12$, $\sigma = 1$, $b = 3.25$, $d = 1$, $\mu_D = 0.001$ and $\theta = 250(\mu_T+\mu_D)$. The equilibrium values of T- and D-players on their own are 9.25 and 3.25 respectively. F is the minimum fraction total initial T-players / total initial D-players needed for invasion. There is strong evidence that $F \to 0$ as $\mu_T \to 0$, so for a sedentary population ($\mu_T = 0$), an arbitrarily small distribution of T-players may invade.)

μ_T	F
0.004	0.093
0.002	0.052
0.001	0.033
0.0005	0.022
0.00025	0.016
0.000125	0.011

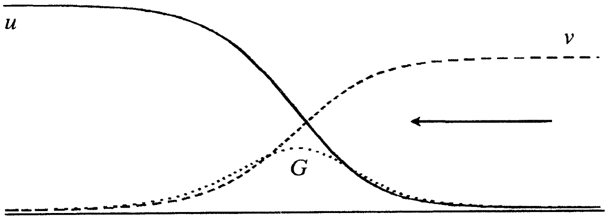

Figure 7. A typical travelling wave. With $c < 0$ as shown, the D-equilibrium is dominant.

Suppose next that μ_D is fixed but μ_T is made very small. Then computations show that (assuming that the peak of the distribution, 0.75 in this example, is not too small) the total number of T-players needed for invasion also becomes small and in fact tends to zero as μ_T tends to zero, see table 3. This is a remarkable result and shows that an arbitrarily small initial distribution of almost sedentary T-players can invade a population of D-players. This appears to provide a partial answer to the paradox of how T-players may 'get off the ground'.

A consideration of the direction of travelling waves shows that unless $(\mu_T+\mu_D)$, and so θ, is fairly large, the all-T equilibrium is dominant and a small distribution of T-players will always invade. However, for large enough θ (corresponding to high mobility of at least one type) the all-D equilibrium becomes dominant.

Of course the situation is not always so favourable for T-players. For different interaction parameters the effect of the diffusion coefficients can be rather different. It is conventional (see the discussion in Dugatkin & Wilson 1991 for example) to expect that increase of mobility, that is of μ_T or μ_D or both, will be 'bad' for T-players and 'good' for D-players, for mobility reduces the number of T–T interactions. This fits with the above example. However, it turns out that the full picture is a great deal more complex. We do not attempt a complete resolution of the problem here, but give a further example which in conjunction with the previous example suggests the complexity of the problem.

We start this example from the point of view of waves of invasion. That is we discuss the speed of solutions of the form $(u(x-ct), v(x-ct), G(x-ct))$, known as travelling waves. Figure 7 represents a typical 'photograph' of the wave at a given time. The sign of c decides which of the equilibria is dominant.

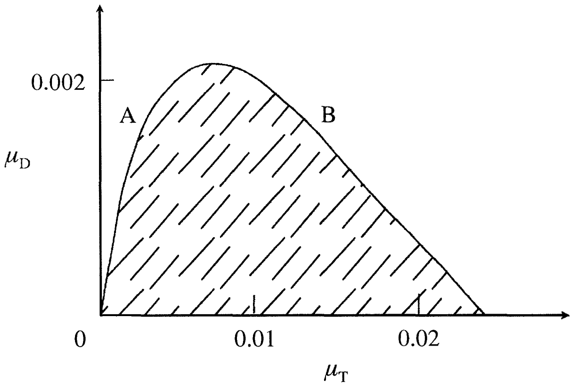

Figure 8. $\alpha = 1.3725$, $\beta = 0.9$, $\gamma = 1.45$, $\delta = 1$, $k = 4$, $b = 0 = d$, $\sigma = 1$, $\theta = 100(\mu_T + \mu_D)$. The hatched region is $c > 0$ when all-T is dominant and in the rest of the region all-D is dominant. Notice that the effect of increasing μ_T is first good for T-players but eventually bad for them.

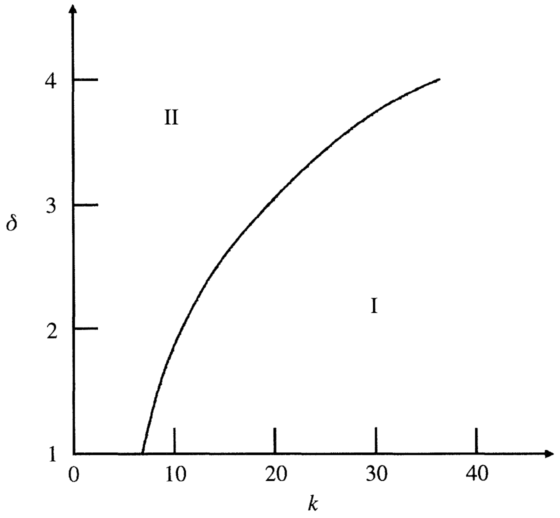

Figure 9. Parameters $\alpha = 7$, $\beta = 0$, $\gamma = 9$, $\sigma = 1$, $b = 3.25$, $d = 1$, $\theta = 0$ (corresponding to very small μ_T and μ_D). In region I a decrease in μ_T favours T-players with invasion possible from very small populations. In region II, small μ_T is bad for T-players and leads to the situation described in figure 8.

If all parameters except μ_T and μ_D (and θ) are fixed, the situation is represented by the graph in figure 8. Notice that for small or large μ_T the defectors win whereas for medium μ_T it is the T-players that win. There is a double reversal of the wave direction. This is not consistent with the simple view that to be sedentary favours T-players; a measure of mobility can favour T-players. In order to interpret this result, consider what happens as μ_T is reduced at each of the points A and B marked on figure 8. These points are where there is a stalemate, i.e. there is a travelling wave with zero speed. When μ_T is reduced, the wave front of the T-players steepens so that u is reduced (from a small value) in the leading edge of the wave. The key factor is that at A the number of encounters between any given two players is determined by the death rate while at B it is by mobility. At B, G and g increase significantly, causing the *per capita* birth rate in u to increase. Thus the T-players have the advantage and a wave develops in which the T-players advance. By contrast, at A there is little change in G and g so that the per capita birth rate of u is mainly influenced by the decrease in u (see equation 9), and this causes the D-players to have the advantage. The same type of argument suggests that reducing μ_D is always bad for D-players and this is consistent with figure 8.

It may be noted that the implications of the results in the two cases treated above differ in one crucial respect. This is that if μ_T becomes very small, in the first case T-players are favoured, whereas in the second case it is bad for these players.

For ease of description we shall use Cases I and II to denote the situations described above. Thus Case I will be taken to mean that decreasing μ_T favours the invasion by T-players (as described in table 3), whereas Case II means that we are in the situation described in figure 8, where small μ_T is disadvantageous for T-players.

The distinction between these two cases clearly has important biological implications. It is an extremely difficult mathematical problem to give an explicit criterion in terms of the parameters for the switchover between the two cases. However, in view of the importance of the distinction, we present some further computations which illustrate the influence of two key parameters δ, k on this switchover. The switchover point is calculated for very small μ_T and μ_D (so that effectively $\theta = 0$). The results are given in figure 9 where the regions are correspondingly marked.

Some remarks on the biological interpretation are as follows. The first parameter used is δ, which governs the population level of D-players on their own. The second k is the encounter rate; clearly large k tends to favour T-players (because many of their encounters with D-players will be second encounters). The expected effects of changing these parameters are well reflected in the diagram. An increase in δ will be bad for T-players in the sense that they will be unable to invade if their initial population is small. However, this effect will be compensated for by increasing the encounter rate k. For example, when δ is 3, the switchover occurs when k is about 19; for k above this value T-players may invade from rarity.

As mentioned in the previous section, a measure of the importance of the game versus the background population pressure is provided by the ratio $\bar{u}/\bar{v}$. For the largest value of δ in figure 9, $\delta = 4$, this ratio is about 1.5, which is probably quite modest in a biological context. One may also remark that the values of k do not seem excessive. Rather than attempting to interpret the size of k directly, it is more illuminating to consider its effect upon g. Recall that g represents the proportion of pairs of T- and D-players which have met. In the case just discussed g is about 0.7 on average over the spatial region. This means that the expected number of encounters between any particular pair of T- and D-players (conditional upon them having met) is about 3. This appears to be a very reasonable figure in many situations where the IPD may be visualized to operate.

Finally, the results in this section partly confirm and partly differ from those of Dugatkin & Wilson (1991) and Ferriere & Wilson (1994). We find that an increase in μ_D is always good for D-players (which agrees with the first but not the second of these references) whereas increasing μ_T may in certain circumstances favour T-players (which agrees with the second but not the first). It is probable that the differences are due to an incorporation of the function g in our model.

6. DISCUSSION

The model presented here has two important novel features: (i) the explicit inclusion of diffusion; and (ii) the incorporation of a 'getting-to-know' field variable, g. This function keeps track of players' encounters and depends on position and time; in our view the incorporation of such a function in the governing equations is essential for a proper understanding of the model. The combination of these two aspects allows us to calculate explicitly the expected number of rounds in each encounter viz $1/(1-g)$ for meetings between a T- and a D-player. Thus the model is self-contained in that this number is part of the solution rather than part of the specification. It is also important to include diffusion explicitly because the relative mobility of the two types of players (tit-for-tat and always-defect) has a crucial effect upon the outcome, partly by limiting the number of rounds to an encounter and partly by allowing players to move to new regions of space.

The behaviour of the model has several significant aspects and they are briefly commented upon below.

1. It is possible to have stable coexistence of T- and D-players with spatially homogeneous densities.
2. A Turing bifurcation of this stable equilibrium may occur leading to a spatial pattern. Indeed there may be a pattern when there is no homogeneous coexistence state.
3. Invasion of a population of D-players by a very small number of T-players is possible, especially if the mobility of the T-players is low.
4. As μ_T, the diffusion rate of the T-players, is increased the direction of the travelling wave connecting the equilibrium points all-T and all-D may undergo two-sign changes.

The above phenomena demonstrate that the present model has a remarkable richness of behaviour. The stable coexistence referred to in point 1. does not seem to be possible if there is also a stable equilibrium point corresponding to all-T. There is always – when viable – a stable equilibrium point where all the players are defectors. That sometimes the coexistence point may be susceptible to a Turing bifurcation as one of the parameters is varied is not particularly remarkable (in so far as any Turing instability is unremarkable) but that it may persist after the equilibrium point that gave birth to it has disappeared is quite curious. It implies that it is all but impossible to state with any confidence the range of behaviour that this type of model may exhibit if only the homogeneous (spatially independent) equations have been analyzed. There are occasions when the inclusion of space introduces entirely new and unsuspected behaviour.

Consider finally the fate of potential invaders. Suppose first that these appear as a small, spatially limited clump. If there are two stable equilibria, all-D and all-T, then it is not surprising that, if the basin of attraction of all-D is small, it can be invaded by a small number of T-players. But note that their maximum density has to be sufficiently high even if their total numbers are low. What does not seem to have been remarked upon previously is that if the invaders have a very low diffusion rate then the total number can be made arbitrarily small, although again the maximum density must be high enough to overcome the invasion barrier of the residents.

The direction of travelling waves (broadly equivalent to the success of potential invaders) depends in a very complex way upon the diffusion coefficients, the pay-offs and the population limitation term. The intriguing question is, what determines their direction? The present model is too involved for us to separate out the relative importance of the various factors. Suffice it to say that we have performed other numerical experiments on simpler models (not involving IPD) and it is clear that the complexities, although compounded here, are certainly not due to any odd feature of this model. This is a fascinating area of study and the authors will, in subsequent publications, investigate it further. The most noteworthy point here is that large or small diffusion cannot unambiguously be claimed to be good for the strategy tit-for-tat, rather it depends upon the entire game and the mobility of the other contestants. Remarkably, in some circumstances an increase of μ_T is advantageous for T-players when μ_T is small, but too large an increase is then to their disadvantage.

We are grateful to R. Ferriere and R. Michod for a preprint of their very interesting paper and also to K. Sigmund, R. Ferriere and R. Michod for most helpful discussions in connection with this investigation. We are also grateful to the referees for many constructive suggestions which have led to an improvement in the paper.

REFERENCES

Axelrod, R. 1981 The emergence of cooperation among egoists. *Am. Polit. Sci. Rev.* **75**, 306–318.

Axelrod, R. 1984 *The evolution of cooperation.* New York: Basic.

Axelrod, R. & Hamilton, W. 1981 The evolution of cooperation. *Science, Wash.* **211**, 1290–1296.

Brown, J., Sanderson, M. & Michod, R. 1982 Evolution of social behaviour by reciprocation. *J. theor. Biol.* **99**, 310–339.

Cressman, R. 1992 *The stability concept of evolutionary game theory.* Berlin: Springer-Verlag.

Cressman, R. & Dash, A. 1987 Density dependence and evolutionarily stable strategies. *J. theor. Biol.* **126**, 393–406.

Dawkins, R. 1989 *The selfish gene.* Oxford University Press.

Dugatkin, L. & Wilson, D. 1991 Rover: a strategy for exploiting cooperators in a patchy environment. *Am. Nat.* **138**, 687–701.

Enquist, M. & Leimar, O. 1991 The evolution of cooperation in mobile organisms. *Anim. Behav.* **45**, 747–757.

Eshel, I. & Cavalli-Sforza, L. 1982 Assortment of encounters and evolution of cooperativeness. *Proc. natn. Acad. Sci. U.S.A.* **79**, 1331–1335.

Feldman, M. & Thomas, E. 1987 Behaviour-dependent contexts for repeated plays of the prisoner's dilemma II: dynamical aspects of the evolution of cooperation. *J. theor. Biol.* **128**, 297–315.

Ferriere, R. & Michod, R. 1994 The emergence of cooperation among mobile organisms. (In the press.)

Fife, P. 1979 *Mathematical aspects of reacting and diffusing systems*. Berlin: Springer–Verlag.

Houston, A. 1993 Mobility limits cooperation. *Trends Ecol. Evol.* **8**, 194–196.

Huberman, B. & Glance, N. 1993 Evolutionary games and computer simulations. *Proc. natn. Acad. Sci. U.S.A.* **90**, 7716–7718.

Hutson, V. & Vickers, G. 1992 Travelling waves and dominance of ESS's. *J. math. Biol.* **30**, 457–471.

Maynard Smith, J. 1982 *Evolution and the theory of games*. Cambridge University Press.

Mesterton-Gibbons, M. & Dugatkin, L. 1992 Cooperation among unrelated individuals: evolutionary factors. *Quart. Rev. Biol.* **67**, 267–281.

Mesterton-Gibbons, M. 1992 On the iterated prisoner's dilemma in a finite population. *Bull. Math. Biol.* **54**, 423–443.

Murray, J. 1989 *Mathematical biology*. Berlin: Springer-Verlag.

Nowak, M. A. & May, R. 1992 Evolutionary games and spatial chaos. *Nature, Lond.* **359**, 826–829.

Nowak, M. A. & May, R. 1993 The spatial dilemmas of evolution. *Int. J. Bifurc. Chaos* **3**, 35–78.

Nowak, M. A. & Sigmund, K. 1990 The evolution of stochastic strategies in iterated games. *Acta Applic. Math.* **20**, 247–265.

Nowak, M. A. & Sigmund, K. 1994 The alternating prisoner's dilemma. *J. theor. Biol.* **168**, 219–226.

Nowak, M. A., Sigmund, K. & El-Sedy E. 1994 Automata, repeated games and noise. *J. math. Biol.* (In the press.)

Sigmund, K. 1992 On prisoners and cells. *Nature, Lond.* **359**, 774.

Trivers, R. 1971 The evolution of reciprocal altruism. *Quart Rev. Biol.* **46**, 35–57.

Received 3 June 1994; accepted 8 November 1994