

Spatial games and the maintenance of cooperation

MARTIN A. NOWAK, SEBASTIAN BONHOEFFER, AND ROBERT M. MAY

Department of Zoology, University of Oxford, South Parks Road, OX1 3PS, Oxford, United Kingdom

Contributed by Robert M. May, January 28, 1994

ABSTRACT The Prisoner's Dilemma (PD) is a widely employed metaphor for problems associated with the evolution of cooperative behavior. We have recently proposed an alternative approach to the PD, by exploring "spatial games" in which players—who are either pure cooperators, C, or pure defectors, D—interact with neighbors in some spatial array; in each generation, players add up the scores from all encounters, and in the next generation a given cell is retained by its previous owner or taken over by a neighbor, depending on who has the largest score. Over a wide range of the relevant parameters, we find that C and D persist together indefinitely (without any of the complex strategies that remember past encounters, and anticipate future ones, which characterize essentially all previous work on the iterated PD). Our earlier work, however, dealt with symmetric spatial lattices in two dimensions, deterministic winning and discrete time. We show here that the essential results remain valid in more realistic situations where the spatial distributions of cells are random in two or three dimensions, and where winning is partly probabilistic (rather than being determined by the largest local total). The essential results also remain valid [*pace* Huberman and Glance [Huberman, B. A. & Glance, N. S. (1993) *Proc. Natl. Acad. Sci. USA* 90, 7716–7718}] when interactions occur in continuous rather than discrete time.

One of the enduring problems of evolutionary theory, from Darwin's time to today, is the evolution of altruism: cooperative behavior may benefit the group, but it is inherently unstable to exploitation by cheats, who achieve greater reproductive success than other members of the group by enjoying the benefits of cooperation without paying the associated costs. The paradoxical Prisoner's Dilemma (PD) is a frequently used metaphor for these problems (1, 2). In its standard form, the PD is a game played by two players, each of whom may choose (independently but simultaneously) to cooperate, C, or defect, D, in any one encounter. If both players choose C, both get a payoff of magnitude R; if one defects while the other cooperates, D gets the games' biggest payoff, T, while C gets the smallest, S; if both defect, both get P. With $T > R > P > S$, the paradox is evident. In any one round, the strategy D is unbeatable (being better than C whether the opponent chooses C or D). But by playing D, both players end up scoring less than they would by cooperating (because $R > P$).

There is by now a large literature—several thousand papers—seeking to elucidate the relative merits of different strategies (such as tit-for-tat) when the game is played many times between players who remember past encounters. These theoretical analyses, computer tournaments, and (to a small extent) laboratory experiments continue, with the answers depending on the extent to which future payoffs are discounted, on the ensemble of strategies present in the group of players, on the degree to which strategies are deterministic or partly stochastic (either by design or by accident), and so

on (3–11). Fascinating though this work is, it is clear that simpler biological entities (self-replicating molecules, bacteria, and arguably most non-human animals) that exhibit cooperative interactions cannot obey the restrictions—recognizing past players, remembering their past actions, and anticipating future encounters whose payoffs are not significantly discounted—which necessarily underpin such strategic analyses.

We have recently proposed an alternative approach, which considers only two kinds of players: those who always cooperate, C, and those who always defect, D (12, 13). No explicit account is taken of past or likely future encounters, so no memory is required and no strategic considerations arise. Our compensating new complication is to place the players—individuals or organized groups—in a two-dimensional spatial array or lattice of patches, with each lattice-site thus occupied either by a C or a D. In each round, every individual "plays the game" with a defined set of immediate neighbors. After this, each site is occupied either by its original owner or by one of the neighbors, depending on who scores the highest total in that round. And so on, round by round. We have shown that this simple and purely deterministic, spatial version of the PD can generate self-organized spatial patterns (some of which are relatively static, and others of which are chaotically changing) in which C and D both persist indefinitely—that is, these spatial PDs show that polymorphisms of C and D persist, over a wide range of values for the parameters that characterize the immediate advantage of cheating, without any reference to memories, anticipated future encounters, or strategic niceties.

Our initial work on spatial versions of the PD and other evolutionary games is, however, based on several simplifying assumptions. In this paper, we show how the results can be generalized in several ways, using more realistic assumptions (with the original work representing a limiting case). First, we generalize the deterministic assumption that a given site is "won" by the neighboring player with the largest total score, to allow for "probabilistic winning": the current site-holder or any relevant neighbor may win the site, with probabilities that depend to a specified extent on the relative score. Second, we go beyond the earlier analyses based on sites arranged as regular lattices, to spatially random distributions of sites (the game now being played with neighbors within some specified distance). We think these two generalizations are important, because real situations are likely to involve probabilistic winning (rather than the largest score always triumphing) and irregular arrays (rather than strictly symmetric lattices). A third extension is to continuous time (with individual sites "playing the game" with neighbors and being updated one by one) in contrast to discrete time (with the entire array simultaneously being updated each round). Huberman and Glance (14) have already analyzed this contrast between continuous and discrete time, but only for a single value of the "cheating-advantage" parameter; as spelled out below, we think the conclusion they draw from this restricted analysis is misleading.

The publication costs of this article were defrayed in part by page charge payment. This article must therefore be hereby marked "advertisement" in accordance with 18 U.S.C. §1734 solely to indicate this fact.

Abbreviation: PD, Prisoner's Dilemma.

We now outline the results of each of these three generalizations of our earlier work and also sketch some other extensions. As summarized by Figs. 1 and 2, the essentials of our original conclusion—that spatial considerations can be enough to enable C and D to persist polymorphically—remain intact. A more detailed account is given elsewhere (15).

Probabilistic Winning

As before, we begin by rescaling the PD payoffs to have the canonical values $R = 1$, $T = b$ ($b > 1$), $P = \varepsilon$ ($1 > \varepsilon \geq 0$), $S = 0$. For a one-parameter exploration of the dynamics of our spatial game, we further simplify by putting $\varepsilon = 0$; this preserves the essentials of the PD, and numerical studies show that none of our findings is qualitatively altered if instead ε is finite (although significantly smaller than unity). The character of the systems' behavior then depends on the magnitude of the cheating-advantage parameter, b . For the illustrations in Figs. 1 and 2, the PD is played with the eight nearest neighbors and with one's own site (which is reasonable if the players are thought of as organized groups occupying territory). The arena consists of an $n \times n$ cartesian array, and the boundaries are fixed, so the players at the boundaries have fewer neighbors; the essential results are unchanged if instead we use periodic boundary conditions.

We now introduce a degree of stochasticity into the contest for ownership of sites or cells (15). At any given time, define $s_i = 0$ if site i is occupied by D, and $s_i = 1$ if C. Let A_i denote the payoff to the occupier of site i , from playing the PD game with itself and with the $v_i - 1$ neighbors with which it is defined to interact. Then we define P_j , the probability that site j is occupied by a C in the next round, to be

$$P_j = \sum_{i=1}^{v_i} A_i^m s_i / \sum_{i=1}^{v_i} A_i^m. \quad [1]$$

The parameter m then characterizes the degree of stochasticity in the contest for sites. In the limit $m \rightarrow \infty$, we recover the deterministic limit studied earlier: site j will be C in the next round if the largest score among the sites $\{v_i\}$ is from a C-owned site, and D otherwise. In the opposite limit of $m \rightarrow 0$, we have random drift: the probability that site j will be C or D in the next round depends on the proportions of C and D in the current set of neighbors $\{v_i\}$. For $m = 1$, the probability for site j to be C or D is linearly weighted according to the scores of the relevant contestants ("proportional winning").

Fig. 1 summarizes our results. The horizontal axis represents the cheating-advantage parameter, b , which increases from left to right. The vertical axis represents the "degree-of-stochasticity" parameter, m , decreasing from top to bottom; the top line is for $m \rightarrow \infty$ and so corresponds to our previous deterministic analysis. The color coding is as follows: blue represents a C site that was C in the preceding generation; red, a D following a D; green, C following D; and yellow, D following C (so that a primarily red and blue pattern is static—note that the polymorphic patterns are indeed mainly static). These results speak for themselves. For all m , from deterministic winning to random drift, we see a clear qualitative pattern: for values of b close to unity, the system becomes all C (all blue); for relatively large b (approaching 2), the system becomes all D (all red); but for a wide range of intermediate b values, there are persisting polymorphisms of C and D. In these polymorphic cases, the proportions of C versus D tend to depend on the starting proportions for relatively small b , but for larger b the proportions are essentially independent of the initial configuration. Of course, the details do depend on the degree to which winning is probabilistic (as measured by m). As m decreases from

very large values (essentially deterministic winning), D fares somewhat better in that b values that for larger m gave polymorphisms now give all D (see, e.g., the columns toward the right in Fig. 1). But as m decreases below unity, moving toward random drift, the band of b values producing polymorphisms again widens.

Spatial Irregularities

We showed earlier that the basic results were independent of whether players interact with eight neighbors or with the four orthogonal neighbors in square lattices, or with six neighbors in hexagonal lattices. But spatial arrays in nature will rarely, if ever, have strict symmetry. We have therefore made extensive computer simulations of our spatial PD when the individual sites or players are distributed randomly on a plane (15). Players interact with those neighbors who lie within some defined radius of interaction, r ; this means, *inter alia*, that different sites can interact with different numbers of neighbors. Specifically, we generated the random array by starting with a 200×200 square lattice and then letting some proportion (say, 5%) of the cells, chosen at random, be occupied by players; these "active" cells thenceforth defined the random array. The "interaction radius," r , varied from 2 to 11 (measured in units of the original lattice) in different simulations.

We explored our spatial PD for these irregular arrays, for various values of r and b , in the limit of deterministic winning ($m \rightarrow \infty$). As for the symmetric lattices—the top row of Fig. 1—we found persistent polymorphisms of C and D for a range of intermediate b values, provided r was not too big (conversely, if players interact with too many neighbors, the system became all D). The specific limit to r consistent with maintaining polymorphism, r_c , depends on b (e.g., for $b = 1.6$, $r_c \sim 9$), with r_c decreasing as b increases. In all cases, the patterns settled to become relatively static (mainly red and blue), and the proportion of C and D tended to depend on the initial configurations, especially for relatively small b .

Continuous Versus Discrete Time

Our original studies of spatial PDs were for discrete time, in the sense that the total payoffs to each site were evaluated, and then all sites were updated simultaneously (12, 13). This corresponds to the common biological situation where an interaction phase is followed by a reproductive phase; although the game is usually played with individual neighbors in continuous time, at the end of each round of game playing the chips are cashed in, and the cashier pays in fitness coinage. Similar things happen in many other contexts (host parasitoid interactions, or prey-predator models where dispersal and territory acquisition is followed by raising the young), resulting in biological situations in which individual events—like challenges for territories, eating and being eaten, and so on—occur in continuous time, yet the appropriate simple model is one with discrete time (16, 17). There are, however, some situations where it may be more appropriate to work in continuous time, choosing individual sites at random, evaluating all the relevant scores, and updating immediately. Huberman and Glance (14), indeed, suggest that "if a computer simulation is to mimic a real world system . . . it should contain procedures that ensure that the updating of the interacting entities is continuous and asynchronous." We strongly disagree with this extreme view, believing that discrete time is appropriate for many biological situations, and continuous time for others.

Be this as it may, Fig. 2 represents the simulations summarized by Fig. 1, but now using continuous time (sequential updating of individual sites) rather than the discrete time (synchronous updating) of Fig. 1 (15). The broad features of




FIG. 1. Spatial dynamics of cooperators, C, and defectors, D, in discrete time simulations. The simulations were performed on an 80×80 square lattice with periodic boundaries. Interaction occurs with the eight nearest neighbors and also includes self-interaction. Each picture shows the spatial distribution after 200 generations. The parameter b denotes the payoff that D obtain by interaction with C (C obtain payoff 1 from interactions with other C). The columns show the outcome for different values of b . From the left we have $b = 1.05, 1.13, 1.16, 1.35, 1.42, 1.55, 1.71, 1.77, 1.9, 2.01$. The rows show different values of m , the parameter that describes the degree of stochasticity in the contest for sites, Eq. 1. For $m \rightarrow \infty$ a cell is always given to the most successful neighbor; $m = 1$ denotes an intermediate case, where cells are given to neighbors with probabilities being proportional to the payoffs of individual neighbors; $m = 0$ denotes spatial drift irrespective of achieved payoffs. From the top we have $m = \infty, 100, 20, 10, 1, 0.5, 0$. The color code is as follows: blue, a C that was a C in the previous generation; red, a D that was a D in the previous generation; green, a C that was a D in the previous generation; yellow, a D that was a C in the previous generation. The amount of green and yellow shows how many cells change from one generation to the next. A picture that is only red and blue is completely static.

Fig. 2, showing the proportions of C and D found for various values of b and m , are qualitatively similar to those of Fig. 1.

There are some differences of detail. For a small range of relatively large b values, continuous time leads to all D, whereas discrete time gave polymorphisms of C and D (compare some of the boxes in the columns to the right in Figs. 1 and 2). In particular, for deterministic winning and b between 1.8 and 2 (the box second from the right in the top row in Figs. 1 and 2), we find polymorphism for discrete time and all D for continuous time. Huberman and Glance (14) consider only this single case, and from it they draw sweeping conclusions about the differences between continuous versus discrete time models. The broad similarities between Figs. 1 and 2, however, show plainly that these conclusions are mistaken.

Other Extensions

Elsewhere, we have considered several other extensions (15). One is to three-dimensional arrays, either symmetric or irregular and with deterministic or probabilistic winning; the results are similar to the two-dimensional ones. If some sites may become unoccupied (“death”), and remain so if surrounded by sites with low payoffs, then C is easier to maintain. Suppose

payoffs to self-interactions are weighted by a parameter a , relative to payoffs from interactions with neighbors (we have been considering $a = 1$, but we could have $a > 1$ or $a < 1$): polymorphisms can be maintained in the absence of self-interaction ($a = 0$), but only for relatively large values of m and for smaller values of b ; for $m = 1$ (proportional winning), C cannot persist in the absence of self-interaction. A simpler “mean field” analysis is possible for models in which sites disperse propagules globally, in proportion to their total payoff; polymorphisms of C and D can now be maintained for all $b > 1$, with the fraction of C being linearly proportional to the self-interaction parameter a .

Conclusions

In our earlier papers (12, 13), we could not resist extended discussion of the beautiful patterns that unfolded if the starting configuration was symmetrical. As we emphasized at the time, however, these “Persian Carpets” are of aesthetic and mathematical interest but have no direct biological significance (although mathematical understanding of such special cases helps illuminate the dynamics of the game more generally). They obviously cannot arise if there is any asym-

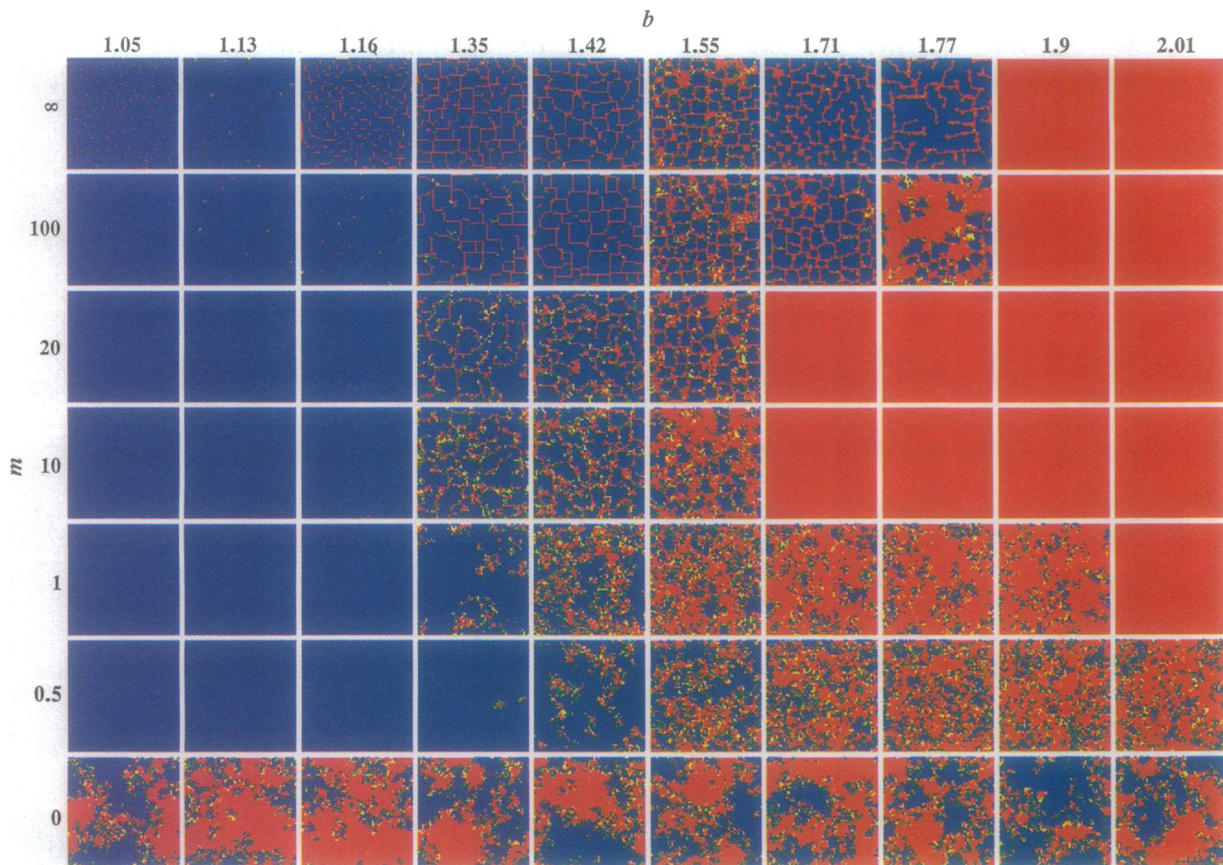


FIG. 2. As Fig. 1, but for continuous time. In each elementary time step a cell is chosen at random, its payoff is compared to the neighbors' payoffs, and the cell is updated immediately. Exactly the same parameter regions as in Fig. 1 are shown. Note that the interesting behavior for $1.8 < b < 2$ in the deterministic game has disappeared, but the general phenomenon of coexistence between C and D for a range of different b values remains completely unchanged. In the stochastic game ($m = 1$) C seem to do better in the continuous time simulations than in the discrete time simulations.

metry or departure from determinism, as there will be if we have probabilistic winning, or spatial irregularity, or continuous updating (much less, of course, an asymmetric initial configuration). Less trivially, in our original study with deterministic winning in symmetrical lattices, we identified a particularly interesting domain of dynamical behavior within the larger polymorphic domain, in which the C–D polymorphism showed chaotic variation in space and time (lots of green and yellow along with the red and blue; see, e.g., the “dynamic fractals” in the box second from the right in the top row of Fig. 1, corresponding to $2 > b > 1.8$). This particular regime of dynamical behavior involves C winning along straight-line boundaries between local clusters of C and D, while D gains along irregular boundaries. As a more local version of the global considerations just noted, any kind of randomness that disturbs the straight lines will destroy this particular dynamical regime. Thus the small regime of dynamic fractals behavior, which we emphasized earlier, will not arise once we have probabilistic winning, or spatial irregularities, or continuous time.

The main conclusion of our earlier, and much more restricted, analysis remains intact. As Figs. 1 and 2 show, C and D can persist together for a broad band of values of the cheating-advantage parameter b . This is broadly true for deterministic or various degrees of probabilistic winning, for symmetric or irregular arrays in two or three dimensions, and for discrete and continuous time. Cooperators can persist in the spatial PD, without any need to invoke memory of previous encounters, anticipation of future ones, or strategic complexities.

More generally, spatial effects can confound intuition about evolutionary games. Thus, for example, it can be seen that equilibria among strategies are no longer necessarily characterized by their having equal average payoffs; a strategy with a higher average payoff can converge toward extinction; and strategies can become extinct even though their basic reproductive rate (at very low frequency) is larger than unity. This is because the asymptotic equilibrium properties of spatial games are determined by “local relative payoffs” in self-organized spatial structures, and not by global averages. Although we have focused here on the PD, our overall conclusion is that interactions with local neighbors in two- or three-dimensional spatial arrays can promote the coexistence of strategies, in situations where one strategy would exclude all others if the interactions occurred randomly and homogeneously.

1. Trivers, R. L. (1971) *Q. Rev. Biol.* **46**, 35–57.
2. Axelrod, R. & Hamilton, W. D. (1981) *Science* **211**, 1390–1396.
3. Axelrod, R. (1984) *The Evolution of Cooperation* (Basic Books, New York); reprinted (1989) by Penguin, Harmondsworth, U.K.
4. Axelrod, R. & Dion, D. (1988) *Science* **242**, 1385–1390.
5. May, R. M. (1987) *Nature (London)* **327**, 15–17.
6. Selten, R. & Hammerstein, P. (1984) *Behav. Brain Sci.* **7**, 115–142.
7. Milinski, M. (1987) *Nature (London)* **325**, 433–435.
8. Nowak, M. & Sigmund, K. (1989) *Appl. Math. Comput.* **30**, 191–213.
9. Lindgren, K. (1991) in *Artificial Life II*, eds. Langton, C. G., Farmer, J. D., Rasmussen, S. & Taylor, C. (Addison-Wesley, Redwood City, CA), pp. 295–312.

10. Nowak, M. A. & Sigmund, K. (1992) *Nature (London)* **355**, 250–253.
11. Nowak, M. A. & Sigmund, K. (1993) *Nature (London)* **364**, 56–58.
12. Nowak, M. A. & May, R. M. (1992) *Nature (London)* **359**, 826–829.
13. Nowak, M. A. & May, R. M. (1993) *Int. J. Bifurcation Chaos* **3**, 35–78.
14. Huberman, B. A. & Glance, N. S. (1993) *Proc. Natl. Acad. Sci. USA* **90**, 7716–7718.
15. Nowak, M. A., Bonhoeffer, S. & May, R. M. (1994) *Int. J. Bifurcation Chaos* **4**, 33–56.
16. Godfray, H. C. J. & Hassel, M. P. (1989) *J. Anim. Ecol.* **58**, 153–174.
17. Maynard Smith, J. (1974) *Models in Ecology* (Cambridge Univ. Press, Cambridge, U.K.).