

A Voter Model of the Spatial Prisoner's Dilemma

Marcus R. Frean and Edward R. Abraham

Abstract—The prisoner's dilemma (PD) involves contests between two players and may naturally be played on a spatial grid using voter model rules. In the model of spatial PD discussed here, the sites of a two-dimensional lattice are occupied by strategies. At each time step, a site is chosen to play a PD game with one of its neighbors. The strategy of the chosen site then invades its neighbor with a probability that is proportional to the payoff from the game. Using results from the analysis of voter models, it is shown that with simple linear strategies, this scenario results in the long-term survival of only one strategy. If three nonlinear strategies have a cyclic dominance relation between one another, then it is possible for relatively cooperative strategies to persist indefinitely. With the voter model dynamics, however, the average level of cooperation decreases with time if mutation of the strategies is included. Spatial effects are not in themselves sufficient to lead to the maintenance of cooperation.

Index Terms—Coexistence, competition for space, evolution of cooperation, prisoners dilemma, voter model.

I. INTRODUCTION

IN THE prisoner's dilemma (PD), two agents make independent choices about whether they will cooperate with one another or defect. The payoffs are structured so that an individual agent does best if it defects while the other player cooperates. At the same time, the combined payoff to both agents is highest for mutual cooperation and lowest for mutual defection. This gives rise to the paradox that even though the two agents would be better off to fully cooperate with one another, self-interest leads to them both defecting. This dilemma elegantly captures the tension between group and individual incentives for behavior, an issue of crucial importance in understanding how functional wholes arise from their constituent parts under selection. Such tensions are a key feature in several facets of cognitive science, most notably the formation of signaling systems, including language.

The relationship between payoffs is shown in Fig. 1 for two players denoted X and Y having degrees of cooperation x and y , respectively. In a PD situation, the payoff for Y has the form indicated by the lower surface in Fig. 1: the slope of the surface is toward the rightmost corner $(1, 0)$, meaning that Y always benefits from increases in X 's cooperation, but is always better off decreasing its own level. The payoff surface for X is the same as this but with axes reversed, so the maximum is in




Fig. 1. Payoff surfaces for two agents X and Y caught in a PD. Each player adopts a degree of cooperation between zero and one and the payoff received depends on both agents' choices. X 's individual surface is not shown, but is the mirror image of Y 's. The slope of either individual's surface favors defection, despite the fact that the total payoff is highest for mutual cooperation.

the leftmost corner (i.e., X prefers a situation where Y is fully cooperative while X itself is not). These two summed together give the total payoff, which is the upper surface in the figure. This has its maximum at the uppermost corner $(1, 1)$, meaning that mutual cooperation yields the maximum total payoff. Any decrease in cooperation by either agent decreases the total value of the joint state. In what follows, we take the payoff surface for Y to be $(4x - y + 1)/5$, where x and y are the respective levels of cooperation of the two players (this is equivalent to the fairly conventional PD payoffs of $R = 4$, $T = 5$, $P = 1$, and $S = 0$ [1], rescaled to lie between zero and one).

Consider a population of agents with fixed degrees of cooperation in the face of the PD payoff structure. In every pairwise interaction, the least cooperative party gets a higher payoff. If individuals mix freely and payoffs translate directly into fitness, the most cooperative individuals are the least fit on average. Conventional evolution by preferential reproduction of fit individuals would lead to the system being dominated by noncooperative individuals. The PD, therefore, raises the question of how evolution allows cooperation to be maintained in environments where defection or free-riding is possible. Not surprisingly, the ubiquitous nature of cooperation in the real world has inspired a number of replies to this challenge. One answer lies in the theory of kin selection [2] and, more generally, trait group selection [3]. Another plausible explanation is reciprocity, which applies to cases where the game is played many times over the lifetime of the agents. If contacts between players are liable to be repeated, it can be advantageous to cooperate if doing so engenders the cooperation of the other agent. The generic example

Manuscript received January 17, 2000; revised November 2, 2000. This work was supported by the Marsden Fund, administered by the Royal Society of New Zealand.

M. R. Frean is with the School of Mathematical and Computing Sciences, Victoria University, Wellington, New Zealand.

E. R. Abraham is with the National Institute of Water and Atmospheric Research (NIWA), Kilbirnie, Wellington, New Zealand.

Publisher Item Identifier S 1089-778X(01)03283-0.

of such a successful behavior is the strategy of “tit-for-tat” [1], which always copies its coplayer’s behavior and its noise-tolerant cousin “generous tit-for-tat” [4], which sometimes cooperates even though its coplayer does not.

II. COOPERATION AND SPACE

Nowak and May [5] introduced the additional idea that spatially structured interactions could promote cooperative behavior, even without reciprocity. Their model consists of an array of sites, each of which is occupied by either a cooperator or a defector. The individual at each site is assumed to interact with each of its immediate neighbors and itself, thereby accruing a total payoff. Once these have been found, the individual with the highest payoff invades the central site.

Nowak and May showed that for a range of payoff values, cooperators (C) could coexist indefinitely with defectors (D). While a C player, which is surrounded by D players, is invaded, a cluster of C players can fend off invasion because they accumulate payoffs from one another. With deterministic dynamics, fractal patterns of cooperators and defectors are seen while non-determinism in the form of noise leads to spatial chaos. The basic result was very striking: even naïve cooperators have a chance against defectors in a spatially organized system. In addition, the makeup of the population remained nontrivial with two strategies persisting indefinitely rather than degenerating into a monoculture. Many studies have explored this model further, also showing that spatial structure enhances cooperation [6]–[13].

This paper investigates an alternative model for the spatial PD. The rules we use are as follows: at each time-step, a site is chosen at random from an array of strategies. This site encounters one of its immediate neighbors (chosen at random) and the two strategies play a PD game. With a probability that increases with its payoff¹ from the interaction, the strategy in the first site then spreads to occupy the second. This process can be viewed as either the invasion of the second site by the occupant of the first or as the original occupant of the second site surviving, but adopting what appears to be the more successful strategy. Of course, the effect is the same: the strategy of the first spreads to that of the second. Because at any time-step this model only involves two players, it is a natural way of embedding PD games onto a grid.

The rules described above are equivalent to one of a family of voter models [14], [15]. These have been used to model situations as diverse as the spread of human voter preferences or the spread of biological species where space is a finite resource and have received considerable attention from theorists because of their generality and simplicity [16]–[18]. In the simplest voter model, each species has a given probability of invading each other species and sites are chosen at random, exactly as above. In our case, the matrix of invasion probabilities is determined by the payoffs the PD strategies accrue when they play one another, given the usual payoffs.

III. SIMPLE LINEAR STRATEGIES

Several results from the voter model are useful here. If there are just two species present and they invade one another at different rates, the species with the higher invasion probability will eventually drive the other to extinction² [16]. In this case, the model is said to cluster. If neither is dominant because they invade one another at the same rate, the eventual outcome is still the extinction of one species rather than stable coexistence [17], provided the dimensionality of the grid is one or two. If the grid is three-dimensional or more coexistence of multiple species is possible due to a result concerning random walks [18]: two random walkers will eventually hit each other with complete certainty in one or two dimensions, whereas they can permanently avoid this (with some nonzero probability) in three or more dimensions. If more than two species are initially present, then, provided there are no cycles in the dominance relations between species pairs, one species will eventually take over the entire domain.

Consider the PD game with players having a fixed degree of cooperation $0 < x < 1$. In this case, a lower degree of cooperation always means a higher payoff and the higher payoff translates into a higher rate of invasion. It follows that there is a clear hierarchy of strategies, with the least cooperative ($x = 0$) at the top and the most cooperative at the bottom. Of any pool of fixed strategies present in equal amounts initially, the least cooperative is most likely to drive all the others to extinction. Alternatively, if we begin with a population of pure cooperators and allow x to mutate a small amount upon each successful invasion, the lowest x -valued players spread the fastest in any given population. Evolution will drive the mean level of cooperation down to its minimal value of zero.

Unlike the model considered by Nowak and May, space is not sufficient for fixed cooperators to persist. We next consider the outcome if simple linear strategies for the iterated prisoner’s dilemma (IPD) are allowed. These strategies take account of their coplayer’s previous move and their degree of cooperation in the next iteration is a linear function of their coplayer’s move. Each strategy can be drawn as a straight line crossing the unit square whose axes are the degrees of cooperation adopted by two players, as the example in Fig. 2 shows. For strategy Y , the degree of cooperation is $y = Y0 + (Y1 - Y0)x$, where the parameters $Y0$ and $Y1$ are the levels of cooperation Y takes for $x = 0$ and $x = 1$. For example, if Y was “always defect,” it would be a horizontal line at the bottom of the square. The rising diagonal corresponds to tit-for-tat. When iterated, the joint state of the two players rapidly approaches the intersection of the two lines over successive iterations [19] independently of whether or not players make their decisions at the same time [20]. The value of a long iterated game for either party is the value of its payoff at this intersection point. The square shown in Fig. 2 is the base upon which the payoff surface shown in Fig. 1 is plotted. Along the diagonal $y = x$, both players adopt the same degree of cooperation and the game is a draw. From Fig. 1, if the point of intersection of the two lines is below the diagonal, Y gets a higher

¹The particular mapping from payoffs to invasion rates does not play a significant role, so long as it is monotonic.

²If the array is small, it is possible for the slower invader to prevail against the odds by chance alone.




Fig. 2. Simple linear players in the IPD. Strategy Y takes the current x value (the degree of cooperation of its opponent) and outputs a new level y . Likewise, X takes the current y value and returns a new x . Hence, every update by the players moves the joint state (x, y) closer to the intersection of the two lines. For simplicity, we ignore the transient and give the agents their payoffs according to the position of the intersection point alone.

payoff than X and wins the game. Conversely, if the intersection is above the diagonal, X wins. Of particular importance is the intersection of a strategy line with the rising diagonal: this is the outcome of a game in which the strategy plays another copy of itself. Since tit-for-tat lies along the diagonal it neither wins nor loses, but draws every contest.

Simple linear strategies lead to identical outcomes (both dynamical and evolutionary) to the more familiar case of probabilistic all-or-nothing cooperation [19]. Nonspatial populations playing IPD become dominated by generous tit-for-tat: reciprocity results in cooperation. Incorporating spatial effects with the IPD leads, as a general rule, to greater levels of generosity than the IPD on its own [9]–[13]. In the case of the voter model, however, we next show that space, even with reciprocity, is still insufficient for cooperation to be sustainable.

As before, the strategies present at any one time have pairwise dominance relationships and the most important property is whether the graph so defined contains a cycle or consists of a clear hierarchy. It can be shown geometrically that no cycles of the form $Y \Rightarrow X \Rightarrow Z \Rightarrow Y$ or $Y \Rightarrow X \Rightarrow Z \Rightarrow Y$ are possible, where $X, Y,$ and Z are simple linear strategists and the arrows indicate either dominance or a draw. To show this, consider a strategy Y that wins against a strategy X . This means the intersection of the two lines when drawn as in Fig. 2 is below the diagonal $x = y$. Denote the value at which the X line crosses the diagonal by x^* . Clearly, the Y line must intersect the diagonal at a point $y^* < x^*$. Thus, there is an order on the linear strategies defined by where they cross the diagonal, which is the outcome of the self-game. From this, it can be seen that there can be no cycles involving any number of players. Because invasion probability is a monotonically increasing function of payoff, this rules out intransitivity in the matrix of dominances and the clustering result holds: after sufficient time, only a single strategy will remain. All strategies with $Y_0 = 0$ draw with one another and defeat those with $Y_0 > 0$; so, if initially present, one of these will be the eventual winner. Apart from tit-for-tat ($Y_0 = 0, Y_1 = 1$), all these strategies defect against themselves and the end result is likely to be a population trapped in

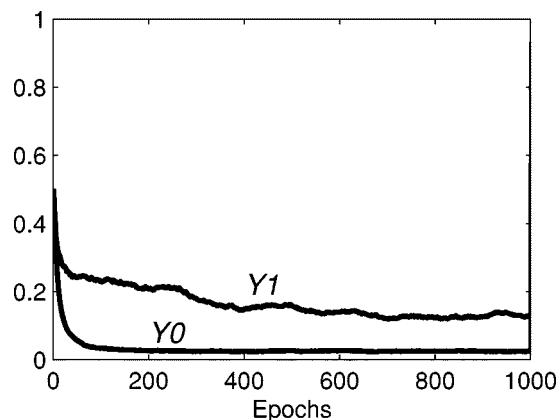


Fig. 3. Evolution of simple linear strategies under voter model dynamics. A 50×50 array of sites was initialized with random strategies by choosing Y_0 and Y_1 randomly between zero and one, independently for each site. The model was then run for 1000 epochs, where one epoch corresponds to choosing 2500 pairs of sites, as dictated by the voter model dynamics. For each pair, the payoffs that result from a long IPD game are calculated and an invasion occurs with a probability given by the payoff for the first site. The graph shows the time course of the population averages of parameters Y_0 and Y_1 . To allow new strategies to enter, each invasion is accompanied by the addition of 1% random noise to the parameters and, on average, 0.1% of invasions result in a totally new random strategy. All the cooperative strategies are wiped out over time.

perpetual defection. If the linear strategies are subject to mutation as they invade, then the system evolves, becoming increasingly defective over time (Fig. 3).

IV. CYCLIC COMPETITION AMONG IPD STRATEGIES

Although linear strategies do not allow cooperation to persist, a simple nonlinearity suffices to allow cycles (intransitive relationships) between strategies. Consider the following, for example.

- 1) X cooperates at a level of 80%, regardless of the other player's behavior.
- 2) Y cooperates at a level of 20%, regardless of the other player's behavior.
- 3) Z cooperates fully if the other player is above 50% and otherwise not at all, making it a highly "opinionated" form of tit-for-tat.

Y gains a higher payoff than X , exploiting its cooperative nature. Z beats Y because Z cooperates the least of the two when they play one another. Despite this, Z loses to X because X does not cooperate fully when Z does. No doubt, there are many ways cycles could be constructed given flexible strategies, but this is perhaps the simplest possible: two of the strategies are unreactive players and the third has only a trivial nonlinearity. Converting these payoffs into invasion rates³ shows that the strategies form a competitive loop.

In its general form in the voter model, such loops have been shown to have some interesting properties. First, the voter model permits multiple species to coexist indefinitely if they are in a cycle of three [16]. This is most easily shown by considering the same system without any spatial component to the dynamics.

³A minor technicality is that we must assume some noise to find Z 's payoff against itself: two Z players spend half their time in mutual full cooperation and the other half in mutual defection.




Fig. 4. Simulation of IPD with three strategies that form a competitive cycle. Strategies X , Y , and Z are soft cooperators, soft defectors, and opinionated tit-for-tat. They are shown by white, black, and grey, respectively.

In the limit of infinite system size, the mean-field theory ordinary differential equation (ODE) yields a family of periodic orbits around a fixed point, which means there are regular oscillations in the proportion of each species. However, without spatial structure, the magnitude of these oscillations increases, leading inevitably to extinction of two out of the three species [21]. By contrast, Durrett and Levin [16] showed that the spatial case exhibits stable coexistence. When viewed at a local scale, there are semiregular oscillations, well separated regions oscillate out of phase, and the result is a stable equilibrium for each of the three types at large scales. This dynamical coexistence of multiple species cannot occur at all without a loop and cannot persist indefinitely if the system is not spatially structured.

Fig. 4 shows the spatial distribution of strategies on the lattice for the competitive cycle formed by the IPD strategies X , Y , and Z given above. A lattice of 256×256 sites was initialized randomly with these three behavioral rules and run for 2000 epochs. The three strategies continually invade one another, forming a fragmented pattern of clumps and clusters. With the particular strategies and payoffs used here, opinionated tit-for-tat succeeds in occupying the largest proportion of the grid. Most of the sites which are occupied by opinionated tit-for-tat form a large connected cluster that spans the domain.

V. EVOLUTION

We have shown that there are sets of strategies, including cooperative ones, that are dynamically stable in a spatial system, but not otherwise. An additional question is whether such behaviors are evolutionarily stable. For instance, what happens if the parameters determining these behaviors suffer mutations as they spread into new sites?

In the cyclic voter model, if each species has control over how quickly it invades the species it dominates, evolutionary pressures strongly reinforce such cycles [22]. An initially sluggish cycle will tend to accelerate under the influence of natural selection because it pays each member of a species to invade more

aggressively. In an intriguing echo of the PD, this acceleration occurs despite the fact that it is more advantageous for the population as a whole to invade less aggressively.

In the case of the IPD however, the invasion rate is not a trait of the player alone but of the interaction between two players since it depends on both their levels of cooperation. It is not difficult to see that this interaction is enough to make the above example of a competitive loop X , Y , and Z unstable under evolution because all three strategies will become less cooperative over time. For example, individual mutants of the partial cooperator X would be more successful in encounters both with Y and Z if they cooperated at a lower value (but still above 50%). Y would similarly be better off against both X and Z by cooperating less than 20%. There are two ways that Z can be less cooperative: by setting a higher threshold in deciding whether to cooperate or by being less cooperative above the existing threshold. Both of these lead to higher payoffs against X without changing those against Y , so mutants with these qualities will be favored by evolution. All these tendencies reduce the level of cooperation and the end result is that all three strategies converge on unconditional defection.

More generally, it is enough to note that no strategy, no matter how complex, can outscore pure defection (“AllD”) in a one-to-one game. Any strategy which cooperates with AllD gets a lower payoff and, hence, is dominated by it. Among a set of strategies that do not cooperate at all with AllD, there may be dominance relationships, however, leaving a reduced set of strategies that draw with AllD and with each other. Evolution will always tend to favor these strategies over their more cooperative counterparts.

VI. DISCUSSION

To be successful in a spatial competition system, a strategy must be able to invade others. The probability of a successful invasion naturally depends on the payoff the individual at a given site obtains through its interactions with its neighbors. The voter model explicitly requires the payoff from a given site to be “spent” on the invasion of that same site—it is not possible for an individual to use this payoff against another neighbor, for instance. We have seen here that while cooperation may persist through the occurrence of cycles of nonlinear strategies, evolution will lead to the prevalence of defecting strategies. By contrast, in spatial models that do give rise to sustained levels of cooperation, we always find some form of transfer of payoff between where it is obtained and where it is “spent” as an invasion. Nearly all the work on spatial PD allows payoffs to be accrued by playing games with a number of sites. This cumulative payoff is then used to invade a single other site, implying a transfer of payoffs. With these rules, cooperation may persist as the higher payoff obtained from playing a cooperative site may be used to invade a defecting one.

If the payoff from interacting with a site must be spent against that same site, only the relative gain—the difference between the two payoffs—matters. The remainder of the payoff leads to an equal expected number of invasions in either direction. On the other hand, if an individual can spend a locally obtained payoff on some third party, it is the overall value that matters. These




Fig. 5. Evolution of simple linear strategies under modified voter model dynamics. Figure shows the same simulation as that in Fig. 3, but run with the slightly altered dynamics as described in the text. The result is evolution of generous tit-for-tat.

two can be markedly different motivations. Indeed, it is almost a defining characteristic of the PD that cooperators forego the relative gain in favor of a higher absolute gain for both parties.

A straightforward way to test this idea is to make a slight change to the voter-IPD model dynamics. As before, a site is chosen at random from the array and a second is chosen from its immediate neighborhood. A game of the IPD is then played and the payoff used to determine the ability of the first site to invade. However the site that is actually invaded is not necessarily the same as that played against, being chosen instead *at random* from the immediate neighborhood of the first site. In all other respects, the dynamics are the same. Fig. 5 shows the outcome of such a simulation for the same class of simple linear strategies discussed earlier. The result is now the rapid rise of the cooperative strategy of generous tit-for-tat, the same behavior which prevails in a randomly-mixed "soup" model [4], [19].

In general, spatial effects are thought to enhance the prospects of cooperative behaviors. With the voter model rules, cooperative and defecting strategies can coexist indefinitely provided there are competitive cycles, but once evolution is included, defecting strategies prevail. The spatial clustering of strategies is not sufficient to favor the survival of cooperation. The study of the voter model shows that without transferability of payoffs, spatial effects do not, on their own, result in the promotion of cooperation.

REFERENCES

- [1] R. Axelrod, *The Evolution of Cooperation*. New York: Basic Books, 1984.
- [2] W. D. Hamilton, "The genetical evolution of social behavior. I and II," *J. Theor. Bio.*, vol. 7, pp. 1–32, 1964.
- [3] E. Sober and D. S. Wilson, *Unto Others: The Evolution and Psychology of Unselfish Behavior*. Cambridge, MA: Harvard Univ. Press, 1999.
- [4] M. Nowak and K. Sigmund, "Tit for tat in heterogeneous populations," *Nature*, vol. 355, no. 6357, pp. 250–253, Jan. 1990.
- [5] M. Nowak and R. May, "Evolutionary games and spatial chaos," *Nature*, vol. 359, no. 6398, pp. 826–829, Oct. 1992.
- [6] M. Nowak, S. Bonhoeffer, and R. May, "More spatial games," *Int. J. Bifurcation Chaos*, vol. 4, no. 1, pp. 33–56, Feb. 1994.

- [7] P. Grim, "The greater generosity of the spatialized prisoner's dilemma," *J. Theor. Bio.*, vol. 173, no. 4, pp. 353–359, Apr. 1995.
- [8] G. Szabo and C. Toke, "Evolutionary prisoner's dilemma game on a square lattice," *Phys. Rev. E*, vol. 58, no. 1, pp. 69–73, July 1998.
- [9] G. Pollock, "Evolutionary stability of reciprocity in a viscous lattice," *Soc. Net.*, vol. 11, no. 3, pp. 175–213, Sept. 1989.
- [10] V. C. L. Hutson and G. T. Vickers, "The spatial struggle of tit-for-tat and defect," *Phil. Trans. Roy. Soc. Lond. B*, vol. 348, no. 1326, pp. 393–404, June 1995.
- [11] M. Nakamaru, H. Matsuda, and Y. Iwasa, "The evolution of cooperation in a lattice-structured population," *J. Theor. Bio.*, vol. 184, no. 1, pp. 65–81, Jan. 1997.
- [12] K. Brauchli, T. Killingback, and M. Doebeli, "Evolution of cooperation in spatially structured populations," *J. Theor. Bio.*, vol. 200, no. 4, pp. 405–417, Oct. 1999.
- [13] M. D. Cohen, R. L. Riolo, and R. Axelrod, "The emergence of social organization in the prisoner's dilemma: How context preservation and other factors promote cooperation," Santa Fe Inst., Santa Fe, NM, Working Paper 99-01-002, 1999.
- [14] R. Durrett, *Lecture Notes on Particle Systems and Percolation*. Belmont, CA: Wadsworth, 1988.
- [15] T. S. Mountford, "Generalized Voter Models," *J. Stat. Phys.*, vol. 67, no. 1/2, p. 303, Apr. 1992.
- [16] R. Durrett and S. Levin, "Spatial aspects of interspecific competition," *Theor. Population Bio.*, vol. 53, no. 1, pp. 30–43, Feb. 1998.
- [17] J. T. Cox and D. Griffeth, "Mean field asymptotics for the planar stepping stone model," in *Proc. London Mathematical Soc.*, vol. 61, July 1990, pp. 189–208.
- [18] R. Durrett and S. Levin, "Stochastic spatial models: A users guide to ecological applications," *Phil. Trans. Roy. Soc. Lond. B*, vol. 343, no. 1305, pp. 329–350, Feb. 1994.
- [19] M. R. Frean, "The evolution of degrees of cooperation," *J. Theor. Bio.*, vol. 182, no. 4, pp. 549–559, Oct. 1996.
- [20] —, "The prisoner's dilemma without synchrony," in *Proc. Roy. Soc. Lond. B*, vol. 257, July 1994, pp. 75–79.
- [21] F. J. Weissing, "Evolutionary stability and dynamic stability in a class of evolutionary normal form games," in *Game Equilibrium Models I: Evolution and Game Dynamics*, S. Reinhard, Ed. Berlin, Germany: Springer-Verlag, 1991.
- [22] M. R. Frean and E. R. Abraham, "Rock-scissors-paper and the survival of the weakest," *Proc. R. Soc. Lond. B*, 2000, to be published.



Marcus R. Frean received the Ph.D. degree in neural networks from Edinburgh University, Edinburgh, U.K. in 1990.

He was a Post-Doctoral Fellow at Otago University, Dunedin, New Zealand, where he studied the prisoner's dilemma. He was also with Cambridge University, Cambridge, U.K., and the University of Queensland, Victoria, Australia, where he researched machine learning. He is currently with Victoria University, Wellington, New Zealand. His current research interests include the evolution of

cooperation and artificial intelligence.



Edward R. Abraham received the Ph.D. degree in theoretical physics from Cambridge University, Cambridge, U.K., in 1992.

He is currently an Oceanographer with the National Institute of Water and Atmospheric Research (NIWA) in Wellington, New Zealand, where his recent research has been on the effects of upper ocean physics on plankton ecology. His current research interests include the way in which biological systems are organized by the interaction of physical and biological processes.