

Game-theoretic Models in Biology

Anatol Rapoport
 Departments of Psychology
 and Statistics
 University of Toronto
 Toronto, Ont., Canada

Received January 24, 1979

The immense deductive power of mathematics has been amply demonstrated in the physical sciences. It is not surprising, therefore, that many investigators outside those sciences, for example biologists, psychologists, sociologists, and linguists, impressed by the prowess of mathematical methods, have made repeated attempts to extend those methods into their own domains. The degree of success of these attempts is difficult to assess for want of a proper standard of comparison. Certainly the achievements of the mathematical approach outside the physical sciences appear very modest when compared to what those methods have achieved in the physical sciences. On the other hand, if we juxtapose what has been done against what was once thought impossible to do, the achievements appear much more impressive.

Of one thing we may be fairly sure. Arguments that mathematics is not suited as an investigative tool in this or that field because of the "nature of the subject matter" carry little force. At one time Sir William Hamilton, a Scottish metaphysician (not to be confused with Sir William Rowan Hamilton, the Scottish mathematician) dismissed Boole's mathematical model of symbolic logic as useless on the ground that logic, being a science of the "laws of thought", was qualitative, not quantitative, whereas mathematics was quantitative, not qualitative, hence irrelevant to logic. Resistance to the mathematization of psychology and of other behavioural sciences often stems from similar sentiments. What can or cannot be done, or to what extent something can be done, can be decided only by attempts to do it. Logic has been mathematicized in the sense of being reduced to manipulation of symbols in accordance with prescribed rigid rules. Experimental psychology has incorporated mathematical techniques at least in its inductive phase, that is, in critical and rigorous evaluation of results through statistical inference, and occasionally even in the deductive phase, that is, drawing observable consequences from hypotheses about non-observable events. Much of the same applies to empirical sociology. Thus, the blanket dismissal of mathematical methods as in principle inapplicable beyond their natural habitat in the physical sciences need not be taken seriously.

Nevertheless, the solid fact remains that the exquisitely edifice that rose on mathematical foundations in the physical sciences has nowhere been duplicated. In summary, we might say that the extension of mathematics beyond the physical sciences has been a partial or limited success.

Opinions about the future of mathematics outside the physical sciences remain divided. The sceptics still insist that mathematics is powerless to deal with behavioural phenomena (some even say life processes of any kind) either in principle, for example because of interventions of "free will", or because those phenomena are too complex to be represented by the sort of highly simplified or idealized models that are tractable by mathematical methods. The optimists expect that mathematical methods or techniques, perhaps aided by computer technology, will eventually become sufficiently powerful to deal with problems of arbitrary complexity.

* * * * *

It should be pointed out, however, that those who see the difficulties of applying mathematical methods outside of the physical sciences primarily in the "complexity" of the subject matter, miss an important point. Namely, it is not technical difficulties that limit the prowess of classical mathematical methods in the new contexts; that is, not the difficulties in solving huge systems of differential equations, but rather conceptual difficulties. Technical difficulties may well be overcome by technical means, for example, by increasing tremendously facilities for numerical computations. Conceptual difficulties are those that keep us in the dark about what sort of equations we should set up. In fact, we are by no means certain that any important theoretical questions in biology or in the behavioural sciences can be answered at all in terms of solutions of equations like the ones that are the physicist's or the chemist's stock-in-trade.

Accordingly, there has appeared a third group of opinions on the future of mathematics outside the physical sciences. Some contend that mathematical methods may well be relevant in the biological or social sciences, but only if new forms of mathematics are developed, especially suited to treat non-physical phenomena.

Now in recent decades two new branches of applied mathematics have been developed that are geared to problem areas quite outside physics. By that I mean that the models to which these types of mathematics are applied are not derived from physical laws. One is information theory, an outgrowth of the theory of stochastic processes with special emphasis on measures of uncertainty. The other is the mathematical theory of games, a generalization of a mathematically rigorous theory of rational decision. It is with the latter that we shall be here concerned. I mention the former in order to point up the somewhat misguided enthusiasm that greeted these developments or at least the publicity that accompanied them.

For here were indeed instances of "new mathematics", moreover mathematics geared not to physical phenomena determined by well-known immutable laws but to phenomena that have apparently direct relevance to human activities such as communication and strategic conflict.

The spectacular technological achievements of the past half century or so have ushered in an era of unlimited expectations. One now expects not merely "new inventions" (that's old stuff) but "break-throughs", a revolution around every corner. And these expectations are no longer confined to technology. The mood is catching. People in academe, pursuing more or less conventional lines of more or less routine investigations, are on the look-out for "methodological break-throughs". So for a number of years we had a lot of talk about "applying information theory" to this or that -- to psychology and sociology and political science and physiology and even ecology and anthropology. I would not go so far as to say that all this talk was nonsense. Far from it. It may make sense to measure the channel capacity of a neural element. The deciphering of the genetic code was certainly a deep and challenging problem, doubtless inspired by information-theoretical concepts. However, whatever mileage was obtained from these concepts was derived through strictly technical applications, not metaphorical ones. Most of the talk about measuring the volume of "information flow" is a society as if it were something analogous to current flowing through an electrical system has now mercifully subsided. These phantasies were inspired partly by wishful thinking (a "methodological break-through"); mainly, however, by metaphorical transformations of terms used in information theory, a dilution of these terms, from their rigorous technical meanings to their diffuse everyday meanings or to metaphorical meanings with the usual penumbrae of connotations.

The much publicized terminology of game theory has undergone similar treatment. Game theory can be reasonably defined as a theory of rational decisions in conflict situations. Well, then! Isn't that what we are looking for? Here is something that must be eminently practical. In the burgeoning militarist literature published in the era of the Cold War in the United States a central question has been posed: "How can the United States use its military might rationally in

Game-theoretic Models in Biology

the pursuit of its national interests?" So let's see what we can learn from game theory to answer this question. The management of a corporation certainly thinks in strategic terms. What should be our advertising strategy? Our strategy of future investments? Of penetration into new markets? All of these questions must be answered in a competitive environment, that is, in the face of similar and antithetic attempts by opponents. Now the "scientists" have come up with a new strategic science. Might be worth looking into. Business is a game. Politics is a game. Diplomacy is a game. Courtship and competition for mates is a game. Welcome, another breakthrough clearing the way for scientific management of life.

These are, surely, the most glaring examples of a vulgarized conception of a scientific discipline, induced partly by wishful thinking, partly by confusing technical concepts with everyday notions coded into the same vocabulary. But both of these unfortunate tendencies pervade also the thinking of serious investigators. In our day, "interdisciplinary approaches" enjoy great prestige. At least great hopes are pinned on them. Consequently, a habit of thought has developed of being on the look-out for applications to one context of ideas developed in another. Cross-fertilization, they call it. It is well to remember, however, that cross-fertilization can lead to sterility as well as to hybrid vigour. Let us, therefore, subject the conjecture that game theory can be fruitfully applied in biology to critical scrutiny. Above all, let us be on the look-out for pitfalls as well as for opportunities.

The context in which game theory was originally developed was that of strategic conflict. A strategic conflict, in contradistinction from other kinds of conflict, is one where the participants make conscious choices of action among available alternatives and moreover base their choices on anticipated consequences. More than that, strategic thinking, at least as it is defined in game theory, implies awareness of the fact that other participants in the conflict are also making choices among alternatives available to them and are guided by anticipated consequences. Finally, a participant in a strategic conflict is aware of still another circumstance, namely, that the "consequences" depend not only on his choices but also on the choices of others and that the preferences of the various participants for the possible outcomes of choices do not, as a rule, coincide. In fact, in the simplest types of games there are exactly two participants (players), whose preferences for the possible final outcome of a sequence of decisions are diametrically opposed.

It seems natural to define game theory as a theory of rational decisions in conflict situations. If "rational decisions" is indeed what game theory is concerned with, what possible relevance can it have for biology? Do organisms other than man make "rational decisions", that is, decisions guided by anticipated consequences, and moreover decisions that take into account the effect of others' decisions on consequences, where these others are trying to bring about consequences that are preferred by them but not by the organism in question? It seems that to assume this much would amount to admitting the crassest sort of teleology into biological theory. There may be biologists willing to do this, but others would surely reject this radical shift of method, which might be thought to imply a return to a naive, pre-scientific outlook.

Fortunately, the admission of some game-theoretic ideas need not mean a re-introduction of a frankly teleological outlook. Recall how notions of teleological guidance of evolution were expunged from evolutionary theory by the theory of natural selection. This was done by showing how apparently "purposeful" or foresightful adaptations could be accounted for without recourse to actual purpose or foresight. The giraffe got his long neck not because his purpose was to reach the leaves on trees but because the genes that made for progressively longer necks remained in the gene pool, while others were eliminated.

Evolutionary adaptation is to a species what unconscious learning is to the individual organism. The learning of muscular skills (typing, driving, etc.) is "unconscious" in the sense that we have no idea of the sequences of the muscular contractions that determine the eventually acquired skilled performance. As we practice, some sequences of contractions are "rewarded" in the sense that the performance is more successful and some are "punished". Therefore in a very real sense natural selection (namely of particular sequences of minute actions) guides the learning process.

Now game theory provides algorithms for singling out "optimal" decisions in certain conflict situations. If we can design learning procedures in which these optimal decision would be "selected for" and, moreover, if we could show how such learning operates in well-defined conflict situations, we would have a basis for incorporating some techniques of game theory into biology, where learning is interpreted as evolutionary adaptation.

I keep stressing the importance of including a conflict situation into any game-theoretic model. This is because I prefer to define game theory as dealing with decisions of at least two players, that is, participants with generally different sets of interests. Others extend the notion of a game to any situation where the outcomes depend partially on circumstances other than the choices made by the actor in question. In particular, if these circumstances can be specified as the "states of the world", one is said to have a "game against nature". What distinguishes a game against nature from a genuine game is that Nature has no interests. Therefore neither foresight nor "learning" tending toward some sort of survival-insuring adaptation can be attributed to her. In fact, the problem of deducing an optimal decision in a game against nature is altogether different from that deduced in a genuine game. Nevertheless, as an introduction to the subject of our discussion, we can examine an instance or two of models where evolutionary adaptation is deduced as an optimization of a decision in a game against nature. The following model is due to R. C. Lewontin (1961).

The player is a population of self-fertilized plants, which may be homozygous for an allele A or a. Then "to be a genotype AA" or "to be a genotype aa" are the two "strategies" available to an individual. Suppose, further, that the aa individuals produce 5 seeds in dry weather and 10 seeds in wet weather. The corresponding numbers of seeds produced by AA individuals are reversed. Consider a time span of t seasons (generations). The possible "states of the world" during those t seasons can be represented by integer r ($0 \leq r \leq t$), meaning that r out of t seasons are wet and, consequently, $t-r$ are dry. The "game" then has the form shown in Table 1. The payoffs are numbers of seeds produced. The entries in the matrix represent the number of seeds produced by an individual (AA or aa) in each state of nature.

* * * * *

Table 1 (after Lewontin)

Nature's strategies ($r =$)

		0	1	2	3		t
Organism's strategies	AA	5^t	$5^{t-1}(10)$	$5^{t-2}(10)^2$	$5^{t-3}(10)^3$		10^t
	aa	10^t	$10^{t-1}(5)$	$10^{t-2}(5)^2$	$10^{t-3}(5)^3$		5^t

* * * * *

In a sexually reproducing species the situation is more complicated, because the segregation of genes in meiosis insures the appearance of heterozygous individuals, Aa. In that case, the strategies would correspond to the three genotypes

AA, Aa, and aa, and the matrix would include the number of seeds produced by Aa individuals in each state of nature.

Clearly, the number of seeds is not necessarily the "utility" to the species. What the "utility" is taken to be depends on a particular definition of "fitness". This is a separate problem, one related to biological theory or even, perhaps, to the philosophy of biology, if you will. For the time being, we shall by-pass the problem, assuming the utilities as given. To fix ideas, let us assume them to be those proposed by Lewontin and consider the game defined by a single season that may be wet or dry. The game matrix with assumed utilities is

		Nature's strategies	
		Dry	Wet
Organism's strategies	AA	8	0
	Aa	3	4
	aa	2	7

If probabilities can be assigned to the states of the world, say $\text{Pr}(\text{Dry}) = 1/4$, $\text{Pr}(\text{Wet}) = 3/4$, optimization of utility is attained by choosing strategy aa, which has expected utility $(.25)(2) + (.75)(7) = 5.75$. Note that this is a "pure" strategy in the sense that it prescribes the genotype (aa) to every individual in the population. Under a different probability distribution of the states of nature, AA or Aa might be the optimizing strategy. Each of these would also be a pure strategy, because it would prescribe a single genotype to every individual in the population. An optimizing population would accordingly be monomorphic, consisting of a single genotype.

A number of questions arise in connection with this model. First, how is "utility", supposedly maximized by an organism's "decision" to be one or another genotype, to be interpreted? Second, how is a "decision" by an organism, which may in no way be subsumed under "sentient beings" (e.g., a plant), to be interpreted? Finally, what is the significance of the result which prescribes monomorphism as the "optimal" genetic composition of a population?

In our context, unless we are willing to introduce ad hoc teleological notions, we can interpret "decisions" only as results of segregation and random recombination of alleles. Once we do this, "optimization" must be interpreted as a long-range effect of natural selection, and this, in turn, forces the interpretation of "utility" in terms of differential survival potential. This view seems to contradict one proposed earlier by Slobodkin and myself (1974) to the effect that the size of surviving progeny might be dysfunctional to a species, for instance if it leads to an irreversible exhaustion of a food supply. We argued, in effect, that the only outcomes of such games against nature that have any meaning as "utilities" are those of species survival or species extinction. It is as if one were playing a gambling game for chips that could never be cashed in. The only thing that matters to such a gambler is whether he still has chips or is out of them. More chips is not necessarily better than fewer chips, because the number of chips in his possession may influence the random devices used in the game. It must be noted, however, that that argument does not apply to the present case. We are asking what mechanism can insure an approach to an "optimal" strategy. The answer is that any strategy will supersede all others in consequence of differential survival rates of the various genotypes. Consequently, the "victorious" strategy (in our case the finally remaining sole genotype), must be interpreted as optimal. The statement to the effect that the "optimal strategy" will be selected for is devoid of empirical content. It is a tautology.

We turn to the final question: What is the significance of the fact that monomorphism turns out to be the "optimal" genetic composition? Are there cases

when this is not so? We know, of course, of many instances when polymorphic populations appear to be in equilibrium. How does one deduce this result? The most obvious example of genetic polymorphism is sex. A sexually reproducing population must consist of at least two genotypes, males and females. Other examples are also known. To argue that a polymorphic population (aside from sexual differentiation) is ipso facto not in equilibrium is to beg the issue.

Game-theoretic models provide an answer of how a polymorphic population can be optimal. If Nature, instead of being neutral, were hostile and omniscient, then the optimizing strategy for a species would be not the strategy that maximizes expected utility but rather a maximin strategy. If Nature were hostile but not omniscient (that is, unable to guess the strategy chosen by the species), the optimal strategy would still be a pure strategy if the game matrix had a saddle point. But if the game matrix had no saddle point, the optimizing strategy would be mixed. In terms of genotype selection, this would mean that the optimizing population would be polymorphic. This was pointed out by R. A. Fisher in 1957. In the same article, Fisher writes that he solved a game involving mixed strategies back in 1934, ten years before the famous treatise by Von Neumann and Morgenstern appeared, although he subsequently learned that the general two-person zero-sum game was solved by Von Neumann in 1928.

We see that sex differentiation provides the most obvious example of polymorphism in nature. However, if it is meaningful to speak of optimal mixtures of the sexual genotypes, there must be some rationale for different sex ratios. Such a rationale is found in a paper by W. D. Hamilton (1967).

In most sexual organisms, the two sexes are usually found in a 1:1 ratio. R. A. Fisher offered the following explanation. Suppose that male births are less numerous than female. Then the males of that generation have a better chance than females of finding mates and, subsequently, will have more offspring. Thus, parents genetically predisposed to produce males will have more grandchildren than those genetically predisposed to produce females. The situation will continue as long as males are in the minority. Thus the 1:1 ratio of the sexes will be approached. Mutatis mutandis, the argument applies to a scarcity of females. It is noteworthy that neither the occurrence or non-occurrence of polygamy nor a differential mortality of males and females affects this argument. However, the argument does depend on special assumptions. For instance, it applies where sex-ratio control is by genes acting on the homogametic sex, or in the females of organisms where males are haploid, or by genes on the autosomes in the heterogametic sex. The argument fails, as Hamilton points out, in the case of sex-linked genes acting on the heterogametic sex.

Hamilton goes on to describe conditions that would make anomalous sex ratios advantageous. He also cites many examples of extremely skewed sex ratios found in nature. Once the plausibility of varying sex ratios is admitted, we can very well examine a game between two competing species, each of which has a continuum of strategies, namely that of a sex ratio ranging from overwhelming predominance of one sex to that of the other.

There is a class of games in which each of two opponents chooses a point on the interval $(0, 1)$, thus determining a point (x, y) on the unit square. These games are appropriately called games on the unit square. The payoffs of such a game are represented by the two functions $f_1(x, y)$ and $f_2(x, y)$ to the respective players. If $f_2(x, y) = -f_1(x, y)$, the game is zero-sum. The general theory of such games is quite complex. We shall examine a very special (simplest) sub-class called 2×2 matrix games, represented by the typical matrix

		Player 2	
		S_2	T_2
Player 1	S_1	a	b
	S_2	c	d

For player 1, to choose a point x on the interval $(0,1)$ means to choose S_1 with probability x and consequently T_1 with probability $1-x$. For player 2 to choose point y on the unit interval means analogously to choose S_2 with probability y and T_2 with probability $1-y$. The payoff functions are not arbitrary but special: $f_1(x,y) = axy + bx(1-y) + x(1-x)y + d(1-x)(1-y)$ and, because the game is zero-sum, $f_2(x,y) = -f_1(x,y)$.

There are two cases. In one, the game matrix has a saddle point, which is to say that one of the four payoffs is simultaneously minimal in its row and maximal in its column. It can be easily shown that in such a game, the best that each player can do is to choose the row (or column) in which this saddle point is located. Such a game is said to have a pure-strategy solution. That is to say x and y are either 0 or 1, so that each player chooses one of his two strategies with certainty. Games without saddle points have no pure-strategy solutions. They do have mixed-strategy solutions. That is to say x and y are numbers between 0 and 1, exclusive of the end points. To illustrate, we shall examine a game of each kind. At the same time, the demonstration will illustrate what we mean by "the best that a player can do". Consider the following two games:

	S_2	T_2		S_2	T_2
S_1	-1	5	S_1	-2	5
T_1	2	3	T_1	2	-1
Game 1			Game 2		

We shall examine Game 1 from player 1's point of view. Henceforth player 1 will be called Row (since he chooses between the two rows), and player 2 will be called Column. We shall assume that both players are "rational" in the sense that each tries to maximize his payoff and, because the games are zero-sum, to minimize the payoff of the other player. Moreover each player attributes the same sort of "rationality" to the other. Thus, if a player comes to a decision that he "ought" to choose one or the other strategy, he should assume that the other has figured out that he has come to that decision.

If Row is rational, he will ask himself what Column would do if he were certain of what Row would do. Thus, if Row were to choose S_1 in Game 1, Column would choose S_2 so as to minimize Row's payoff. If Row were to choose T_1 , Column would likewise choose S_1 for the same reason. Thus Row associates -1 with S_1 and 2 with T_1 . Since he prefers 2, he should choose T_1 . Moreover, even if Column knows that Row will choose T_1 and consequently will choose S_2 , T_1 is still Row's best choice. Similar reasoning leads to the conclusion that S_1 is Column's best choice. Neither player, facing a rational opponent, can improve the outcome from his point of view.

The situation in Game 2 is quite different. In that game, if Row decides tentatively to choose T_1 , he must realize that Column, guessing his decision, would choose T_2 . But if Row were sure that Column would choose T_2 , he should certainly choose S_1 . However, if Column knew this, he would choose S_2 , in which case Row should choose T_1 . Analogous reasoning puts Column in a similar dilemma. The notion of mixed strategy shows a way out of this impasse. Suppose Row chooses the point 0.3 on the unit square. That is to say, he chooses S_1 with probability 0.3 and T_1 with probability 0.7. Then according to our rules, if Column chooses S_2 with certainty Row gets a payoff of $(-2)(.3) + (2)(.7) = 0.8$. That is Row gets 0.8 regardless of how Column chooses. Even if Column chooses some mixed strategy $(y, 1-y)$, Row will still get his 0.8. Thus if Column chooses S_2 with probability 0.6 and T_2 with probability 0.4, he will get -0.8 regardless of how Row chooses. We see that each player can guarantee himself a certain payoff, namely Row can be sure to get 0.8 and Column can be sure that he gets no more than 0.8, because Column is sure to get -0.8. In sum, this is the best that each can do under the constraint that the opponent is also trying to do his best.

Imagine now two competing species. Interpret the strategies as "choices" of genotypes and the payoffs in terms of some survival value. If the resulting game model has no saddle point, the "optimal strategies" of each of the species is mixed. The optimal mixtures represent the proportions of the genotypes. We have the species in equilibrium when each is polymorphic. In particular, the genotypes may represent the sexes. The optimal mixtures represent optimal sex ratios in the particular competitive situation.

The task remains of interpreting "choice" of genotype in a context that excludes "rationality" or foresight. If we can somehow conceive of a "learning mechanism" (where by "learning" we mean an adaptation brought about by natural selection), we can dispose of rationality and foresight.

Suppose Row chooses some mixture other than the optimal (0.3, 0.7), say (0.5, 0.5). Then if Column chooses S_2 , Row will get $(-2)(.5) + (2)(0.5) = 0$. Consequently, Column will also get 0. On the other hand, if Column chooses T_2 , Row will get $(5)(.5) + (-1)(.5) = 2$, and consequently Column will get -2. Natural selection working on the Column species will shift the genotype of Column toward S_2 . But this shift will start natural selection working on the Row species, shifting it toward T_1 , i.e., toward a decrease of S_1 in favor of T_1 . Thus constant adjustment will take place until the equilibrium at the mixtures (.3, .7) and (.6, .4) will be established.

It is possible to arrive at the same result by introducing a differential-equation model, where the rates of change of x and y are respectively proportional to the gradients of the survival potential with respect to them. The solution of this system is typically periodic, where x and y fluctuate around the optimal values, with time averages tending to the equilibrium point. In this way, the two-person zero-sum game model is a conceptually adequate representation of two competing species in equilibrium, where optimal strategies are represented by proportions of genotypes. Polymorphism is associated with a game without a saddle point, monomorphism with a game that has a saddle point.

Let us now turn to the two-person non-zero-sum game. Here the interests of the two players are not diametrically opposed as in the zero-sum game. That is to say, some outcomes (resulting from strategy choices of the two players) may be preferred by both players to other outcomes.

Imagine two organisms, X and Y . Each produces a nutritive substance. X produces an amount x of its substance, and Y produces an amount y of its. Further, the substances are shared between the two organisms. Each keeps a fraction p of what it produces and shares the remaining fraction q with the other organism. Thus X receives the amount $px + qy$, and Y receives the amount $qx + py$. These amounts will be incorporated into the payoffs in the game to be presently described.

The utility of the total amount of nutrient received is assumed to be an increasing function of the amount and to have a negative second derivative, reflecting "diminishing returns". To fix ideas, suppose this function is logarithmic, calibrated to vanish when no nutrient is received. The total utility accruing to the organism, however, is diminished by an amount proportional to the amount of its own substance produced, reflecting an expenditure in producing the substance. Accordingly, we have the following utilities as "payoffs" accruing to the respective organisms as functions of x and y (Rapoport, 1947):

$$u_x = \log_e(1 + px + qy) - \beta x \quad (1)$$

$$u_y = \log_e(1 + qx + py) - \beta y \quad (2)$$

We will now assume that the organism "chooses" the amount to produce so as to maximize its own utility, given that the other produces a fixed amount of the substance. Equilibrium will obtain if the partial derivatives of u_x and u_y with respect to x and y respectively are set equal to zero. Thus

$$\partial u_x / \partial x = \frac{p}{1 + px + qy} - \beta = 0 \quad (3)$$

$$\partial u_y / \partial y = \frac{p}{1 + qx + py} - \beta = 0 \quad (4)$$

Solving these equations simultaneously, we obtain

$$x^0 = y^0 = \frac{p}{\beta} - 1 \quad (5)$$

as the point of equilibrium.

Now there are two questions that arise concerning the equilibrium that are of interest in biology.

First, is the equilibrium stable? We note that the equations represent two straight lines. Let us plot these lines in the (x, y) plane. Both have negative slopes, and it can be easily shown that the lines intersect in the positive quadrant if β is sufficiently small. Since only non-negative values of x and y are physically significant, we shall suppose that this condition is satisfied throughout. Biologically, this means that the "cost of production" must not be excessively large in order for the organisms to produce x and y at all.

Now one of these lines, namely $px + qy = 1/\beta - 1$, represents the "optimal line" which X is "attempting to reach" while maximizing u_x . The slope of this line is $-p/q$. The equation of y 's optimal line is $qx + py = 1/\beta - 1$, with slope $-q/p$. Graphical analysis shows that the equilibrium is stable if $q < p$; otherwise unstable. In the first case, the equilibrium is attainable and represents the end results of maximizing utilities by the individual organisms independently. In the second case, the equilibrium is fictitious, for the slightest deviation from it generates a movement away from it until either x or y attains the smallest permissible value, namely zero. If this happens, one of the organisms will produce nothing, while the other, "attempting to maximize its utility", will produce an amount given by the solution of the corresponding equation, namely $y = 1/\beta_p - 1$, if $x = 0$, or $x = 1/\beta_p - 1$, if $y = 0$. A biological interpretation of this result is that one of the organisms has become a parasite on the other.

The second question about the equilibrium, assuming that it is stable, is whether it represents the maximum utility attainable by the two organisms in this situation. In the stable case, we have seen that $x^0 = y^0 = p/\beta - 1$ represents the equilibrium. Hence

$$u_x = \log_e(p/\beta) - p + \beta = u_y. \quad (6)$$

However, if each organism produced $1/\beta - 1 > p/\beta - 1$, we would have

$$u_{x^*} = \log_e(1/\beta) - 1 + \beta = u_y. \quad (7)$$

It is easy to see that $u_{x^*} - u_x = -\log_e p + p - 1 > 0$. It follows that in "maximizing" their utilities independently, the organisms do not attain as much as they could under the circumstances.

This leads to still another question. Under what circumstances could the two organisms attain the maximum? This could happen if the amounts produced were functionally dependent. In fact, suppose $x = y$. Then

$$u_x = u_y = \log_e(1 + x) - \beta x = u$$

$$du/dx = (1 + x)^{-1} - \beta$$

Setting $du/dx = 0$ and solving for x , we get $x = x^* = 1/\beta - 1$, as expected. The same result would have been obtained if the utility of each organism were the sum of the two utilities defined above. It follows that some sort of "physiological linkage" between the organisms precludes parasitism, or establishes symbiosis, if you will, and so insures maximum utility for both.

Admittedly, it would be difficult or impossible to find a biological situation adequately represented by the extremely simple model just discussed. The value of such models is heuristic rather than representational. They bring out some concepts that could well be central in theoretical biology: conditions under which organisms linked by shared production may end up in a symbiotic or a host-parasitic relationship; conditions under which "maximization of utility" (which must, of course be interpreted in terms of some evolutionary dynamic process) leads to a "true optimum" or only to an "apparent" one, etc.

I investigated the class of models of the sort just described in 1947, the same year that Von Neumann and Morgenstern's famous treatise Theory of Games and Economic Behaviour came out in its second edition and attracted world-wide attention. But it wasn't until 1954 that I became acquainted with that theory and discovered that in the model of utility maximization that leads to either parasitism or symbiosis I was posing a problem that would engage my attention for many years to come: What are the mechanisms that bind individual utilities into a collective utility or, to put it in human terms, integrate individual rationalities into collective rationality?

The heart and soul of this problem is depicted by the most famous of two-person non-zero sum games, called Prisoner's Dilemma. In a non-zero sum game, the interests of the two players are only partially instead of diametrically opposed. This feature is reflected in the non-constant sum of the player's payoffs in the various outcomes; so that both players can prefer some outcome to some other. The general format is represented by the following matrix

	C_2	D_2
C_1	1, 1	-10, 10
D_1	10, -10	-1, -1

where the first entry in each cell is the payoff to Row, the second to Column. Individual rationality dictates the choice of D to both players. But both prefer the outcome C_1C_2 to D_1D_2 . Consequently, collective rationality dictates C to each player.

One might suppose that in repeated plays, the choice of C might be selected for by some reinforcement procedure. But this will not occur if the reinforcement is monotonically related to the individual payoffs resulting from successive paired choices. C will be learned if the utilities of the players are somehow linked. In human contexts this might happen through an extension of the awareness of "self" to the awareness of others and eventually at least partial identification of self to others. One might even suppose such a mechanism operating in some non-human social or family-raising species. Such a hypothesis, however, is clearly ad hoc, a verbal explanation that begs the question, so to say. Can we postulate a specific reinforcement schedule that leads to the fixation of the C choice in repeated plays of Prisoner's Dilemma? The following is an example (cf. Rapoport and Chammah, 1965).

Label the payoffs of Prisoner's Dilemma R, T, S, and P, where R is the payoff accruing to each player if both choose C, i.e., "cooperate"; T is the payoff to the "defecting" player, who chooses D while the other chooses C; S is the payoff to the "cooperating" player, who chooses C while the other chooses D; P is the payoff to both if both choose D. The relation that makes this 2 x 2 game Prisoner's Dilemma is $T > R > P > S$.

Next, introduce the following conditional probabilities of choice:

$$\begin{aligned} x_i &= \Pr [C_i / C_i C_j] \\ y_i &= \Pr [C_i / C_i D_j] \\ z_i &= \Pr [C_i / D_i C_j] \\ w_i &= \Pr [C_i / D_i D_j] \quad (i, j = 1, 2, i \neq j) \end{aligned}$$

In words, x_i is the probability that player i chooses C, after both chose C on the previous play; y_i is the probability that player i chooses C after he chose C while the other chose D on the previous play, etc.

If each player is characterized by given values of x, y, z, and w, each has a certain expected payoff over a series of plays. We postulate a learning model in which these conditional probabilities are reinforced or inhibited by the

associated expected payoffs. To illustrate the model in the simplest case, suppose $x = 1, y = z = 0, w = 1$ for both players. This means that when a player receives either R or T, that is, one of the two larger payoffs, he repeats the previous choice; when he receives P or S, one of the two smaller payoffs, he switches his choice. It is easy to see that two players characterized by these "learning parameters" will fix on the choice C_1C_2 after at most two plays.

To take a somewhat more interesting case, suppose $y = 0 = z; w = 1$ as above but x_i is variable. Then the expected payoffs will be given by

$$G_1(x_1, x_2) = \frac{Rx_1x_2 + T(x_2 - x_1)}{2 + x_1 + x_2 - 3x_1x_2}$$

$$G_2(x_1, x_2) = \frac{Rx_2x_1 + T(x_1 - x_2)}{2 + x_2 + x_1 - 3x_2x_1} = G_1(x_2, x_1).$$

We now assume the same learning (or natural-selection) model as in the previous case of the zerosum game, namely

$$\frac{dx_i}{dt} = k_i \partial G_i(x_i, x_j) / \partial x_i \quad (i = 1, 2).$$

Thus

$$\frac{dx_i}{dt} = k_i \frac{x_j^2 (3T + R) + 2x_j(R - T) - 2T}{(2 + x_i + x_j - 3x_ix_j)^2} \quad (i, j = 1, 2)$$

It can be shown that if the initial values of x_1 and x_2 are high, learning or natural selection will "drive" both of these conditional probabilities toward 1. If the initial values are low, the same process will drive both toward zero. In our case, if $x = 1$, the choice of C will become fixed; if $x = 0$, there will be alternation between C_1C_2 and D_1D_2 . This effect is an artifact of our assumption that $w = 0$. It is easy to construct models of the same type with parameter values that will divide the phase space into two regions, two drainage plains, if you will, the one leading to fixed cooperation, the other to fixed non-cooperation.

Note that if we regard evolution as an analogue to learning, then in the present case we are dealing with individual, not group, selection, because it is the individual's progeny that "learns" by modifying its conditional response parameters. But the learning is of "higher order", as it were, because not the probability of response is reinforced but rather the conditional probability of a response, given the outcome resulting from a paired response. The model just presented is the bare bones of the model presented by J. M. Smith and G. R. Price (1973) in the context of intraspecific combats. The latter model was also inspired by the Prisoner's Dilemma game. The so called metagame model of Prisoner's Dilemma developed by N. Howard (1966) is exactly of the same sort.

To summarize, applications of game-theoretic models to evolution are most interesting where the evolutionary process can be viewed as formally analogous to learning, but where reinforcement schedules depend not on what happens to an individual but on what happens as a result of interactions between at least two individuals.

REFERENCES

- Fisher, R. A. 1957. polymorphism and natural selection. Bull. Inst. Internat'l. Statistique 36: 284-289.
- Hamilton, W. D. 1967. Extraordinary sex ratios. Science 156:
- Howard, N. 1966. The mathematics of metagames. General Systems 11: 187-200.
- Lewontin, R. C. 1961. Evolution and the theory of games. J. theoret. Biol. 1: 382-403.
- Rapoport, A. 1947. Mathematical theory of motivation interactions of two individuals. Bull. Math. Biophys. 9: 17-27 and 41-61.
- Rapoport, A., and A. M. Chammah. 1965. Prisoner's Dilemma. University of Michigan Press, Ann Arbor.
- Slobodkin, L. B., and A. Rapoport. 1974. An optimal strategy of evolution. Quart. Rev. Biol. 49: 181-200.
- Smith, J. M., and G. R. Price. 1973. The logic of animal conflict. Nature 246:15-18.

OPEN DISCUSSION

Van Valen: With respect to this higher-order rationality as applied to evolutionary processes: you have a referee which is I gather an essential part of this mechanism, with players playing simultaneously. Does this have an analogue outside human processes?

Rapoport: To begin with, in order to apply game theory to biology, one must, in my opinion, view biological evolutionary processes as a sort of learning process. There is a mathematical isomorphism between learning and evolution. For example, one can postulate mechanisms in the brain that eliminate some sequences of neuro-impulses and facilitate others. The learning of motor skills, for example, can be viewed in this way: certain neural pathways are progressively inhibited, others facilitated. Facilitated ones "survive", to use the analogy with evolution; the inhibited ones do not, just as in biological evolution, certain genotypes are selected for, others selected against. Coming back to your question, how would this scheme apply to the evolution of certain behavior patterns? I refer to papers of Maynard Smith which treat of the selection of organisms that use certain kinds of fighting strategies, e.g., "I will attack him, but he attacks back, I will flee"; or "I will use mild fighting, but if he resorts to severe fighting, I will reply in kind." These are conditional strategies. What Maynard Smith has shown is that in a population where selection works on individuals, the equilibrium population will consist largely of individuals who use mild fighting. Have I answered your question?

Van Valen: Not really.

Rapoport: Have I understood your question?

Ronnie De Souza, University of Toronto: Could I have a try at this? It seems to me that you've given an example drawn from Maynard Smith of selection of conditional strategies, but what you've set up here is a conditional meta-strategy. What would be the biological equivalent of selecting a conditional strategy conditional on the other's conditional strategy?

Rapoport: What can be done once can be done again. Once you break through the ordinary responses and introduce conditional responses, there is no reason why you can't have conditional responses of higher order.

De Souza: Maybe you can give us a picture.

Rapoport: Suppose there are two things you can do: C and D. A conditional response (or strategy) might be: "If he does C, I'll do C; if he does D; I'll do D." The response is conditioned by what the other does.

De Souza: Right, but you can see how that could be neurologically determined.

Rapoport: Perception of the other's use of C induces C; perception of the other's use of D induces D. This is just one of four possible conditional strategies, supposing the other player has two strategies at his disposal, C and D. The four conditional strategies are:

1. Choose C regardless of the other's choice.
2. If the other chooses C, choose C, otherwise D.
3. If the other chooses C, choose D, otherwise C.
4. Choose D regardless of the other's choice.

Given a choice of one of the two unconditional strategies by one player and a choice of one of the four conditional strategies by the other, the outcome of this "metagame" is determined. It is possible to construct a higher-order metagame, where one of the players has available the four conditional strategies, while the other has sixteen strategies conditional upon these. Again, if each player chooses from the strategies available to him, the outcome of the game is determined. Formally, the structure of metagames with conditional strategies is no different from that of any other game in matrix form.

George Oster, University of California, Berkeley: It seems to me that what you've done here in the metagame is simply introduce an equation of motion; you've introduced some dynamics into the situation. You've introduced a different equation, an equation of motion which determines the sequential strategies. Now you're doing an optimization of the set of dynamic constraints. Yes, what I'm saying is the difference between the game and the metagame is the difference between a static and a dynamic game.

Rapoport: If you want to look at it that way, I would tend to agree, because in order to perceive another's conditional strategy, one would have to perceive his choices in repeated situations. As a matter of fact, repeated situations are involved in all game-theoretic models in biology, because in order for a strategy to be selected genetic "learning" must be involved.

Question: If that is true, then it's also true that there exist dynamic games for which there are no equilibria.

Rapoport: It has been shown that all finite games representable in matrix form have equilibria. However, behavior based on some learning models do not reach them. I will give an illustration.

Consider a game played on the unit square. That is, the strategies available to the two players are respectively points x and y on the interval $(0,1)$ including the end points. This is an extension to the mixed-strategy space of a 2×2 game, in which each player has just two strategies. When the first player chooses x , this means he chooses his first strategy with probability x and the other with probability $(1-x)$; similarly for the second player. In repeated plays of the game these probabilities can be interpreted as frequencies.

Let the four possible payoffs to Player 1 be g_{11} , g_{12} , g_{21} , and g_{22} . Then, if Player 1 chooses x and Player 2 chooses y , the expected payoff to Player 1 will be

$$G_1 = dxy - ax - by + g_{22} ,$$

where $d = (g_{11} - g_{12} - g_{21} + g_{22})$; $a = (g_{22} - g_{12})$; $b = (g_{22} - g_{21})$. The expected payoff to Player 2 will be $G_2 = -G_1$.

If the game has no saddle point, that is, if the optimal strategies are mixed, then Player 1's optimal strategy mixture is given by $x = b/d$, and Player 2's optimal mixture is given by $y = a/d$. Player 1's expected payoff will be accordingly

$$G_1 = g_{11}b/d + g_{21}(1 - b/d).$$

Player 2's payoff will be

$$G_2 = -g_{11}a/d - g_{12}(1 - a/d) = -G_1.$$

Assume now a long sequence of plays, in the course of which the players "learn" in the following manner.

$$dx/dt = \partial G_1 / \partial x = dy - a$$

$$dy/dt = \partial G_2 / \partial y = -dx + b$$

That is to say, the rate of change of x is proportional to the gradient of G_1 with respect to x : the more is to be gained by increasing x , the more rapidly x increases, and analogously for y .

The solution of the above differential equations is a circle in x - y space around the point (b, a) as center and with radius determined by the initial values of x and y . Thus at each moment t when the game is played, $G_1(t)$ and $G_2(t)$, which are sinusoidal functions of t , will be the payoffs to the two players respectively. If we now integrate $G_1(t)$ and $G_2(t)$ over a complete cycle, that is over the time interval $(0, 2\pi)$ and choose our units properly (these depend on the number of plays per unit time), we shall find that both players get the same payoffs per play as they would have got if they chose optimal strategies. Nevertheless, if the radius of the circle is not zero, the players never use their optimal strategies simultaneously. That is to say, an equilibrium never occurs during the sequence of plays.

Question: The third possibility is that it can wander aperiodically, never approaching even a limit cycle?

Rapoport: The model implies that the two players go around the circle periodically, never reaching the equilibrium. This is the derived consequence of the differential equations as they have been set up. One can, of course, add a "damping" term, in which case instead of a circle we shall have a spiral converging toward the center, which is the equilibrium. According to this model, the players will gradually learn to choose their x and y closer and closer to the optimal value.

Mathematically you can cook up anything. You can imagine any sort of situation and represent it by a mathematical model. The problem becomes that of finding something in the real world to fit the model. It reminds me of a story, which I often tell my audiences. I hope those that have heard it before will forgive me. There was a man who liked to fix things around the house, but the only tools he could use were a screw driver and a file. When he saw a screw that wasn't tight, he tightened it with his screw driver. Finally there were no more screws to tighten. But he saw some protruding nails. So he took his file and made grooves in the caps of the nails. Then he took his screw driver and screwed them in. To paraphrase Marshall McLuhan's famous remark, "The medium is the message", the mathematician could well say, "The tool is the theory."

Question: By speaking of higher orders of strategies, does this imply that mathematical models of this sort are hierarchically organized and can be represented that way?

A second question: By having mathematical models like this, can they have any bearing on conceptual models, say in the working processes of the human mind, for example information processing and things of this nature?

Rapoport: The question you are raising, if I understand you, is what conceptual value is involved in mathematical formalism? What kind of ideas can it suggest to the mathematician? Can he get some insights from mathematical models? Sometimes he can. I hope I have mentioned some. But it is sometimes dangerous to think non-mathematically about mathematically obtained results. This is because ordinary language imposes connotations on words quite foreign to the precise meanings of terms used by the mathematician. It is sometimes very tempting to go off on a tangent and to fall off the deep end, and I constantly warn against this sort of extrapolation. People philosophize and speculate about all kinds of implications that they believe to see in mathematical results. You know what the

much publicized Uncertainty Principle suggested to philosophers and theologians, a "proof" of the existence of Free Will and notions of this sort. People readily find support for their preconceived notions and for wishful thinking. I do not wish to discourage speculations. It's difficult, however, to draw the line between fruitful philosophizing and free-wheeling phantasies. I remember a paper I read once on the sexual connotations of war. There is such a thing as a "dirty" hydrogen bomb, and some people think of sex as "dirty". So this play on words was used as support for the idea that war and sex are somehow connected. I am sometimes appalled by the lack of responsibility or discipline that induces people to link words together and present the results as "insights". That is why I say that philosophizing, although it is a tempting and engaging pastime, is at times a dangerous one.