PART II. HUMANITY

§5. Competitive Equilibrium

In Section 3 we applied the Darwinian Selection Principle to the problems of evolution and ecology, and in Section 6 we will make some attempt to apply it to human history. This section will study another classical idea. Like the Selection Principle, it is almost a tautology, and yet has many applications.

> **The Equilibrium Principle (A. Cournot, J. Nash, J. Maynard Smith).** *Consider a collection of entities which are in competition or conflict with each other. In an equilibrium situation, no single competitor can improve its own position by changing only its own strategy.*

This idea was introduced implicitly by Augustin Cournot in 1838 in the context of economic bargaining. It was formalized by John Nash in 1949 in the context of non-cooperative game theory, and was discovered independently by John Maynard Smith (together with G. R. Price) in 1973 in the context of evolutionary biology. Several examples will be given in §5.4.

Some comments are necessary. First, note that this is a theory for the *non-cooperative* situation in which the contestants are either unable or unwilling to negotiate with each other in order to reach some mutually beneficial outcome. If we want to apply this theory to a case where there is a possibility of limited communication or even binding agreements, then this would have to be incorporated into the description of possible moves by the players. Note also that there is an unwritten subtext: For this Equilibrium Principle to be useful, it must be fairly common for an equilibrium or near equilibrium situation to arise.

Since rather different interpretations of this principle are needed in different contexts, it will be convenient to discuss it first in the study of human interactions and then in evolutionary theory.

§5.1. Economics, Politics, Crime, and War. There are many different ways that groups of human beings can interact with each other. Many of these interactions involve a struggle for competitive advantage which may be either subtle or completely overt, and may be either friendly, or aggressive, or even deadly. In all such interactions, the Nash equilibrium theory provides a fundamental tool for interpretation. (Compare Exercises 5.4a-c.) According to Ordeshook:

> *"The concept of a Nash equilibrium n-tuple is perhaps the most important idea in noncooperative game theory. ⋯ Whether we are analyzing candidates' election strategies, the causes of war, agenda manipulation in legislatures, or the actions of interest groups, predictions about events reduce to a search for and description of equilibria. Put simply, equilibrium strategies are the things that we predict about people."*

Sylvia Nasar, in her biography "A Beautiful Mind"[1], illustrates the dollars and cents impact of game theoretic ideas by describing "The Greatest Auction Ever" in 1994, when the US government sold off large portions of the electromagnetic spectrum to commercial users. A multiple round procedure was carefully designed by experts in the game theory of

[1] Nasar's book has served as the inspiration for a film of the same name. The film is excellent as entertainment, but has little relation to actual events.

auctions to maximize both the payoff to the government and the utility of the purchased wavelengths to the respective buyers. The result was highly successful, bringing more that \$10 billion to the government, while guaranteeing an efficient allocation of resources. By way of contrast, a similar auction in New Zealand, without such a careful game theoretic design, was a disaster in which the government realized only about 15% of its expected earnings, while the wavelengths were not efficiently distributed. (In one case, a New Zealand student bought a television station license for one dollar.)

I am indebted to Hector Sussmann for the following two examples, which show that the equilibrium concept can play a role even in everyday social life.

Example 5.1a (compare SCHWARTZ). At a boring party, all of the guests would like to go home early, but no one is willing to leave before midnight unless someone else leaves first. There is just one equilibrium point: everyone stays until midnight.

Example 5.1b. A group of twenty is going to dinner, and each one has the choice of an adequate meal for ten dollars or an excellent meal for twenty dollars. If paying individually, each one would choose the cheaper meal. However, they have decided to split the bill. Since the marginal cost of the more expensive meal for each person is only fifty cents, everyone chooses it.

§**5.2. Nash's Theorem.** A major contribution by Nash was the proof that equilibria exist under very general conditions. In order to make sense of his statement, we must first introduce probabilistic strategies, that is strategies which may involve random choices by the contestants. These can be desirable in many different kinds of competition. One well known example occurs in the game of poker: A player who always bluffs (placing a substantial bet even when holding a terrible hand) will surely lose drastically; yet a player who never bluffs is much too predictable and will also lose in the long run. The way out of such dilemmas, introduced by Emile BOREL in 1924, was the use of randomized strategies, with some definite probability of bluffing, even on a terrible hand. These were later popularized by VON NEUMANN AND MORGENSTERN, who introduced the term "mixed strategy".

We will usually assume that each player has only a finite number of ***pure*** (that is non-probabilistic) strategies. By definition, a ***mixed strategy*** assigns a non-negative probability to each of these pure strategies, where the sum of these probabilities must be equal to one. As an explicit example, one possible mixed strategy for a very simple poker game would be as follows. On each hand:

With probability $1/10$, place a bet no matter what cards you hold.

With probability $9/10$, bet only if holding a pair of kings or better.

Here is a formal statement of Nash's result. (For a proof, see Appendix B1.)

> **Theorem 5.2a.** *Suppose that there are finitely many competitors, each with a choice among finitely many pure strategies. Suppose that each one is trying to maximize an individual "**payoff**" function which depends on the choices of all of the players. Suppose also that this payoff function is measured in units which are linear with respect to probabilities. Then there exists at least one way (but often many ways) for the players to choose mixed strategies so that no player can increase his expected payoff by changing his own strategy while the mixed strategies of the other players remain fixed.*

Some words of explanation may be needed:

The *expected payoff* can be defined roughly as the average payoff over a very large number of repeats of the contest, assuming that each of the contenders follows the prescribed strategy.

The *linearity* assumption means for example that a player would be equally satisfied with a certainty of gaining five units, or with a fair coin toss to decide whether the gain will be ten units or zero units. The use of payoff functions with this linearity property has a long history, going back to von Neumann. It is essential for the theory, but in many cases its realism can be questioned. One case in which it certainly makes sense is for a win-lose game, with only two possible outcomes for each player. For such cases, each player simply wants to maximize the probability of winning, so the linearity hypothesis is completely reasonable. Similarly, in the evolutionary context which is discussed below we can identify the payoff with the expected number of fertile offspring, or with some similar measure of evolutionary *fitness*. With this definition, linearity again becomes clear. (Of course this is the payoff only from the evolutionary point of view. It has little to do with the wishes of the individual, except to the extent that these wishes are shaped by evolution to maximize fitness. On the personal level, most people would surely prefer to have two healthy children, rather than having four children with a fifty per cent chance that each one would die early.) Even in an economic context, the payoff should not be simply measured in money, although it will usually be an increasing function of the monetary award. Lotteries can exist only because many people are willing to pay more than one dollar for one chance in a million of winning a million dollars. This is not simply a statement that people are irrational. The insurance industry is built on similar probabilistic estimates, and most people believe that it is quite sensible to buy suitable insurance coverage.

Another implicit hypothesis is that each player cares only about his own payoff. If some player cares not only about his own score but also about the other scores, then this must be taken into account in constructing the payoff function.

Note. For the special case of a two person *zero-sum* game, where any gain to one contestant results in an equal loss to the other,[2] Theorem 5.2a had been proved already in 1928 by von Neumann, and in special cases even earlier by Borel. In this classical case, the equilibrium is essentially unique, so that each player has a uniquely defined optimal payoff when optimal strategies are followed. Nash's contribution was to introduce and study the equilibrium concept for more general games. In this more general context, von Neumann's ideas were quite different, and were too convoluted for any useful application.

Nash's Theorem leaves many open questions which have been studied by subsequent workers. For example, which equilibrium points are likely to be reached in practice? To what extent are these equilibrium configurations stable? (See §5.5, as well as 5.3b.)

A more serious problem with equilibrium theory is the question as to whether humans really act like rational agents, trying to maximize some linear utility function. Actually, this seems rather unlikely. In the economic context (and even more in the evolutionary context which is discussed below), perhaps the more accurate interpretation is that *the equilibrium principle is a corollary of the selection principle (survival of the fittest)*. When there are many competitors, some of them are bound to exploit any deviation from an equilibrium strategy. Hence any corporation (or any species) which deviates substantially is likely to be punished.

[2] Note that real world situations are almost never zero-sum. For example most commercial transactions are *mutually* beneficial.

Over the short term, this will result in diminished prosperity and, over longer periods of time, species will die out and corporations will fold. The survivors will be those which follow a close-to-equilibrium strategy.

§**5.3. Ecology and Evolution.** In an ecological or evolutionary context, consider a population of organisms which are competing with each other for evolutionary survival. Maynard Smith speaks of an *evolutionarily stable strategy* or *ESS* for such a population when the equilibrium principle is satisfied, with a suitable stability condition. He writes [in 1995]:

> "*I was led to the idea when trying to explain ritualized behavior in animal fights. I did not take the idea seriously until I found that it can be used to explain peculiar sex ratios (as Fisher and Hamilton knew before me), plant growth, animal migration, male mating behavior, and even the evolution of viruses.*"

In this context, the concept of a "strategy" doesn't seem to make any sense at all. Perhaps lions may have some conscious hunting strategy, but insects surely do not. Yet this equilibrium principle is often quite successful in explaining evolution in species as diverse as lions or insects. In fact, the *strategy* of a population of organisms must be defined, not as any conscious set of choices, but rather as the physical capabilities and instincts which are built into its members. (Compare Appendix A.10.) Maynard Smith writes in [1982, p.10]:

> "*A 'strategy' is a behavioral phenotype; i.e., it is a description of what an individual will do in any situation in which it may find itself. An ESS is a strategy such that, if all the members of a population adopt it, then no mutant strategy could invade the population under the influence of natural selection. The concept is couched in terms of 'strategy' because it arose in the context of animal behavior. The idea, however, can be applied equally well to any kind of phenotypic variation, and the word strategy could be replaced by the word phenotype; for example, a strategy could be the growth form of a plant, or the age at first reproduction, or the relative numbers of sons and daughters produced by a parent.*"

Such a strategy must be compared with alternative strategies, that is with different instruction sets or physical capabilities which are attainable by small modifications of the existing ones. Here the word "attainable" must be emphasized. Again from Maynard Smith 1982 [p. 7]:

> "*Clearly not all variations are likely for a given species. ⋯ In some cases, the possible range of phenotypic variation may be quite sharply circumscribed; for example Raup (1966) has shown that the shapes of gastropod shells can be described by a single mathematical expression, with only three parameters free to vary. Further, the processes of development seem to be remarkably conservative in evolution, so that the evolution of legs, wings and flippers among the mammals has been achieved by varying the relative sizes and, to some extent, numbers of parts rather than by varying the basic pattern, or bauplan.*"

Of course it would be advantageous for an organism to be simultaneously stronger and faster, with sharper senses; but in practice a gain in one area is usually balanced by a loss in another. To give one explicit example, L. Donehower and his colleagues have recently discovered that the "p53" protein in mice (and presumably in humans) works to suppress cancer. (Compare Ferbeyre and Lowe.) Yet mice with too much p53 will age prematurely. Thus a gain in one area is balanced by a loss in another. As another example, the sickle cell gene (described in 3.12c) conveys a definite advantage in leading to greater resistance to Malaria, but also a definite disadvantage in leading to a proportion of the population with a debilitating anemia. Thus again a gain in one area is balanced by a loss in another.

To explain the concept of ESS in mathematical terms, let me use the notation $E_1(s\,,\,t)$ [respectively $E_2(s\,,\,t)$] for the expected payoff to the first player [or to the second player] in a two person contest where the first player follows strategy s and the second follows strategy t . For the moment, I consider only games which are completely symmetric between the two players. By definition, this means that they have the same set of possible strategies, and that the expected payoff $E_2(s\,,\,t)$ to the second player is precisely equal to $E_1(t\,,\,s)$.

Consider the situation in which randomly chosen pairs from a large population engage in a two person competition, and where all of them follow the same mixed strategy s . In order for this to be an equilibrium against a possible invasion by the alternative strategy a we must have the following:

> **Equilibrium Condition 5.3a.** *The inequality* $E_1(s\,,\,s) \;\ge\; E_1(a\,,\,s)$ *is satisfied for every possible mutant strategy* a *. (This says that no single player can obtain a larger payoff by switching to the strategy* a *.)*

However, if equality holds in 5.3a, then some fraction of the population may adopt strategy a by random genetic drift. To avoid this, Maynard Smith also requires the following:

> **Evolutionary Stability Condition 5.3b.** *If equality holds in 5.3a, then* $E_1(s\,,\,a) \;>\; E_1(a\,,\,a)$ *. (This says that a player whose opponent plays* a *will definitely do better to adopt the equilibrium strategy* s *.)*

By definition, the strategy s is an ESS if both 5.3a and 5.3b are satisfied.

Maynard Smith emphasizes that there are two radically different ways of interpreting this equilibrium theory. There is the *mixed strategy* interpretation, where all members of the population are identical, and each one follows a randomized strategy. Alternatively there is the *polymorphic* interpretation, which each individual follows a pure strategy, but there is some mixture of individuals with different pure strategies in the population. In any detailed study, for example of stability, it is important to be clear as to which interpretation is intended. Of course, in real world examples, some combination of these two interpretations would usually be closer to the truth.

This study of a single population is only the beginning of the important applications of game theoretic ideas to population biology. In more complicated situations, one often has several quite different populations interacting, for example in a predator-prey relationship. Even within a single homogeneous species, there can be complicated interactions between male and female, and between children and adults, which can usefully be described in game-theoretic terms.

Note also that equilibrium theory is far from the final answer. Hofbauer and Sigmund write [1998, p. xxiii]:

> "*The definition of evolutionarily stable strategy (if the residents adopt it, no mutant can invade) is based on an implicit dynamics. It is easy to make this dynamics explicit, by assuming that like begets like. This yields the replicator equation describing the evolution of the frequencies of different strategies in a population. Thus dynamics is not merely a prop to sustain arguments from equilibrium the-*

ory. For many games, equilibria alone do not suffice to describe what happens, and a static outcome cannot be expected. ⋯ ”

They emphasize that, in an ecological context, actual equilibrium may be too much to expect. We should be equally satisfied with a periodic or chaotic situation, provided that the populations of the various species remain bounded away from zero.

For examples illustrating some of the ways in which equilibrium theory theory can be applied in evolution, see Exercises 5.4d-g below.

§**5.4. Some Examples.** Here are seven mathematical models, presented as exercises for the reader, which illustrate the equilibrium principle in various contexts. All of these models are highly simplified, to illustrate the essence of the competitive situation. The first three are taken from classical economic theory and from political science.

Exercise 5.4a***. Pricing Spring Water (or Oil).** Cournot, in 1838, described a mathematical model which is rather unsophisticated, but which non-the-less may help us to understand the power of cartels. He considered n competitors, each able to produce very large quantities of some commodity at negligible cost. The example he chose was spring water; but perhaps today the pricing of oil in the Middle East would be an even better example. Cournot assumed that everything placed on the market would be sold, but that the price $p(s)$ per unit quantity would be a decreasing function of the total supply s. Thus, if the i-th producer decides to supply the quantity s_i, then his or her profit will be the product $s_i\,p(s)$ where $s = s_1 + \cdots + s_n$ is the total supply. In an equilibrium situation, show that

$$\frac{\partial\,(s_i\,p(s_1 + \cdots + s_n))}{\partial s_i} \;=\; p(s) + s_i\,p'(s) \;=\; 0\ .$$

If the price function $p(s)$ has been specified, this yields the equation

$$s_1 \;=\; s_2 \;=\; \cdots \;=\; s_n \;=\; s/n \;=\; -p(s)/p'(s)\ ,$$

which must be solved for s.

Cournot points out that the producers can usually greatly increase their profits if they are able and willing to cooperate to limit production. Here is a somewhat exaggerated example. Suppose that $p(s) = 10000/e^s$, so that the price $p(s)$ decreases from $p(0) = 10000$ to zero as the total supply s increases from zero to infinity. Show that $p(s)/p'(s) = -1$, so that at equilibrium we have $s_1 = \cdots = s_n = 1$, and $s = n$. Thus each producer gets a profit of only $s_i\,p(n) = 10000/e^n$. On the other hand, with cooperation, show that their best strategy would be to choose $s_1 = \cdots = s_n = 1/n$, with total supply $s = 1$, so that each would obtain a profit of $s_i\,p(1) = 10000/(ne)$. As an example, with twelve producers, for the equilibrium solution each one obtains a profit of $10000/e^{12} \approx 0.06$, but with cooperation each could get $10000/(12e) \approx 307$, or almost 5000 times as much.

Exercise 5.4b. Ricardo's Law of Comparative Cost. This is a fundamental law which helps to explain the economics of international trade. Suppose that two nations X and Y both require, and can both produce, two commodities C_1 and C_2. (As examples, these might be wheat or clothing or computer chips.) Suppose that one unit of C_i takes x_i man-hours of work to produce in X, and y_i man-hours to produce in Y. As a

* Starred exercises require the use of calculus or infinite series. The unstarred ones are more elementary.

numerical example, these production costs might be:

$$\begin{aligned} x_1 &= 2 \text{ and } x_2 = 1 \quad \text{man-hours per unit in } X\,, \\ y_1 &= 3 \text{ and } y_2 = 4 \quad \text{man-hours per unit in } Y\,, \end{aligned}$$

If country X needs the commodity C_1 and country Y needs C_2 , then their manufacture would require 2 hours per unit, respectively 4 hours per unit. But if they simply trade on a one-for-one basis (and if shipping costs can be ignored), then X would need only 1 hour per unit and Y would need only 3 hours per unit. Thus each country can save one hour per unit by such a trade.

More generally, whenever the ratio x_1/x_2 is different from y_1/y_2 , show that both countries can profit from a trade. More explicitly, suppose that both nations have the option only of accepting or rejecting an offer to barter c_1 units of C_1 , shipped from Y to X , for c_2 units of C_2 shipped from X to Y . If the costs of transportation are negligible, show that acceptance is an equilibrium strategy if and only if

$$\frac{x_1}{x_2} \geq \frac{c_2}{c_1} \geq \frac{y_1}{y_2}\,.$$

This has the consequence that it may be worthwhile for both parties to ship a commodity from Y to X even if (as in the example above) it takes more man-hours to produce in Y . This may seem surprising, but helps us to understand why labor intensive products are so often produced in the third world.

In the real world, things are more complicated, since both sides would try to bargain for the most favorable exchange ratio c_2/c_1 . As long as $x_1/x_2 \neq y_1/y_2$ there is ample room for bargaining. (Perhaps the fairest ratio would be the geometric mean

$$\frac{c_2}{c_1} = \sqrt{\frac{x_1}{x_2}\frac{y_1}{y_2}}\,.$$

Of course, in practice one also has to worry about shipping costs, tariffs, and monetary exchange rates. If three or more nations are involved, then the situation becomes even more complicated.)

Exercise 5.4c. The Median Voter Theorem (Ordeshook [1992]). Consider a two candidate election in a perfect democracy where the candidates know the positions of the voters, announce their own positions, and where the voters trust the candidates to carry out their announced policies. *If the object of each candidate is to win, if the possible positions are simply ordered, and if each voter prefers a candidate as near as possible to his own position, show that there is a unique equilibrium strategy where both candidates take the position of the median voter.* Ordeshook comments:

> "$\cdots$ *people who complain about the fact that political parties in the United States often fail to offer the electorate meaningful choices on important issues misconstrue the purpose of elections. That purpose is not necessarily to provide meaningful choices; rather it is to select public policy in accordance with majority rule principles and to assure the rejection of radical candidates.*"

Of course we cannot apply this Median Voter Theorem literally to the real world, since disinformation and moneyed interests (bribes under the guise of political contributions) play such an important role, at least in the United States and presumably in much of the world. Furthermore, the spectrum of possible political positions is not at all simply ordered. (Ordeshook discusses many of these problems. For more about the theory of elections, see SAARI.) Nevertheless, the theorem does have real meaning in helping us to understand actual political positions.

Exercise 5.4d*. "Tit for Tat": the Iterated Prisoner's Dilemma. Perhaps the most famous example of a non-cooperative game has applications to human behavior, but is also important in evolutionary theory. The classical *Prisoner's Dilemma* game can be described as follows. Each the two players has the choice of either *cooperating* (strategy $\mathbf{C}$) or *defecting* (strategy $\mathbf{D}$). The payoff can be described in terms of numbers $T > R > P > S$, as follows. If both players choose strategy $\mathbf{C}$, then each one receives a payoff of R (the reward), while if both choose $\mathbf{D}$, then each one receives P (the punishment). But if the players choose different strategies, then the defecting player receives T (the temptation) while the cooperative player receives S (the sucker's payoff). *Show that there is a unique equilibrium point, with both players defecting.*

Now suppose that this game takes place many times with the same two players. To make this precise, suppose that after each round there is some fixed probability $p < 1$ of there being at least one more round, where p is independent of the number of rounds already played, and is independent of the wishes of the players. Suppose also that the total payoff to each player is just the sum of the payoffs to this player in the individual rounds. Evidently larger values of this probability p are more likely to encourage cooperation. *If the probability of further rounds satisfies* $p \ge (T-R)/(T-P)$ *, show that there is not only the* (*very unfortunate*) *equilibrium point at which both players always choose* $\mathbf{D}$ *, but also a much more satisfactory equilibrium point with the strategy for both players as follows:*

> **Grudging Cooperation.** *Keep playing the cooperative strategy* $\mathbf{C}$ *as long as the other player does so. But as soon as the other player defects, switch to* $\mathbf{D}$ *for all subsequent rounds of the game.*

If both players follow this strategy, note that they will actually cooperate on every play. However, it is a rather terrible strategy in practice, because of its unforgiving nature.

Now suppose that p also satisfies $p \ge (T-R)/(R-S)$. *Show that there is then a more forgiving equilibrium strategy as follows:*

> **Tit for Tat.** *Play the cooperative strategy* $\mathbf{C}$ *on the first round, and henceforth mimic the strategy of the opponent, choosing* $\mathbf{D}$ *or* $\mathbf{C}$ *according as the opponent has chosen* $\mathbf{D}$ *or* $\mathbf{C}$ *on the previous round.*

As an example, suppose that $T=4$, $R=2$, $P=0$, $S=-1$, so that the "temptation" is relatively rather strong. Nevertheless, show that the Grudging Cooperation strategy for both players forms an equilibrium point provided that $p \ge 1/2$, while the Tit for Tat strategy for both players constitutes an alternative equilibrium point if $p \ge 2/3$.

Experimental tests by Robert AXELROD in 1980 showed that the Tit for Tat strategy can be quite effective in practice. Both Allied and German troops seem to have used something like it during trench warfare in the first World War. (For further discussion, see FLAKE.) In fact, animal experiments have shown that even birds and fish can sometimes be observed to follow something like this strategy. (Compare MILINSKI.) Such results are of great interest in evolutionary theory since they help to explain the apparent paradox that animals (and humans) often cooperate with each other, although a naive understanding of the selection principle might predict total selfishness.

Remark. By definition, the players in a non-cooperative game do not communicate with each other. The idea is that, if some communication is possible, then this possibility should be incorporated into the description of the set of pure strategies. In practice, animals surely do communicate with each other, by body language and scent as well as by sound. It must be very difficult to take such effects into account.

Exercise 5.4e. Sex Ratios. Here is a classical example, studied already by R. A. Fisher. Consider a reasonably stable population made up of $P_♂$ males and $P_♀$ females. Is it more advantageous to have a male child or a female child? Assume, to simplify the discussion, that, in the next generation, each female will have an average of two children, and that all of the males will have equal chances of being the father of each child. Show then that the expected number of children for each male is $2P_♀/P_♂$. Thus, if the object is to have as many grandchildren as possible, then it is better to have a male child if $P_♂ < P_♀$, but is better to have a female child if $P_♂ > P_♀$. The only evolutionarily stable strategy for the species is to have equal numbers of males and females.

Note: Many populations do have roughly equal numbers of males and females, but there are certainly exceptions. One notable exception, among the social insects, actually serves as a vindication of the theory, since the entire mechanism for sex determination is quite different, and leads to a different prediction. However, sex ratios in human populations are also sometimes very far from the 50:50 norm. (Compare Hrdy.) The causes appear to be social rather than biological: In some cases, the explanation seems to involve selective abortion or even infanticide, while in others it is simply a matter of better care for male children.

The last two examples, due to Maynard Smith, are described in an evolutionary context, but are also relevant to some kinds of human competition.

Exercise 5.4f. Aggression or Appeasement. Consider the following two pure strategies, for example for a contest between two animals of the same species.

> **"Hawk" Strategy:**[3] Always fight to the finish, gaining a *payoff* of W with a win, but incurring damage of L in case of a loss so that the payoff will be $-L$, where $L > W > 0$. Assuming that each of the two contestants is equally likely to win, the expected payoff for a hawk in confrontation with another hawk is equal to the average of W and $-L$, that is $(W-L)/2 < 0$.
>
> **"Dove" Strategy:** In case of a confrontation between two doves, the two animals glare at each other until one gives up. Thus the winner gains a payoff of W while the loser gets zero, so that the expected payoff to each is $W/2 > 0$. In case of confrontation with a hawk, the dove runs away, so that the hawk wins W and the dove gets zero.

Let h be the proportion of hawks in the population, and let $1-h$ be the proportion of doves. Show that the expected payoff to a dove in an encounter with a random member of the population is $(1-h)W/2$, while the expected payoff to a hawk is $(1-h)W + h(W-L)/2$.

[3] Here the words *hawk* and *dove* are used as political terms, and have little or no relationship to the behavior of the birds with these names.

Conclude that it is better to be a hawk if $h < W/L$, but better to be a dove if $h > W/L$. Thus there is a unique equilibrium situation where the proportion of hawks is exactly $h = W/L$. In this stable situation, note that any animal, hawk or dove, has an expected payoff of $(1 - \frac{W}{L})\frac{W}{2} > 0$ from each random encounter.

Remarks. For fixed W, if we increase the threatened loss L, note that W/L, the equilibrium proportion of hawks, will decrease. This seems intuitively reasonable, since the hawk strategy becomes less desirable as L increases. Paradoxically, it is true to such an extent that the equilibrium payoff $(1 - \frac{W}{L})\frac{W}{2}$ to both players actually gets larger as L gets larger with W fixed, even though the expected payoff $E_j(s_1, s_2)$ to either player is either fixed or gets smaller as L increases with s_1 and s_2 fixed.

Note that the best result for all animals involved would occur if there were no hawks at all, so that the expected payoff in each encounter would be $W/2$. However, as is common in equilibrium theory and in human affairs, such an idyllic strategy would be inherently unstable. In a population consisting mostly of doves, any hawk would prosper dramatically.

Exercise 5.4g*. A "War of Attrition". Two animals of the same species will sometimes compete with each other, without actually fighting, simply by wearing each other down until one of the two gives up, say after time T . Suppose that there is a benefit of $+1$ to the winner, but suppose that *both* animals incur a cost of kT for the time and effort spent in contesting, where $k > 0$ is some constant. In other words, suppose that the net payoff to the winner is $1 - kT$, while the net payoff to the loser is $-kT$. Let us consider a strategy s in which the probability of giving up during a time interval of length dt is equal to $p(t)dt$, where $p(t) > 0$ is some smooth function with $\int_0^\infty p(t)dt = 1$. If animal number 2 follows such a strategy s with probability distribution $p_2(t)dt$, and if animal number 1 chooses the strategy of giving up at time T , let $E_1(T, s)$ be the resulting expected payoff to animal 1. Show that

$$E_1(T, s) = \int_0^T (1 - kt)\, p_2(t)dt - kT \int_T^\infty p_2(t)dt ,$$

where the first and second derivatives with respect to T are given by

$$E_1'(T, s) = p_2(T) - k\int_T^\infty p_2(t)dt , \qquad \text{and} \quad E_1''(T, s) = p_2'(T) + k\, p_2(T) .$$

If the choice of strategy s for both contestants is an equilibrium, then the function $T \mapsto E_1(T, s)$ must be constant, so that $E_1' = E_1'' \equiv 0$. Show then that

$$p_2(t) = ke^{-kt} \qquad \text{and} \qquad \int_T^\infty p_2(t)dt = e^{-kT} .$$

Since $E_1(0, s) = 0$ and since the derivative $E_1'(T, s)$ is identically zero, it follows that the expected payoff $E_1(T, s)$ is identically zero. *Thus, against an opponent choosing the strategy s , the expected payoff will always be zero, regardless of what strategy is chosen.* It follows that there is an equilibrium situation in which all of the contestants follow this strategy s . In fact Maynard Smith [in 1982] notes the sharper statement that this is the unique equilibrium, and that it is evolutionarily stable.

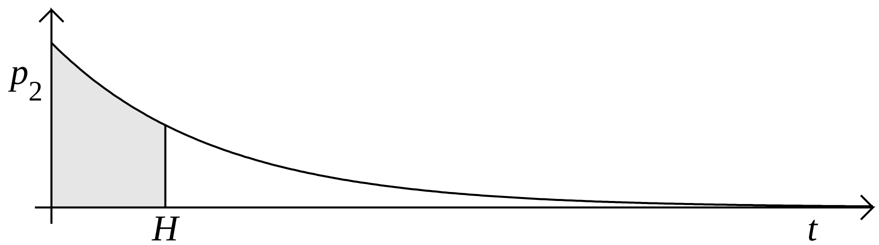

Remarks. This same probability distribution, with exponential decay of probabilities, occurs in the study of radioactive decay. Note that the associated *half-life*, that is the unique number H such that there is a 50% chance of giving up by the time H, is inversely proportional to the constant k. (More precisely, $H = \ln(2)/k$.)

As often happens with non zero-sum games, this equilibrium strategy is not at all advantageous to the contestants. Against this equilibrium strategy s, the expected payoff $E_1(T, s)$ is zero for every possible stopping time T; but if the two animals were willing and able to cooperate so that one of the two would give up immediately, then each could have a fifty percent chance of winning, and the expected payoff would be $0.5 > 0$. A more puzzling problem with this equilibrium solution is that it seems to be rather unstable. Compare the discussion in 5.5c below.

§**5.5. Stability.** Considerations of stability in equilibrium theory are clearly important, but are also rather subtle, and depend very much on context. Thus Maynard Smith's discussion of stability is concerned with the case of a large population of individuals who all choose the same mixed strategy, with pairwise competitions between randomly chosen individuals. Different interpretations would be appropriate in an economic context, and quite different interpretations in a game with only two players.

Let me consider the relatively simple case of a (not necessarily symmetric) game between two opponents who are trying to outwit each other. Intuitively I would like to describe a Nash equilibrium point as "stable" if any deviation from the equilibrium strategy is likely to be punished, and "unstable" if at least one of the players may be strongly tempted to deviate. Here is an attempt to make this precise.

> **Suggested Definition.** A Nash equilibrium point will be called *strategically stable* if any player who deviates from his equilibrium strategy will either have an immediate reduction of expected payoff, or can be punished by a change in strategy by the opponent which decreases the expected payoff for the deviating player while increasing the expected payoff for the opponent.

Thus, with notation as in 5.3, the equilibrium pair (s_1, s_2) is strategically stable if, for every alternate strategy $a_1 \neq s_1$ by the first player, either

$$E_1(a_1, s_2) < E_1(s_1, s_2) \qquad (5:1)$$

so that the change is immediately disadvantageous, or else equality holds but there exists a counter-strategy c_2 so that

$$E_1(a_1, c_2) < E(s_1, s_2) \quad \text{and} \quad E_2(a_1, c_2) > E_2(a_1, s_2)\,; \qquad (5:2)$$

with similar inequalities when the roles of the two players are interchanged. This is quite different from Maynard Smith's evolutionary stability condition 5.3b. Here is a simple example which illustrates the difference.

Example 5.5a. Trying to be Different. Each of two players chooses either **M** or **F**. The payoff is $+1$ to both players if they make different choices, but -1 if they make the same choice. This game has three equilibrium points, namely the symmetric equilibrium with an expected payoff of zero, where each players flips a coin to decide, but also two pure strategy equilibria where one player chooses **M** and the other chooses **F**. If we think of this as an evolutionary game, then the symmetric equilibrium is the unique ESS. (This can be thought of as a simplified version of the sex ratio game of 5.4e.) However, if we think of it as a simple contest between two players, then the two non-symmetric equilibria are much more satisfactory, and are the only "strategically stable" equilibria.

Example 5.5b. Rock-Scissor-Paper. Consider the children's game in which each of the two players simultaneously chooses either **R**, **S** or **P** (Rock, Scissors, or Paper); where **R** wins against **S**, while **S** wins against **P**, and **P** wins against **R**. The payoff is to be $+1$ to the winner and -1 to the loser; or zero to both players in case of a tie. Suppose that Player 2 follows an announced mixed strategy, choosing:

$$\mathbf{R} \text{ with probability } r\,, \quad \mathbf{S} \text{ with probability } s\,, \text{ or } \mathbf{P} \text{ with probability } p\,,$$

where $r+s+p=1$. Then if Player 1 chooses **R** then 1's expected payoff will be $s-p$, and similarly if he chooses **S** or **P**, his expected payoff will be $p-r$ or $r-s$. Thus if any one of these three differences is positive, then 1 can make a choice so as to win on the average. It follows that the only way that 2 can prevent 1 from coming out ahead is to choose the mixed strategy with $r=s=p=1/3$, so that all three differences are zero. There is a unique equilibrium point in which both players use this mixed strategy; and this equilibrium point is clearly "strategically stable", since any deviation from equilibrium can be punished by the other player.

On the other hand, this equilibrium is not an ESS. Since this is a zero-sum game, it follows by symmetry that $E_1(a\,,\,a)$ is zero for every possible a. Since $E_1(s\,,\,a)$ is also identically zero, this shows that the equilibrium solution $(s\,,\,s)$ is not evolutionarily stable in Maynard Smith's sense. Hofbauer and Sigmund comment as follows [1998, p. xxiv]:

> "*It used to be thought that the rock-scissors-paper game was just a conundrum devised for the amusement of theoreticians, until it was found that lizards do play it: one of their species has three different types of male with different mating strategies (they are conveniently distinguished by their throat color). Type A keeps one female and guards it closely; Type B keeps several females and necessarily guards them less closely; and C guards no females at all and looks out for sneaky matings with unguarded females. The three types can invade each other cyclically.*[4] *Similar ratchets can occur in parasitology.*"

Example 5.5c. Stability in the War of Attrition Game. The unique equilibrium described in 5.4g is evolutionarily stable, but is far from "strategically stable". To see this, for every positive constant ℓ let σ_ℓ be the strategy with probability distribution $p_\ell(t)dt$ where $p_\ell(t)=\ell\,e^{-\ell t}$. Thus, for the equilibrium situation both players choose the strategy σ_k, taking $\ell=k$. But suppose that one player deviates from the equilibrium by choosing σ_ℓ for some constant $\ell<k$. (Thus, on the average, the deviating player will hold out for longer than an equilibrium player.) There is no immediate penalty for this, since against

[4] More precisely, C can invade B, which can invade A, which can invade C.

the equilibrium strategy the expected payoff is always exactly zero. However, computing $E_1(T\,,\,\sigma_\ell)$ and its derivative $E_1'(T\,,\,\sigma_\ell)$ as in 5.4g, we find that

$$E_1'(T\,,\,\sigma_\ell) \;=\; (\ell-k)e^{-\ell T} \;<\; 0.$$

In other words, against this deviant strategy, the longer a player waits the lower his expected payoff will be. *The best that the victimized player can do in this situation is to give up immediately, obtaining the equilibrium payoff of zero for himself, but yielding a payoff of* $+1$ *to the deviating player.* Thus a contestant will be strongly rewarded for deviating from the equilibrium strategy. I am quite puzzled as to what choice a rational contestant should make in a game of this sort, where the only known equilibrium point is so strongly unstable.

§**5.6. References and Remarks.** The theories described here have their origins in the classical works by Cournot, by Nash, and by Maynard Smith and Price. For some refinements of Nash's work, see Weibull §1.4; and for more about Nash, see Nasar and also Milnor 1995. For a fuller presentation of Maynard Smith's work, see his 1982 book which includes studies of the underlying dynamics, information transfer, bargaining, cooperation, and stability. More recent presentations emphasizing both the dynamic and the game theoretic approaches to evolution have been given by Weibull, and by Hofbauer and Sigmund. For a less technical discussion of game theory, see Flake, and for less technical expositions of evolution, including some game theory, see the various books by Dawkins, for example "River out of Eden". For the evolution of cooperation, see Axelrod, or Nowak, May and Sigmund, or Ridley 1996, as well as other publications of Sigmund. For the complicated dynamics associated with predator-prey situations, see May, and for an observed example see the discussion of lynxes and hares in Leaky and Lewin. For applications of equilibrium theory to political science, see Ordeshook; and for an exposition of the view that human economic behavior is actually highly irrational and chaotic, see Ormerod, or Thaler, as well as Lewis. For stability in classical economic theory, see Arrow and Hahn.

If there is a moral to this section, it is the following fairly obvious one: Unrestrained competition or conflict is usually bad for all concerned—that is why we have laws. Ideally, a major focus of game theory, in the analysis of economics and politics, should be to study the extent to which changes in the rules of the game, that is changes in law, can mitigate this effect and help to promote the general welfare. Of course, in practice things are not so easy. Even if one could develop a clear and almost universally accepted theory, it would be very difficult to put such changes into effect. Any change in rules is bound to disadvantage someone; and in practice, much more effort is spent in studying how to bend or modify the rules in order to benefit particular individuals, or corporations, or political organizations.