



International Institute for
Applied Systems Analysis
Schlossplatz 1
A-2361 Laxenburg, Austria

Tel: +43 2236 807 342
Fax: +43 2236 71313
E-mail: publications@iiasa.ac.at
Web: www.iiasa.ac.at

Interim Report

IR-12-069

Game theory

Karl Sigmund (ksigmund@iiasa.ac.at)
Christian Hilbe

Approved by

Ulf Dieckmann
Director, Evolution and Ecology Program

February 2015

Interim Reports on work of the International Institute for Applied Systems Analysis receive only limited review. Views or opinions expressed herein do not necessarily represent those of the Institute, its National Member Organizations, or other organizations supporting the work.

Game Theory

Karl Sigmund^{1,2} and Christian Hilbe¹

1 Faculty of Mathematics, University of Vienna

2 International Institute of Applied Systems Analysis (IIASA)

Game theory was developed as a tool for rational decision-making. Its basic concepts were later used in evolutionary game theory to describe the evolution of behavioral phenotypes. In the hands of evolutionary biologists, this merger of game theory and population dynamics became an important tool for analysing frequency-dependent selection and social interaction.

I. Game theory

Game theory, as originally created by mathematicians and economists, addresses problems confronting decision makers with diverging interests (such as firms competing for a market, staff officers in opposing camps or players engaged in a parlor game). The 'players' have to choose between strategies whose payoff depends on their rivals' strategies. This interdependence leads to mutual outguessing (she thinks that I think that she thinks...). There usually is no solution which is unconditionally optimal, i.e., which maximizes a player's utility function, no matter what the co-players are doing. In contrast to such mutual dependence, monopolists can optimize their budget allocations without having to worry that others will anticipate their decisions. An optimization problem may be fraught with uncertainty, or computationally complex, but usually, what is meant by a

solution stands beyond doubt. In game theory, this need not be the case. Even in the simple game of 'matching pennies' (two players I and II choose independently between two alternatives, I wins if the two agree, and II if they differ), no outcome can leave both players satisfied.

A player can choose between alternative moves, or strategies. Since it is often useful to be unpredictable, a player may also choose a mixed strategy, i.e., opt with specific probabilities for this or that alternative. It can be shown that for any game, there exist at least one set of strategies (one for each player) which are best replies to each other (see Box 1). In this case, no player has an incentive to deviate from his or her strategy, as long as the other players stick to theirs. This defines a Nash equilibrium. (In the matching pennies game, both players have to choose with probability 1/2 between the two alternatives; as this example shows, Nash equilibria need not exist if mixed strategies are not admitted).

The notion of a Nash equilibrium satisfies a minimal consistency requirement for the 'solution' of a game (since otherwise, at least one player would deviate from it), but it presents a series of pitfalls. Consider, for instance, the following 'helping game', where two players have independently to decide whether or not to confer a benefit b to the other player, at a cost c to themselves. If $b > c$, they would both earn $b - c > 0$ by cooperating. But since it is better to defect, i.e., not to incur the cost, each player's best reply, irrespective of the other's decision, is to defect. The unique Nash equilibrium, in the helping game, is thus mutual defection. This game thus displays a 'social dilemma': the pursuit of self-interest is self-defeating. In other games, there exist several Nash equilibria, and the choice of the right can be a tricky issue. A

large part of classical game theory deals with equilibrium refinements and equilibrium selection.

II. Evolutionary game theory

In the context of evolutionary biology, the two central concepts of game theory, namely strategy and payoff, have to be re-interpreted. A strategy is not a deliberate plan of action, but an inheritable trait, for instance a behavioral program. Payoff is not given by a utility scale indicating subjective preferences, but by Darwinian fitness, i.e., average reproductive success. The 'players' are members of a population, competing for a larger share of descendants. If several variants of a trait occur in a population, then natural selection favors the variants conferring higher fitness. But if the success of the trait is frequency-dependent, then an increase of the frequency of variant may lead to a composition of the population for which other variants do better. Similar situations are studied in population ecology. Thus, if prey is abundant, predators increase for a while. But this increase reduces the abundance of prey, and therefore leads to a decrease of the predators. Evolutionary game theory can be viewed as the ecology of behavioral programs.

A classical example, which led Maynard Smith to develop evolutionary game theory, is provided by inner-specific contests. Assume that there are two behaviorally distinct types: 'Hawks' escalate the fight until the injury of one contestant settles the issue, whereas 'Doves' stick to some form of conventional display (a pushing match, for instance, where injuries are practically excluded), and give up as soon as the adversary escalates. If most contestants are 'Doves', 'Hawks' will be able to settle every

conflict in their favor, with a corresponding gain in fitness. Hence, 'Hawks' will spread. If most contestants are 'Hawks', however, then escalating a conflict will lead with probability one-half to injury. If the object of the fight is not worth the injury, then the 'Dove' trait will spread. No trait is unconditionally better than the other. 'Hawks' can only spread if their frequency is below G/C , where G is the value of the contested object and C the cost of an injury (both measured in terms of fitness). If their frequency is higher, it will diminish. Oversimplified as it is, this thought experiment shows that heavily armed species, for which the risk of injury is large, are particularly prone to conventional displays, i.e., ritual fighting. This fact had been observed empirically, but before the advent of evolutionary game theory, it was erroneously interpreted as benefitting the 'good for the species'.

[Place Fig. 1 near here]

A large number of behavioral traits, but also of morphological or physiological characters, such as the length of antlers, or the height of trees, are subject to frequency dependent selection. Trees invest considerable resources into growth, for instance, because neighboring trees do. To fall behind, in such an 'arms race', means to give up a place in the sun. Traits subject to frequency-dependent selection occur in many types of conflicts between two individuals, for instance concerning territorial disputes (between neighbors), division of parental investment (between male and female), or length of weaning period (between parents and offspring). Moreover, frequency-dependent selection also occurs without antagonistic encounters, as when individuals are 'playing the field'. The sex ratio is a well-studied example. In the simplest scenarios, the rule is simple: if the sex-ratio is biased

towards males, it pays to produce daughters, and vice versa. Under specific conditions, however, occurring with inbreeding or local competition for males, the sex-ratio may evolve away from 1:1. Other examples of frequency-dependent selection concern the dispersal rate among offspring, the readiness to emit an alarm-call, or the amount of time spent on the look-out for predators.

The evolution of cooperation is one of the best-studied chapters of evolutionary game theory. Traditionally, this is modeled by the helping game described above. If an individual is equally likely to be potential recipient or donor in a given encounter, then a population of cooperators would earn, on average, $b-c > 0$ per interaction, and be better off than a population of defectors earning 0. But an individual would always increase its fitness by refusing to help, and hence we should not see cooperation.

Game theorists have encapsulated this social dilemma in the Prisoner's Dilemma (PD) game. In this game, each player can choose between the two strategies C (to cooperate) and D (to defect). Two C players will get a reward R which is higher than the punishment P obtained by two D players. But a D player exploiting a C player obtains a payoff T (temptation to defect) which is higher than R, and this leaves the C player with the sucker's payoff S which is lower than P. A rational player will always play D, which is the better move no matter what the co-player is doing. Two rational players will each end up with payoff P instead of R (see Fig 2).

Many species engage in interactions which seem to be of Prisoner's Dilemma type. Vampire bats feed each other, monkeys engage in allogrooming, vervet monkeys utter alarm calls, birds join in anti-

predator behavior, which includes vigilance and mobbing, guppies and stickleback cooperate in predator inspection, hermaphroditic sea bass alternate as egg-spenders, many species of birds engage in nest helping, lions in cooperative hunting or joint territorial defense. It is difficult, however, to measure the lifetime fitness of free-living animals, and in many cases, it remains doubtful whether a given type of encounter is really of the Prisoner's dilemma type, i.e., satisfies the inequalities $T > R > P > S$. Some of the afore-mentioned examples could be instances of by-product mutualism, in which both players are best served by cooperating and none is tempted to defect. Other types of encounters (for instance, the Hawk-Dove game) may have the structure of a so-called Chicken game (with $T > R > S > P$), in which the best reply to the co-player's C is a D, but the best reply to a D is a C. In both cases, cooperation (at least by one partner) is no paradox.

There are several ways in which the Prisoner's dilemma can be overcome. In general, any form of associative interaction favors cooperation. Such association may be due to kinship, to partner choice, to the ostracism of defectors or simply to spatial structure and limited dispersal. Indeed, if players can only interact with their nearest neighbors, then clusters of cooperators can grow. This spatial aspect of game theory is likely to operate for many sessile organisms.

Moreover, if interactions of the Prisoner's dilemma type are repeated between the same two individuals, players can have the option to break up partnerships, or vary the amount of cooperation, depending on past experience. But even without these options, the strategy of always defecting is not invariably the best option in the Iterated Prisoner's dilemma (IPD game). If the probability of a further round is sufficiently high, then even a small amount of conditional cooperators suffices to favor cooperation. The best known example of such a

discriminating strategy is Tit For Tat (TFT). A TFT-player cooperates in the first round and from then on always repeats whatever the co-player did in the previous round (see Fig.3).

The best examples for reciprocation may be found in human societies. Among humans, moreover, reciprocation is often indirect. An act of assistance is returned, not by the recipient, but by a third party. A prerequisite is that players know enough about each other. This condition is likely to hold if groups are close-knit and individuals can exchange information about each other.

III. Game dynamics

The major new tool of evolutionary game theory consists in using population dynamics. This 'technology transfer' from population ecology relies on the assumption that successful traits spread. If there are only two possible types A and B, for instance, then essentially only three scenarios are possible, depending on whether a minority of one type can invade a resident population consisting of the other type only (see Fig.4):

(a) A can invade B but B cannot invade A. In this case, the dominant strategy A will always out-compete B. This happens with the Prisoner's dilemma, if A-players defect and B-players cooperate.

(b) A can invade B and B can invade A. This leads to the coexistence of both types in stable proportions as, for instance, if A are 'Hawks' and B are 'Doves'.

(c) no type can invade the other. This is a bi-stable situation; whoever exceeds a certain threshold will outcompete the other. This happens with the Iterated Prisoner's Dilemma if A is TFT and B always defects.

With three types A, B and C, the game dynamics becomes considerably more complex, in part because 'rock-paper-scissors'-cycles can occur: A is dominated by B, B by C, and C in turn by A. Several such situations have been documented. In cultures of *E. coli*, for instance, the wild type A can be superseded by a mutant strain B killing the competitors by producing colicin, which acts as a poison. Simultaneously, this mutation produces a protein conferring immunity against the poison to its bearer. A population of type B can be superseded by a further mutant type C which produces the immunity protein but not the colicin (since this poison is inefficient in a population consisting of types B and C). In turn, type C can be invaded and eliminated by type A. Another rock-paper-scissors cycle has been found among males of the lizard *Uta stantibus*. The three types correspond to inheritable male mating strategies. Type A forms no lasting bonds but looks for sneaky matings; type B lives monogamously and closely guards the female; and C guards a harem of several females, of course less closely.

Depending on the parameters, evolutionary models of rock-paper-scissors games either lead to the stable coexistence of all three strategies or to oscillations with increasing amplitude which lead to the recurrent elimination of the three types (see Fig.5). The competition of male lizards displays the former type of dynamics, and that of *E.coli* bacteria the latter.

With four or more types competing, game dynamics can become yet more complex. The frequencies of the different types can keep oscillating in a regular or chaotic fashion. In addition to the dynamics describing frequency-dependent selection among a given set of types, mutations can produce new types occasionally. This usually proceeds at another time scale. Evolutionary game theory allows studying both short-term and long-

term evolution. For the latter, it is often convenient to assume that the transient effects following a random mutation have settled down before the next mutation occurs. As long as the population consists of one type only, this leads to a trait substitution sequence: the fate of a mutant, i.e., its fixation or elimination, is settled before the next mutation occurs. The path of the corresponding 'adaptive dynamics' can lead to evolutionary stable states immune against further invasion (see ESS) or to 'branching points' where the population splits up and becomes polymorphic.

Game dynamics can also be used to analyze the interactions between different subpopulations (such as males and females, or territorial owners and intruders). A fast-growing branch of evolutionary game theory deals with structured populations: here, the assumption of random encounters is replaced by that of interaction networks.

Evolutionary game theory deals with phenotypes, and usually assumes that 'like begets like'. With sexual replication, however, this assumption can fail. Mendelian segregation, pleiotropy and sexual recombination can lead to situations where more successful types produce less successful variants. In principle, such features can be integrated into models of frequency dependent selection acting within the gene pool, but this can lead to intractable dynamics. Moreover, arguments from evolutionary game theory can fail, just like optimization arguments from adaptationism, due to genetic constraints. In the absence of specific information on the genotype-phenotype map, however, evolutionary game theory often provides an efficient heuristic tool for understanding frequency-dependent adaptation at the phenotypic level. Moreover, it also proved a suitable tool to describe social learning and cultural evolution.

See also the following articles:

Evolutionarily Stable Strategies

Evolution of Cooperation

Adaptive Dynamics

Glossary:

Strategy: Rule that describes how an individual acts in a given situation. For example, in an inner-specific contest, possible strategies are to fight or to flee.

Replicator dynamics: A model for the dynamics in evolutionary games. When a strategy fares better than the average then this strategy is expected to spread in the population.

Hawk Dove game: A prominent model for animal contests in evolutionary game theory. It is assumed that there are two types: 'Hawks' escalate a fight, in which case 'Doves' give up. When 'Hawks' are frequent it is better to be a 'Dove', in order to avoid serious injuries. Conversely, if the population consists of 'Doves', then escalating a fight pays off.

Prisoner's Dilemma: A famous game that describes the conflict between group-interest and self-interest. Two individuals may either cooperate (C) or defect (D). If both choose C, they are better off than if both choose D. However, individually each player prefers to defect, leading to a dilemma.

Payoff: Number that represents the success of a given strategy. In classical game theory, payoffs are described as utilities, whereas evolutionary game theory interprets the payoff of a strategy as its reproductive success.

Frequency dependent selection: when the reproductive success of an individual does not only depend on its own type, but also on the composition of the population. For example, if the sex-ratio in a population is biased towards females, then males have an advantage.

Nash equilibrium: A game is in equilibrium, if none of the players has an incentive to deviate from its strategy, as long as the other players stick to theirs.

Further Reading

Binmore, K. 2009 *Game Theory: A very short introduction*. Oxford: Oxford University Press

Cressman, R. 2003 *Evolutionary Dynamics and Extensive Form Games*, Cambridge, MA: MIT Press

Dugatkin, L.A. 1997. *Cooperation among animals: an evolutionary perspective*. Oxford: Oxford University Press

Hofbauer, J. & Sigmund, K. 1998. *Evolutionary Games and Population Dynamics*. Cambridge: Cambridge University Press

Fudenberg, D. & Levine K. 1998 *The Theory of Learning in Games*. Cambridge MA: MIT Press

Maynard Smith, J. 1982. *Evolution and the Theory of Games*. Cambridge: Cambridge University Press

Nowak, M. 2006. *Evolutionary Dynamics*. Cambridge, MA: Harvard University Press

Sigmund, K. 2010 *The Calculus of Selfishness*, Princeton NJ: Princeton University Press

Weibull, J.W. 1995. *Evolutionary Game Theory*. Cambridge, MA: MIT Press

Box 1:

A game between two players I and II can be described by its normal form, which consists of a list of all the strategies e_1, \dots, e_n and f_1, \dots, f_m available to player I and player II, respectively, and of their payoff values a_{ij} resp. b_{ij} obtained when I plays e_i and II plays f_j . A mixed strategy for player I is given by the vector x of the probabilities x_i to use e_i . Since $x_1 + \dots + x_n = 1$, the vector $x = (x_1, \dots, x_n)$ is an element of the unit simplex S_n spanned by the vectors of the standard basis in R^n , i.e., the vectors with $x_i = 1$ and $x_j = 0$ for $j \neq i$, which correspond to the pure strategies e_i .

If player I uses strategy x and player II uses y , then the payoff for the former is given by the sum of the terms $a_{ij}x_iy_j$, summed over all i and j , and the payoff for the latter by $b_{ij}x_iy_j$. We denote these terms by xAy and xBy , respectively.

The strategy x is said to be a best reply to strategy y if $xAy \geq zAy$ holds for all z in S_n . In this case, player I cannot expect any gain from using a strategy different from x . Similarly, y is a best reply to x if $xBy \geq xBw$ for all w in S_m . A pair of strategies (x, y) is said to be in Nash equilibrium if both conditions are satisfied, i.e., if each strategy is a best reply to the other. In this case, both players have no incentive to deviate unilaterally from their strategy. In the special

case of a zero sum game (i.e., when $a_{ij} = -b_{ij}$ holds for all i and j), these strategies are maximin strategies, i.e., each maximizes the minimal payoff and thus guarantees the best security level.

One speaks of a symmetric game if the players have the same sets of strategies and payoff values and thus cannot be distinguished. Formally, this means that $a_{ij} = b_{ji}$ holds for all i and j . In this case, a strategy x is said to be a Nash equilibrium if the symmetric pair (x, x) is a Nash equilibrium pair, i.e., if $zAx \leq xAx$ for all z in S_m .

Box 2

In the simplest formal setup for evolutionary game theory, the e_1 to e_n correspond to different types of individuals in a large, well mixed populations, and the x_i are their relative frequencies (thus, the state of the population is given by x in S_n). The game is assumed to be symmetric. Since an individual of type e_i randomly meets an e_j -individual with probability x_j , and obtains payoff a_{ij} from the interaction, the average payoff for e_i -players is given by $(Ax)_i = a_{i1}x_1 + \dots + a_{in}x_n$, and the average payoff in the population by xAx . The frequencies x_i evolve as a function of time t , according to their success. If one assumes that the per capita growth rate of type e_i is given by the difference between its payoff and the average payoff in the population, one obtains the replicator equation $\frac{dx_i}{dt} = x_i[(Ax)_i - xAx]$ on the state space S_n . Every Nash equilibrium is a fixed point of the replicator equation, and every stable

fixed point is a Nash equilibrium, but the converse statements need not hold.

Figure Captions:

Fig. 1: Payoffs for the Hawk-Dove game: If a hawk encounters another hawk, there is an equal chance to win the contest or to get injured, resulting in an expected payoff of $(G-C)/2$. Against doves, a hawk always comes off as the winner, leading to a safe payoff of G . The payoffs for doves are derived analogously.

Fig. 2: Payoffs for the Prisoner's Dilemma (with $T > R > P > S$): Irrespective of the opponent's strategy, it is always better to defect, since $T > R$ and $P > S$. If both players follow this logic they end up with payoff P instead of R .

Fig. 3: Payoffs for the Iterated Prisoner's Dilemma (IPD): When a TFT player meets a co-player of the same type, both will cooperate mutually, leading to an average payoff of R . Against a co-player who defects always (All D), a TFT player stops cooperating after the first round and plays D subsequently. If the number of rounds is random and the probability of a further round is w , this results in the payoffs displayed in the matrix.

Fig. 4: Different scenarios for the evolutionary dynamics between two strategies:

(a) Dominance: The blue strategy always out-competes red. Evolution leads to the state in which every individual adopts blue.

(b) Coexistence: Red invades blue and blue invades red. Eventually, there is a stable coexistence of both strategies.

(c) Bi-Stability: Both, red and blue are stable. The eventual outcome depends on the initial population.

Fig. 5: Dynamics of the rock-paper-scissors game. Paper beats rock, scissors beats paper and rock beats scissors. Depending on the exact payoff values, this may either result in closed cycles (left), a stable

coexistence of all strategies (middle) or never-ending oscillations
(right).

Figure 1:

		Type of the opponent	
		Hawk	Dove
Type of the player	Hawk-Dove		
	Hawk	$(G-C)/2$	G
Dove	0	$G/2$	

Figure 2:

		Type of the opponent	
		Cooperator	Defector
Type of the player	Prisoner's Dilemma		
	Cooperator	R	S
Defector	T	P	

Figure 3:

		Type of the opponent	
		TFT	All D
Type of the player	IPD		
	TFT	R	$(1-w)S+wP$
All D	$(1-w)T+wP$	P	

Figure 4




Figure 5

