PETER HAMMERSTEIN

1

Games and markets: economic behaviour in humans and other animals

1.1 Introduction

Economists think they know how humans ought to behave if only they were smart enough. Biologists have some knowledge of how animals actually do behave. It does not seem a feasible question to ask how animals ought to behave. Yet, there is a conceptual link between normative economic theory and its empirical biological counterpart. Darwinian evolution often creates animal traits that look to an observer as if the animal did care about the economist's advice. Therefore, economic analysis of animal behaviour has become a flourishing field of biological research in which games and markets play an important role. This chapter discusses fundamental concepts in human and animal economics. They are illustrated with examples from both disciplines. Furthermore, it is shown that the facts often do not meet theoretical expectations. Some hints are given as to why this may be so. Theory development in animal and human economics is far from being completed.

1.2 Evolutionary adaptation and bounded rationality: are animals better economists than humans?

Classical economic theory based most of its thoughts on the idea that decision makers are rational in the sense of maximising subjective expected utility. Savage (1954) axiomatised this Bayesian approach to decision making and Harsanyi & Selten (1988) explored it further in the game theoretic context of strategic interaction. However, the more the economic notion of rationality had been made precise, the less it seemed adequately to reflect properties of the real world. Rational decision makers are assumed to possess unlimited cognitive and computational power and to solve correctly every mathematical problem in zero time at no cost.

Furthermore, they are supposed to have clear, invariant preferences that can be described by a utility function. In contrast, we know from research in psychology and experimental economics that (a) human preferences are far from being stable and (b) most decisions are made without the help of sophisticated mathematics. If mathematics is used at all, it is often used incorrectly. In particular, humans in their everyday life seem to avoid dealing with probabilities and are poor statisticians from a Bayesian point of view (Kahneman *et al.* 1982).

To illustrate this, let us consider a phenomenon called 'the winner's curse' (Thaler 1988). This phenomenon was found to occur in real-life auctions of oil-fields (Capen *et al.* 1971; Levinson 1987). Businessmen who overestimated the value of an oil field most were the ones who received it in the end. As a result, many of them ended up with an unprofitable deal – the winner's curse. It is possible to design a simple experiment in which a similar phenomenon occurs even if only one bidder is involved. The experimental scenario is as follows. A dealer has a car that he wants to sell. He knows x which denotes what the car is worth to him. The car is worth 50% more to a potential buyer. However, the buyer does not know x. He only knows that x is uniformly distributed between 0 and 1. At the beginning of the experiment, a random number generator is used to determine x. The buyer has one bid y and the dealer, played by a computer, will sell if y is greater than x. Therefore, if $x<y$, the object will be traded and the buyer's net payoff is $1.5x-y$, otherwise the payoff is zero.

How much should a buyer offer in this experimental setting? When faced with the decision problem in a single-shot game without any possibility for learning, bids near 0.5 are frequently observed (Ball *et al.* 1991). At first sight this may look like a reasonable bidding decision. Bidding 0.5 seems to imply a positive expected payoff of 0.25 because the average value of the object to the buyer is 0.75. This argument, however, contains a fatal error in that it overlooks the *conditional expectation* that really matters in the scenario under discussion. One has to take into account that the trade only takes place if $x<y$. The expected value of x *given that* $x<y$ is $0.5y$ (as opposed to 0.5). Therefore, in an actual trade, the expected value of the object to the buyer is only as small as $0.75y$. His expected net payoff then is $-0.25y$. On average the buyer, therefore, makes losses if he comes up with any positive bid. This is the winner's curse and the only rational way of playing the game is to bid zero.

Why do so many humans fail to avoid the winner's curse? We have just seen that in order to foresee the winner's curse, a decision maker would have to calculate his payoff as a conditional expectation. If not trained by mathematicians, humans do not seem to have the cognitive skills to correctly deal with conditional expectations in this and many other experiments. Even if they have a chance to learn their bidding strategy in the scenario described above, they typically do not arrive at the rational solution. This was shown by Selten (1996) in a version of the bidding experiment where subjects were exposed a hundred times to the buyer's situation. Selten also offers an explanation of the phenomenon. In order to see why learning is *not* a powerful mechanism to overcome the winner's curse, consider a person who shifts the bidding tendency according to last experience. Let this person start with the bid $y = 0.5$. With probability 0.5 the person will not get the object and may simply classify the situation as 'I bid too low'. With the same probability he will get the object and reach the opposite conclusion 'I bid too high'. If learning consists of a simple shift of bidding tendencies according to this classification of situations, it will probably not take the learning individual anywhere near the rational solution. This is an example of what Selten (1996) calls 'learning direction theory' and it may to some extent explain the facts.

To summarise the experimental results, both learning and higher cognitive abilities fail to generate economic rationality in the winner's curse example. Let us now conduct the following thought experiment. Suppose that animals of a given species have to act generation after generation in a situation resembling the buyer's situation of the winner's curse problem. Suppose that the animal's payoff is an increment (or decrement in case of a negative sign) of its expected reproductive success (fitness). Assume further that the bid is coded for by genes and genetically transmitted to the next generation. New bids only arise by rare mutations and the animals do not even try to analyse the bidding situation by any cognitive means. They act instead as if they were robots controlled by an inherited computer program. How is the program going to evolve? In the long run, the evolutionary process of mutation and selection will take our genetically coded bid to zero. This means that the evolutionary process itself (not the animal!) generates a behavioural feature that would look to an observer *as if* it had been brought about by rational decision making. We often find such quasi-rationality in models of Darwinian adaptation. The bidding example is particularly interesting because it shows that evolution can 'do better' than both human learning and time-limited cognition. However, in principle it is possible for learning to also outcompete evolution. The reason is that evolution has no anticipatory power, whereas learning combined with some cognitive skills can include elements of anticipation.

1.3 Games and Nash equilibria in economics and biology

Classical game theory studies two or more individuals who interact and whose payoffs depend on how everybody in the game behaves. In such an interactive situation, the rational decision maker has to solve a particularly intricate problem, namely what to expect about other players? Are they also rational and if not what are they like? Game theory in economics has eliminated this problem by ignoring it in an elegant way. It is assumed that all players know that all know that all are rational. Needless to say that this assumption of an entirely rational world made it possible for generations of game theorists to earn their living with writing papers, but it did not bring game theory any closer to reality. The assumption of omnipresent rationality implies that the players in a game will play a so-called Nash equilibrium. To see how this is defined, consider an n-person game. For every player i choose one of his strategies and call it p_i. Looking at all players simultaneously one then obtains the strategy combination $p = (p_1, p_2, \ldots, p_i, \ldots, p_n)$. This combination p is called a Nash equilibrium if it has the following property for every player i: Player i gets maximum expected payoff by playing p_i (compared to his other strategies) given that everybody else plays his strategy of p. In other words, in a Nash equilibrium every player's strategy is a best response to the joint strategic action of all other players. If a player assumes the others to play according to a Nash equilibrium, then he has no incentive to deviate from this equilibrium himself. In this sense a Nash equilibrium is self-enforcing. However, the problem is that (a) playing a strategy of a Nash equilibrium can turn into a payoff disaster if at least one other player is not rational, and (b) there often exist several Nash equilibria.

Let us look at the example of a coordination game in which cooperative exploitation of a resource is possible but may not be achieved. Suppose that 10 persons are asked independently, in separate rooms without seeing one another, to put money in an envelope. The experimenter collects the envelopes and empties their contents into a large box. He counts the total sum collected in the box and checks whether it amounts to at least 100 monetary units. If not, he keeps the money for himself and all subjects have lost what they put into the envelope. However, if the total in the box is 100 or more, the experimenter adds 50 units to this amount and distributes the enriched contents of the box equally among the subjects, no matter what an individual's original contribution had been. Searching for symmetric Nash equilibria (where everybody gives the same amount) we find two alternative solutions: one in which everybody puts 10 units into the envelope and one in which all envelopes are empty. In the first Nash equilibrium, the subjects cooperate and exploit the experimenter. Everybody has a net gain of 5 units. In the second Nash equilibrium, no cooperation takes place and the 'resource' (i.e. the experimenter) remains unexploited. How do these solutions compare? The first one is risky in the following sense: if only one player gives a little less than 10, everybody else has a net loss of 10 units

The Nash equilibrium was originally created as a tool for analysing human behaviour. However, in theoretical evolutionary biology Nash equilibria have an even better foundation than in classical economics. At first sight this appears like a silly statement because animals are not supposed to be rational, and it would be absurd to assume that all animals know that all animals know that all of them are rational. Yet, if animals play a game generation after generation and if strategies are inherited, then natural selection will often drive these strategies towards a Nash equilibrium. Suppose that after evolving for quite some time in the same environment, a population has reached a situation in which the current strategy is maintained by the forces of natural selection. This means that no mutant strategy should be able to invade the population if it initially occurs at a low frequency. John Maynard Smith (1982) called a strategy with this uninvasibility property an evolutionarily stable strategy (ESS). If one studies ESSs in selection models with asexual reproduction, exact inheritance and encounter rates of strategies proportional to their population frequencies, then it can be shown very easily that an ESS will necessarily be a Nash equilibrium (e.g. Hammerstein & Selten 1994).

To some extent, John Nash himself had foreshadowed this result in an unpublished part of his Ph. D. thesis where he briefly talks about the 'mass action interpretation' of game-theoretic equilibrium (Nash 1950). We only know this because, decades after Nash got his Ph. D., the Nobel award made people curious to have a look at his thesis. Historically speaking, the field of evolutionary game theory was initiated by Maynard Smith & Price (1973). Much later, economists developed their own evolutionary game theory (e.g. Weibull 1995; Samuelson 1997). Hofbauer & Sigmund (1998) review evolutionary game theory from a dynamical systems point of view. Hammerstein (1996) discusses the links between population genetics and evolutionary game theory. He compares the phenotypic trajectory of an evolving population with the course of a streetcar. There are temporary stops of the streetcar in which the logic of phenotypic game theory is blurred by genetic constraints. These stops do not correspond to a Nash equilibrium. However, evolution is capable of removing genetic constraints. After successive elimination of such constraints, the streetcar will often reach a final stop that can be characterised (in the sense of a necessary condition) by a Nash equilibrium.

Consider now a fictitious biological example of an evolutionary game that resembles a little the above coordination game. Animals usually do not put money into envelopes, but imagine instead a situation where 10 genetically unrelated predators would be able to kill a large prey item by jointly injecting a large amount of toxin. Suppose that the minimum lethal dose of the toxin is 100 mg and that the energy needed to produce 1 mg of the toxin is 1 joule. Assume further that by communal consumption of the prey each predator would have an energy gain of 15 joules. Let an individual's fitness depend linearly on the difference between energy gain from prey consumption and energy loss from toxin injection (the energetic cost of producing the toxin).

The biological coordination game has two symmetric Nash equilibria. In the cooperative equilibrium everybody injects 10 mg and has a net return of 5 joules. In the non-cooperative equilibrium everybody injects a zero amount of toxin. We can easily convince ourselves that both the 10 mg strategy and the 0 mg strategy are evolutionarily stable. In a population playing the cooperative strategy, a mutant injecting less or more than 10 mg would always have less than the return of 5 joules. In contrast, the average return of cooperators would be 5 joules, at least in an infinitely large population. We conclude that it is an ESS to inject 10 mg. Similarly we can show that zero mg is also an ESS.

What do we learn about animal cooperation in this fictitious biological context? On the one hand, a population of cooperators would be maintained by natural selection. On the other hand, a population of non-cooperators would also be maintained by selection. Therefore, it is difficult to get an evolutionary transition from non-cooperation to cooperation and we may never see the latter evolve. This transition problem occurs in many other biological scenarios. For example, in the context of repeated interactions (see section 5 of this chapter), 'reciprocal altruism' or the strategy 'Tit-for-Tat' cannot easily evolve because of the transition problem.

In the above example of predators injecting a toxin, cooperation is biologically plausible if one assumes the existence of some peculiar historical path that led the population of predators towards the cooperative state with its evolutionary maintenance property. There are other games where cooperation, though highly beneficial to all cooperators, would not be maintained. Needless to say that the Prisoner's Dilemma has this property (Luce & Raiffa 1957). Most readers probably are saturated with discussions of this game which has been 'on the charts' for too long. So, let us be quick. In the Prisoner's Dilemma, there are two strategies, namely 'cooperate' and 'defect'. Both players would be very successful if they both defected. However, due to a payoff incentive for unilaterally exploiting a cooperator, defection is a strictly dominant strategy. The term dominating strategy means that by playing this strategy one always receives a higher payoff than by playing strategic alternatives, regardless of how the other player behaves. Obviously, the only Nash equilibrium, and the only ESS, is to defect.

In experiments of the Prisoner's Dilemma, humans often play the cooperative strategy instead of the Nash equilibrium. After all, our intuition is adapted to a highly social context where defection might have negative consequences outside the narrow setting of the game. However, the term 'tragedy of the commons' would not exist if humans always cooperated in games with the flavour of a Prisoner's Dilemma. Ostrom (1990, this volume) analysed how human communities found ways to resolve the tragedy of the commons. She describes various human policing strategies by which defectors are punished. However, there is a twofold tragedy which unfolds (Hammerstein 1996). Individuals often have an incentive to save the effort that would go into policing. Splitting the cost of policing is yet another tragedy of the commons.

1.4 Games in which players react to what has been done before

In most real-life games, the players act one after another. Simple game theoretical models in normal form – with conflict being represented by a payoff matrix – do not explicitly depict sequences of actions. In contrast, the theory of games in extensive form – where conflict is represented by a decision tree – was especially designed to give an adequate mathematical picture of the course of actions in a game. This makes it possible to analyse more deeply how players should react to one another.

One of the simplest games in which at least one player can react to what has been done before is the ultimatum game. Suppose that the Nobel committee has introduced a new Nobel prize in biology but cannot really afford to pay any substantial amount of money to the winners. This year, the prize goes to evolutionary biology. Let us assume there to be two fictitious winners who share the prize, say Richard and Steven. In order to deal with the shortage of money, Stockholm invents the following rule. Richard will be asked about the percentage he claims of the total award. After he has made up his mind, his decision will be communicated to Steven who may or may not accept the deal. If he rejects, both winners will only get the Nobel dinner and no cash. The total award is $1 million.

What are the Nash equilibria? One of them can be described as follows. Steven rejects unless he gets at least 50%. Richard offers 50%. Both are playing a best response to each other's strategies. Why is this so? As long as both play these strategies, Richard gets half a million. If he goes for more than 50%, his reward will be zero. If he goes for less than 50%, his reward will be less than half a million. So, we clearly understand that Richard's strategy is a best response to Steven's strategy. Conversely, if Richard asks for 50%, Steven cannot improve his payoff by shifting the acceptance threshold. If he raises this threshold, he will get nothing, if he lowers it, he will get no more than half a million. In other words, Steven is playing a best response to Richard's strategy. Now, the fact that both strategies are best responses to each other means that we are talking about a Nash equilibrium.

In real life, the equal-split solution would often occur in an ultimatum game. If the total is something like a dollar, it would be very risky to ask for more than 60%. However, what about the concept of economic rationality? The Nash equilibrium in which Richard claims 50% looks, on the one hand, quite reasonable because both players optimise against each other. However, Steven is not considered to be entirely rational if he is assumed to reject any split that treats him in an unfair way. Suppose Richard demands 80% for himself. By rejecting this, Steven is 'trashing' $200 000 that he otherwise could have brought back home. Would his wife say: 'how rational you are in playing games'? Certainly not. This is my caricature of a discovery by Selten for which he actually got the Nobel prize. He maintained that if players are considered to be rational then they have to be rational in *all situations* of the game, not only before the game starts. Selten (1975) coined the term *subgame perfection* in order to make this idea precise. He looked at substructures of a game that would also qualify as a game. These are called subgames. Obviously, strategies of the original game specify what a player would do in a subgame. Furthermore, a Nash equilibrium specifies a combination of subgame strategies so that one can ask the following question about the equilibrium's consistency with rationality assumptions: will the Nash equilibrium under consideration induce Nash equilibria for all subgames? If the answer is yes, the players remain rational while playing the game. In this case, the Nash equilibrium is called subgame perfect. Many Nash equilibria fail to have this property. For illustration, let us reconsider the fair Nash equilibrium of the 'Richard and Steven' game which is to 'claim 50% in Richard's role and accept a share of 50% or more in Steven's role'. We have already looked at a subgame with only one player, namely Steven, being informed that Richard demands 80% for himself. In the fair equilibrium, Steven rejects the remaining 20% and both players end up with zero payoff. Subgame perfection, however, requires for Steven to make an optimal decision *in this situation*. Obviously he then must accept the unfair deal.

Where does subgame perfection matter in biology? Consider the parental investment game between the sexes. A male and a female have copulated and the offspring needs some care. Who will provide the care? Suppose we are studying an animal example in which both male and female would have an advantage from leaving their partner alone with the offspring. An assumption of this kind can be justified if the time and energy gained by mate desertion has a positive effect on the deserting individual's future reproductive success. The structure of the mate desertion game can be similar to the 'Stockholm' ultimatum game. We begin to look at the interaction after a copulation has taken place. The male is in a role similar to Richard and can either decide to be unfair, by running away, or to be fair and stay in order to help. If he runs away, the female finds herself in a situation similar to that of Steven. She is associated with the fertilised eggs and may either 'trash' them or accept the unfair deal and care for them. In a subgame perfect equilibrium, depending on the parameters, she may have to accept the unfair deal very much like Steven has to accept it in the Stockholm ultimatum game. If subgame perfection really mattered in evolutionary biology, it would help us to understand why in many animal species with uniparental care it is the female and not the male who looks after the offspring.

Does subgame perfection indeed matter? Let us play devil's advocate and discuss the particular Nash equilibrium of the mate desertion game which is not subgame perfect, and in which the female is prepared to abandon the fertilised eggs if left alone by her mate. Her mate stays and cares. Is this an ESS? Note that, if the sexes behave like this, the female is never left alone with the offspring. However, in reality there are external causes of why females might be left alone with the offspring. Let us make the realistic assumption that occasionally a male is killed by a predator, or lightning, or another environmental effect. In such a case, the female will be left behind 'holding the babies'. If this happens frequently enough, compensatory care will evolve. This means, however, that she cares when he is gone. Deliberate male mate desertion will then evolve. In other words, an evolutionary transition takes place towards the subgame perfect equilibrium, imposing the burden of parental care upon the female. We conclude that the concept of subgame perfection has a strong biological appeal in the context of parental investment between the sexes.

1.5 Repeated games

Since the early days of game theory (Luce & Raiffa 1957) it is well known that in a repeated game more cooperation is typically possible than in the single-shot case where the game is played only once. A repeated game, or supergame, consists of a series of interactions between the same two players. In each interaction they play the *same* game, which is the building block of the supergame. The intuition behind cooperation in repeated games is simple. A player can act as a conditional cooperator, relating his behaviour to how 'well behaved' the other player was in previous rounds of the same game. Unlike popular belief in biology, this insight has little to do with the Prisoner's Dilemma because it holds, roughly speaking, for any game used as the building block. Mathematical results capturing the idea of cooperation in supergames are known as 'folk theorems' (e.g. Fudenberg & Tirole 1991). They are called folk theorems because 'all the folks in game theory' knew them for a long time and it would be difficult to pin down the exact origin of ideas. Axelrod (1984) did not create the classical theory of cooperation in repeated games but he succeeded extremely well in making the subject known to scientists outside the field of game theory.

Trivers (1971) was the first to realise that the theory of repeated games might be of great importance to biology. He caused many biologists to start an empirical search for 'reciprocal altruism' in animal behaviour. Axelrod & Hamilton (1981) initiated a similar search for animal strategies resembling 'Tit-for-Tat'. They also discuss some of the theoretical problems that arise in evolutionary theories of repeated games. Looking back at the three decades of research since Trivers advocated the biological supergame, we observe a considerable discrepancy between how excited behavioural biologists were about the subject, and how little evidence