Game Theory - Can We Play?

Carol Boate - Junior Sophister

The effect of agent interactions can all too often lead standard economic models into difficulties. The existence of 'cheats' in the 'economic game' *is* plausible. Carol Boate leads us through an analysis of Game Theory and goes on to develop some interesting twists in the story.

Game theory provides a unifying framework with which to analyse any question involving the behaviour of rational decision makers whose decisions affect each other. One of the most researched areas is that of oligopoly theory. "*Until game theory came along, most economists assumed that firms could ignore the effects of their behaviour on the actions of others*." (1)

This discussion will concern itself exclusively with non-co-operative game theory in oligopoly. In non-co-operative game theory, individuals cannot make binding agreements and the unit of analysis is the individual, who is concerned with doing as well as possible for (him)herself, subject to clearly defined rules and possibilities. The essay will also examine the usefulness of game theory in the analysis and understanding of oligopoly. In order to do this we require a standard by which to judge its usefulness. Kreps (1990) suggests that a useful theory is one which helps us to understand and predict behaviour in concrete economic situations. Hence studying the interactions of ideally rational firms should aim to help our understanding of the behaviour of real firms in real economic situations.

The Theory

The methodology of game theory consists of taking an economic problem, formulating it as a game, finding the game theoretic solution and translating the solution back into economic terms. So how do we formulate an economic problem as a game?

A game is defined by a number of players each facing a set of possible actions, called strategies, and consequent payoffs. The games in oligopoly theory involve some mutuality of interests and so we will only concern ourselves with the solution of non-zero-sum games. If each player in the game has the same information and each knows that the others have this information, then the game is one of common knowledge. If one player has private information, or even if they think that others might, then we have an asymmetric information game. Such games are much more complex to analyse and so the focus here will be on common knowledge games. A critical piece of common knowledge is that all players are rational, in the sense that when presented with the same information they will come to the same conclusion: in particular, they can duplicate each other's thought processes.

Now, how do we find the game theoretic solution? 'Nash equilibrium' (NE) is the most frequently applied solution concept in economic examples. A NE occurs when each player chooses his optimal strategy given their rivals' (expected) choice. A NE is also known as a situation of 'no regret'. That is, neither player has an incentive to deviate from their chosen strategy. For example, consider the payoff matrix below:

	Player B		
		Left	Right
Player A	Top	3, 2	0, 0
	Bottom	0, 2	1, 2

The outcome (top, left) is a NE: if player A chooses 'top', B will choose 'left', if B chooses 'left', A will choose 'top'. There is no incentive for A to switch to 'bottom' or for B to switch to 'right'.

Though there is some debate over the validity of Nash Equilibrium, it does have the useful property that there exists at least one NE for every game. In the example above, the players chose a 'pure strategy' and had to stick to it forever. If a NE does not exist in such a game - one can always be found using 'mixed strategies'. That is, if we allow the players to randomise their strategies - assign probabilities to their possible strategies and then play them according to those probabilities - then it can be shown that a NE may be found.

Applications to Oligopoly Theory

The three main models of oligopoly will be discussed here, restricting our attention to duopoly for the sake of simplicity. We assume that the firms are producing a homogeneous product.

Cournot's Analysis

Cournot proposed a model whereby firms compete against each other in terms of the quantities they simultaneously choose to produce (q1,q2). Their respective payoff functions are their respective profit functions: each firm's payoff, and thus each firm's strategy, depends on their rival's strategy. Cournot showed that there exists a unique solution of this model when the demand functions are linear. For example:

given p = a - q

and assuming constant marginal cost, one can derive the following equations:

q1=(a-q2-c)/2 (R1)

q2=(a-q1-c)/2 (R2)

R1 and R2 are called the reaction functions for firms 1 and 2 respectively. They show the optimal level of output for each firm given their rival's (expected) level of output. While the term reaction function may be misleading, since the firms act simultaneously with no chance to reply or react, they are useful in imagining what each firm would do if the rules of the game allowed the firm to move second. Plotting these reaction functions:

Figure 1:

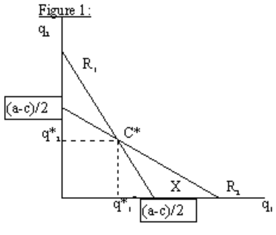

Where the functions intersect, point c* is the Cournot NE. Algebraically, it is found by solving the reaction functions which, in this case, yields the result (q1*=(a-c)/3 , q2*=(a-c)/3).

In the Cournot game, the NE has the desirable property of stability: we can imagine how starting from some other strategy profile the players might reach the equilibrium profile. If the initial profile is point X above, for example, firm1's best response is to decrease q1, (q1), and firm 2's response is to increase q2(q2), which moves the profile closer to the equilibrium, c*. However, we might still be dissatisfied with the Cournot NE. The main objection is that the strategy sets are specified in terms of quantities when, in reality, firms often set their prices and then sell as much as they can at those prices.

Bertrand's Analysis

Bertrand realised that the focus on quantity competition was unrealistic. He proposed an alternative model whereby firms simultaneously choose prices (p1,p2) based on their expectations of each others' prices. In our example with linear demand and constant marginal cost, the payoff function for firm 1, (and analogously for firm 2) is:

$$H1= \begin{cases} p1(a-q1)-cq1 & \text{if } p1<p2 \\ (p1(a-q1)-cq1)/2 & \text{if } p1=p2 \\ 0 & \text{if } p1>p2 \end{cases}$$

And so each firm always has an incentive to cut their price to just below their rival's as this would mean capturing the entire market. The Bertrand NE must therefore be p1=p2=c. Then, for either firm: a higher price would mean zero market share and to charge a lower price would mean losses.

This is the Bertrand paradox. We know that in reality firms do compete in prices and make profits. What it tells us is that duopoly profits do not arise simply because there are two firms. These profits arise from something else such as multiple periods, incomplete information or differentiated products.

Stackleberg's Analysis

Stackleberg proposed a basic derivative of the Cournot model: the leader-follower model. Firm 1 chooses what quantity it wishes to produce, q1, and after this choice is communicated to firm 2, q2 is chosen. Firm 1 is the Stackleberg leader, firm 2 is the Stackleberg follower.

In our example, the Stackleberg NE, S, is as shown in figure 2 below.

Figure 2

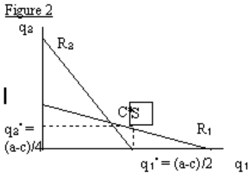

All three models can be seen as the application of the NE concept to games which differ in their strategic variables and timing. The Cournot and Bertrand models are simultaneous moves games where the strategic variables are quantities and prices respectively. The Stackleberg model is a sequential game where quantities are chosen. We can conclude from this that the outcome of an oligopoly situation is very sensitive to the details of the model.

Multiple Periods

All of the above models were one-shot games. In reality, firms compete against each other repeatedly, and the outcome is not always as these models predict. In the example of the Cournot model (a similar analysis may be carried out for Bertrand), if the firms could co-operate with each other, they would produce the monopoly level of output, q* = (a-c)/2. Thus they would be maximising the total payoffs and would both receive higher profits. So, q1 = q2 = (a-c)/4 is the pareto optimal outcome from the firms' perspective.

However, since we have excluded the possibility of making binding contracts there is always an incentive for each firm to cheat. That is, say firm 1 believes that firm 2 will 'co-operate' and produce q2* = (a-c)/4 [=1/2Hm], then firm 1 could 'cheat' by producing a quantity consistent with its reaction curve [>1/2Hm].

The 'dilemma' can be represented by the following payoff matrix.

		Firm 2	
		Co-operate	Cheat
Firm 1	Co-operate	1/2Hm, 1/2Hm,	<Hc, >1/2Hm
	Cheat	>1/2Hm, <Hc	Hc, Hc

(where 1/2Hm>Hc)

(Hc, Hc) is a NE, and so the outcome is (cheat, cheat) and Cournot NE profits are made by both firms. The only way out of this dilemma, usually called a 'prisoner's' dilemma, is if a firm can punish another for cheating. This requires that the game be repeated by both players. Reinhard Selten found that whether or not firms co-operate in repeated games depends greatly on whether or not there is a known fixed end to the repeating.

A Finite Number of Plays

This situation is best analysed through 'backward induction'. That is, we should look ahead to the final play. At this last play there is no future to consider and so, this play is simply a one-shot prisoner's dilemma game and the outcome is (cheat, cheat). Since this outcome is a foregone conclusion, we must turn our attention to the penultimate play, as this effectively becomes the last play. The same logic applies here and once again the outcome is (cheat, cheat). This argument unwinds all the way back so that there is no co-operation, even in the first play. While this logic is impeccable, the argument is severely criticised due to the many real world examples of successful co-operation in such situations. One explanation for this is that the real gains to be made by co-operating for a while may be very high. Thus there may be an initial phase of mutually beneficial co-operation while each side is waiting to take advantage of the other.(p.101)

Indefinitely Repeated Games

Consider a situation where two firms start off on a basis of trust. If one firm cheats now, there are immediate gains to be made (>1/2Hm - 1/2H) followed by losses (1/2Hm -H c) while the trust has collapsed. If the value of these gains exceeds the value to the firm of the subsequent losses, then this firm will cheat. Of course time is money and the value of the future losses will depend on how much weight the firm places on the future. Therefore we must use a discounting technique. In our example, the two alternative profit streams would be:

Co-operate = 1/2Hm + a(1/2 Hm) + a2(1/2Hm) +...

Cheat = >1/2Hm + a(Hc) + a2(Hc) + ...

Thus the decision of whether or not to cheat will depend on the magnitude of Hm, Hc, >1/2Hm and 'a' the discount factor. As the game is repeated indefinitely the decision to cheat now is identical to the decision to cheat in any subsequent period. Thus, given that each firm faces the same Hm, Hc, >1/2Hm and has knowledge of each other's discount factor, the outcome for both firms will be to co-operate forever, or cheat forever.

Once again reality does not conform to the theory and leads us to question the assumption that each firm will punish the other's cheating by playing 'cheat' forever. For example, what if firm 2 cheated once and then made a (credible) promise never to cheat again - regardless of the economic gains to be made. Surely it would be foolish, and more importantly irrational, of firm 1 to play a strategy of 'punish forever'? A more appropriate strategy for any firm to adopt would encourage co-operation while avoiding the possibility of exploitation.

Dixit and Nalebuff identify three criteria which a punishment strategy ought to meet(p.105):

1) it should be easy for one's rivals to calculate the costs associated with the punishment.

2) the rival must perceive the punishment to be certain - that is, the threat of punishment must be credible.

3) the punishment ought to be severe enough to deter cheating but not too severe in the case of mistakes or misperceptions. One strategy which meets these criteria is the famous 'tit-for-tat' strategy. This strategy co-operates in the first period and from then on mimics the actions of its rival in previous rounds. Robert Axelrod conducted a computerised tournament of game theorists' punishment strategies, which tit-for-tat won outright. At best, it ties its rivals, at worse, it is exploited only once. Despite this encouraging evidence, tit-for-tat is still disliked by game-theorists since once taken outside the realm of economic models, the slightest possibility of misperceptions can lead to disastrous results for the players.(p.108)

Thus we must return to our criteria for a good punishment strategy. Most game theorists focus on the second criteria, that of credibility. By making it impossible for the firm to reverse its decision, the firm can commit itself to its threat of punishment and so encourage co-operation. Alternatively, if the firm could change the payoffs of the 'game' they could make it more costly for their rivals to cheat and this too would promote co-operation.

A well known way of guaranteeing punishment is through, 'beat-the-competition' advertisements. For example, if our two firms compete on prices and firm 2 is running a campaign which reads: 'If within X days of purchase you find the same product on sale at a lower price, we will refund you double the difference, subject to written proof.' Then if firm1 decided to 'cheat' and lower its price by say £10, rather then attracting more customers, these customers would purchase the product from firm 2 and claim their £20. Eventually firm 2 will also have to lower its price by £10 and so firm 1 would in fact be worse off than before. Thus this advertising campaign is effectively a punishment strategy which guarantees that firm 1 will not lower its prices, and so both firms can enjoy the profits of an implicit cartel. Perhaps then such a strategy might be better analysed in the framework of co-operative game theory? This question has been addressed by game theorists and has been used by governments to break up such (illegal) cartels.(Dixit, p.103)

Evaluations and Conclusions

Game theory is criticised for its lack of use in predicting the outcomes of many real world oligopolies. Rather, its use lies simply in explaining the past. This feature is due to the sensibility of the outcomes to the details of the situation and to the rules of the game. However, when we compare this to the traditional approach of Sweezy's Kinked Demand Curve we find that game theory has much in its favour in the analysis of oligopolistic industries. Sweezy proposed that the oligopoly market demand curve is 'kinked' at a particular price and so the quantity produced by each firm will depend upon its marginal cost curve. The analysis does not explain or predict where the demand curve will be kinked or what quantities will be produced. Thus game theory's apparent weakness, in its attention to detail, is really a strength. Using game theory enables us to pinpoint a starting position and tells us exactly how, in microeconomic terms, a certain end position will be reached. One must therefore echo Kreps' words that we ought to be "*happily dissatisfied*".(Kreps 1990) Happy that progress has been made, and dissatisfied that we cannot yet predict real world economic behaviour. That said, game theory is being increasingly used by governments, management lecturers and firms world-wide. It is a tool that no economics student can afford to be without.

Bibiliography

Dixit, A.K. and Nalebuff, B.J. (1991) Thinking Strategically, W.W. Norton and Co., New York

Kreps, D. (1990) A Course in Microeconomic Theory, Harvester Wheatsheaf, London

Rasmusen, E. (1994) Games and Information, 2nd ed, Blackwell, London

Varian, H. (1993) Intermediate Microeconomics, 3rd ed. Norton, London

The Economist, The Games Economists Play, October 5th, 1994