INTRODUCTION TO GAME THEORY

Jonathan Cave

April 1987

THE RAND GRADUATE SCHOOL
FOUNDED 1970
CALIFORNIA

P-7336

INTRODUCTION TO GAME THEORY

Jonathan Cave

This Paper was prepared as a classroom note for the Fall 1986 Microeconomics I course at the RAND Graduate School which was co-taught by the author.

The RAND Corporation, 1700 Main Street, P.O. Box 2138, Santa Monica, CA 90406-2138

INTRODUCTION TO GAME THEORY

Jonathan Cave
April 1987

This is a glimpse at the concepts and tools of game theory in the context of microeconomics. It is neither exhaustive nor overly descriptive. At the end of the paper are references to longer works.

GENERAL REMARKS

A *game* is a situation in which several individuals have choices to make. In an interesting game, what one person does affects everyone's welfare. Going one step further, what one person does may affect what others choose to do. To see the difference, consider three situations.

1. **Game 1**: You and I have a case each of Beer and Perrier. Each of us prefers Beer to Perrier, and neither of us cares what the other drinks.
2. **Game 2**: Suppose there is only one bottle of each drink. Taste is more important than quantity, so half a bottle of Beer is better than a whole bottle of Perrier. However, we each prefer more to less of the same drink.
3. **Game 3**: We are both in a bar, and the Beer is really strong. In this case at least one of us must drink Perrier to be able to drive home.

In each of these games we interact in different ways. In game 1, we are independent. In game 2, my happiness depends on what you do, but my preferred action (drink Beer) does not. In game 3, both welfare and choices are interdependent. These interactions come from constraints; if there were two bottles of each drink in game 2, each of us could do as well as in game 1, regardless of the other's choice. In game 3, the drunk-driving constraint makes our choices (as well as our utility) interdependent. Economics deals with scarcity and other constraints, so this model is appropriate to economic situations.

Before going on, we need some apparatus. A **game** consists of three things:[1]

1. A set, N, of **players**; with the typical player called "i";
2. For each i, a set S_i of **strategies** or choices i can make [$S = S_1 x S_2 x \ldots x S_n$ is the set of outcomes -- one choice for each player]; and
3. For each i, a **payoff** (utility) **function**, $U_i(s)$ describing i's preferences over members of S.

Maximization problems, such as game 1, are games in which each player's payoff U_i depends only on s_i and not on the other s_j's.

There are different ways of describing a game. The **extensive** (tree) **form**, shows the explicit sequence of **moves** as a branched structure. A **strategy** is a collection of moves, one for each situation in which the player has to choose. One strategy therefore stands for many moves. In the **normal** (matrix or strategic) **form**, the players' strategies and payoffs are shown explicitly. Consider a game with two players, each of whom has two strategies. We draw a 2x2 matrix, and put a pair of numbers in each cell, corresponding to the payoffs of the two players. By convention, player 1 gets to choose the row, and player 2 the column. Figure A below shows extensive and normal forms for two versions of the game of "matching pennies." In the "perfect information" version, player 1 announces Heads or Tails, followed by a similar announcement on the part of player 2. If their announcements are the same, player 1 wins [payoff pair (+1, -1)]; if they announce different things, player 2 wins [payoff pair (-1, +1)]. Player 2 thus has two components to her strategy: a reply to player 1 if he announces H and a reply if he announces T. Since she can choose h or t in either case, she has 4 = 2x2 strategies. In the "simultaneous" version, both players announce at the same time. The extensive form represents this by an "information set," which is a box around the two points at which

[1] Some games specify the information available to different players. However, this data can be represented by using modified payoffs and strategies.

player 2 has to choose. It means that her choice must be the same no matter where she is in the box -- her choice cannot depend on what she does not know (namely, what player 1 will announce). In other words, the entire information set is a single "choice situation." For the normal form, this reduces the number of her pure strategies to 2 (h or t). The analysis of games rests on the **postulate of rationality**:

Each player i tries to get the highest possible value of U_i.

The meaning of this rule depends on players' views of the game, which determine what each player regards as possible. In particular, a player's best choice depends on what he or she expects other players to do. This general model can be given a variety of interpretations.

The noncooperative interpretation: players select strategies independently, guided by expectations or conjectures about the actions and reactions of other players. Strategies rationally chosen by independent players are in *equilibrium*.

Extensive

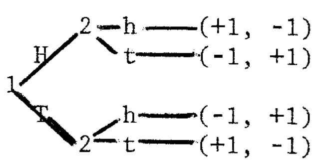

Normal

	(h, h)	(h, t)	(t, h)	(t, t)
H	(+1, -1)	(+1, -1)	(-1, +1)	(-1, +1)
T	(-1, +1)	(+1, -1)	(-1, +1)	(+1, -1)

a) Perfect Information

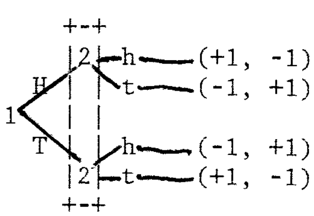

	h	t
H	(+1, -1)	(-1, +1)
T	(-1, +1)	(+1, -1)

b) Simultaneous Moves

Fig. A -- Normal and extensive forms for Matching Pennies

The institutional interpretation: the "rules of the game" themselves reflect individual choices, and players' attitudes and expectations are conditioned by history. This interpretation seeks to explain social institutions. For example, the struggle to control property is a game that can work to the disadvantage of all players, since property is often consumed in the process of establishing ownership. To control this, we have developed legal systems that allow us to make *agreements*. Without the penalties embodied in the law, the agreements would not be carried out. To understand the legal structure and identify improvements we must look at the incentives that guide behavior that laws are intended to promote or control.

The social choice interpretation: at the deepest level, games provide a theory of cooperative, collective, or social choice. The strategy space S contains all things that might be done (collective action is composed of individual actions), and the U_i are individual preferences for collective activity. Following this interpretation, s in S is a *social choice*, and we can ask what sort of choices will be made, whether they will be implemented, and how institutions within the game facilitate or impede the realization of collective preferences.

This interpretation also provides a theory of psychological choice. Economics customarily treats individuals as having well-defined preferences over comprehensive outcomes, but experience and psychological experiment suggest that this is not the case. The choices society offers us are highly complex, and difficult to evaluate. They bear only the most tenuous relation to the primitive "goods" of shelter, nourishment and reproduction. In addition, our preferences reflect a host of different outside influences: parents, society, experience,.... The human mind may act like a single entity, but is best regarded as a sort of society in which many different interests and capabilities struggle for influence. The members of this society play a game, and the outcome is "individual preference."

Let us use the three examples given above to illustrate the formal apparatus. Figure B shows their normal forms, with payoff numbers inserted for the preferences of the players. The strategy B_i (P_i) represents the choice of Beer (Perrier) by player i.

	B_2	P_2
B_1	(1,1)	(1,0)
P_1	(0,1)	(0,0)

1) Independent Payoffs

	B_2	P_2
B_1	(4,4)	(6,2)
P_1	(2,6)	(0,0)

2) Independent Choice

	B_2	P_2
B_1	(0,0)	(8,1)
P_1	(1,8)	(2,2)

3) Dependance

Fig. B -- Game matrices

The first (second) number of each pair represents player 1's (2's) payoff. In game 1, each person's payoff is unaffected by the other's choice, so the postulate of rationality tells us that each will choose Beer. In game 2, each person's payoff is affected by the other's choice, but the best choice for each person is not. Player i prefers Beer to Perrier regardless of which player j chooses, so rationality again tells us that each will choose Beer. In game 3, the choices are interdependent. Rationality predicts that only one will drink Beer, but not which one. If both players choose Beer (B_1,B_2), either can improve his payoff from 0 to 1 by unilaterally switching to P_i. If both players choose Perrier (P_1,P_2), either can improve his payoff from 2 to 8 by switching to B_i. If the players choose either (B_1,P_2) or (P_1,B_2), then neither can improve his payoff by a *unilateral* change of strategy, so we say they are in equilibrium.

This model requires us to pay attention to the degree to which the interests of the players coincide. Two polar cases are shown in Figure C. In game 4, the interests of the players coincide, and the theory predicts choices that further this mutual interest. In game 5, their interests are opposed; whatever one prefers, the other dislikes. It is to this situation that we turn next.

```
   L     R
 +-----+-----+
T|(0,0)|(1,1)|
 +-----+-----+
B|(2,2)|(3,3)|
 +-----+-----+
4) Coincident
```

```
   L     R
 +-----+-----+
T|(3,0)|(1,2)|
 +-----+-----+
B|(0,3)|(2,1)|
 +-----+-----+
5)   Opposed
```

Fig. C -- Confluence of interests

GAMES OF OPPOSED INTERESTS

These are called "zero-sum," "constant-sum," or "strictly competitive" games. **They always involve two players**, since it is impossible for more than two players all to have opposed interests. If player 1 prefers s to t, then player 2 must prefer t to s, and player 3 cannot disagree with both of them. Two-player games for money are games of opposed interests: what one player wins, the other loses. The strategic interaction does not change if they must pay for the right to play, as long as the payment does not depend on the outcome of the game.

The players choose strategies simultaneously or (what is the same thing) in ignorance of each other's choice.[2] To choose a "best" strategy, a player has to make a guess about the opponent's move. Figure D shows three games of opposed interests:

In Game 6, player 1 (2) always wishes to use the strategy T (L). Strategies which are always best are called **dominant**. In this game, neither player needs to consider what the other is doing. In game 7, player 1's best choice depends on what player 2 chooses. However, player 1 can figure that player 2 will always play L so player 1 should play T. In Game 8, neither player can figure out what is best without knowing what the other is going to do. We shall now describe two ways in which the players can pick moves.

[2] In games of timing, a player who moves later has a strategy that depends on the outcome of an earlier choice -- nothing changes if we suppose that the player writes down the contingent choices at the beginning of the game and gives them to a referee to carry out.

	L	R
T	(1,2)	(3,0)
B	(0,3)	(2,1)

6) Dominance

	L	R
T	(1,2)	(2,1)
B	(0,3)	(3,0)

7) Semi-Dom.

	L	R
T	(1,2)	(3,0)
B	(2,1)	(0,3)

8) No Dom.

Fig. D -- Games of opposed interests

Maxmin

The "safest" course is to assume the worst, and to take precautions against it. The rational conservative player i solves the following problem:

(1) $$\max_{s_i} \{\min_{s_j} U_i(s_1,s_2)\} = V_i$$

Each player is guided by the other's **capability** to harm him. No matter what s_i is, player j is expected to pick the s_j that makes player i worst off. In game 8, if player 1 picks T, player 2 will pick L, and if player 1 picks B player 2 will pick R. Accordingly, player 1 should play T which assures a payoff of at least 1. By the same token, player 2 should play L, which guarantees her at least 1.

This is a general principle: if player i plays according to (1), s/he chooses the so-called **optimal** or maxmin strategy s_i^m that guarantees a payoff of at least the **security level** V_i.[3]

[3] Since this is a game of opposed interests, we can say that player 2 maximizes U_2 by *minimizing* U_1. Therefore, we can describe the maxmin strategy for player 2 as:

(2) $$\min_{s_2} \{\max_{s_1} U_1(s_1,s_2)\} = -V_2$$

This suggests that the assumption that the other player is "out to get you" may be appropriate in the circumstances.

Equilibrium

It may be possible to do better by taking account of the opponent's **intentions**. U.S. defense policy vis-a-vis the Soviet Union is dominated by their capability to nuke us back (?) to the Stone Age. It ignores the question of whether it would be in the Soviet interest to use this capability. An alternative is to choose the best reply to the opponent's presumed **intentions**. Formally, player i's **best reply** to strategy s_j is $\beta_i(s_j)$; the s_i that solves the problem:

$$(3) \qquad \max_{s_i} U_i(s_i, s_j) = U_i(\beta_i(s_j), s_j)$$

Strategies that are best replies to each other are in **(Nash) equilibrium**. No player would want to change strategy after finding out what the other player did.

In game 8, the maxmin strategies (T,L) are not in equilibrium: player 1 would switch to B. Then player 2 would switch to R, 1 would switch to T, 2 would switch to L, and so on *ad infinitum*. It appears that this game has no equilibrium. Appearances can be deceiving, however, and we shall see that there is a sense in which we can talk about equilibrium for this game.

To summarize the discussion so far, we have defined two approaches. The maxmin approach is a conservative procedure according to which the player concedes the second-move advantage to his opponent. The equilibrium procedure lets each player react to his opponent's choice. Maxmin or optimal strategies depend on opponent capabilities, and guarantee each player at least his or her security level. They predict how people would act in ignorance of their opponent's motives (payoffs) and intentions. Optimal strategies are prescribed choices. Equilibrium strategies are a self-confirming set of expected choices. Optimal strategy play may not persist after players have a chance to learn about each other and the game. Equilibrium, which is based on opponent intentions (mutual best replies), has the property that no player wishes to change his or her mind after seeing the other player's choice. Equilibrium behavior can persist over time: long enough to allow us to observe it, and long enough to make it socially significant. However, we cannot yet ensure that equilibria exist.

Mixed Strategies

Players do not really react to other players' choices, but rather to what they expect other players to do. These expectations may involve uncertainty. For example, player 1 may believe that player 2 is three times as likely to play L as R. We represent this by assigning **probabilities**: player 2 is expected to play L with probability 3/4 and R with probability 1/4. These beliefs should be consistent with observation and rationality. Beliefs about player 2's strategy should assign non-zero probability to anything that she will actually do. At the same time, player 2's expected choices should be best replies to her beliefs about player 1. These beliefs are called **mixed strategies**, and they must be **common knowledge**.

Common knowledge is a technical term. It says that

> player 2 knows what player 1 expects her (player 2) to do;
> player 1 knows that player 2 knows player 1's beliefs about player 2;
> player 2 knows that player 1 knows that player 2 knows player 1's beliefs about player 2,...

and so on *ad infinitum*.

Mixed strategies can be modeled as random choices. Player 2 could flip a coin twice, and choose R if "Heads" came up both times (probability 1/4), and L otherwise. Of course, there is no reason for her to do this: it is sufficient that player 1 cannot predict exactly what she will do. If the other player can make such a precise prediction, the probability of using a given strategy is 1, and all other strategies get probability 0. This is called a **pure strategy**.

We limit attention to the 2x2 case. The probability attached to player i's first (T or L) strategy is a number p_i between 0 and 1. The other strategy (B or R) must receive probability $1-p_i$. The result of (p_1,p_2) will be:

(T,L) -- expected with probability p_1p_2;
(T,R) -- expected with probability $p_1(1-p_2)$;
(B,L) -- expected with probability $(1-p_1)p_2$; and
(B,R) -- expected with probability $(1-p_1)(1-p_2)$.

To find the **expected payoff**, multiply each payoff by the probability of its occurrence, and add them up:

$$(4)\qquad EU_i(p_1,p_2) = p_1p_2U_i(T,L) + p_1(1-p_2)U_i(T,R) + (1-p_1)p_2U_i(B,L) + (1-p_1)(1-p_2)U_i(B,R)$$

To analyze player i's best reply to p_j, rearrange equation (4) to isolate the effect of changes in p_i on U_i:

$$(5)\qquad EU_1 = p_1\{p_2[U_1(T,L)-U_1(B,L)]+(1-p_2)[U_1(T,R)-U_1(B,R)]\} + p_2U_1(B,L) + (1-p_2)U_1(B,R)$$

$$EU_2 = p_2\{p_1[U_2(T,L)-U_2(T,R)]+(1-p_1)[U_2(B,L)-U_2(B,R)]\} + p_1U_2(T,R) + (1-p_1)U_2(B,R)$$

The best reply β_i is 1 if the term in { } in the equation for EU_i is positive, 0 if { } < 0, and can be anything if { } = 0. For game 8:

$$EU_1 = 2p_2 + p_1[3 - 4p_2]$$
$$EU_2 = 3(1 - p_1) + p_2[4p_1 - 2]$$

Figure E shows the best replies of the two players in (p_1, p_2) space. Where they intersect [at (1/2, 3/4)] the players are in equilibrium. You should satisfy yourself that neither player can improve (or even affect) his own payoff by a unilateral change of strategy.

A simple diagram can be used to find the best replies. Graph the payoff to player i against the (mixed) strategy of the other player. Note that:

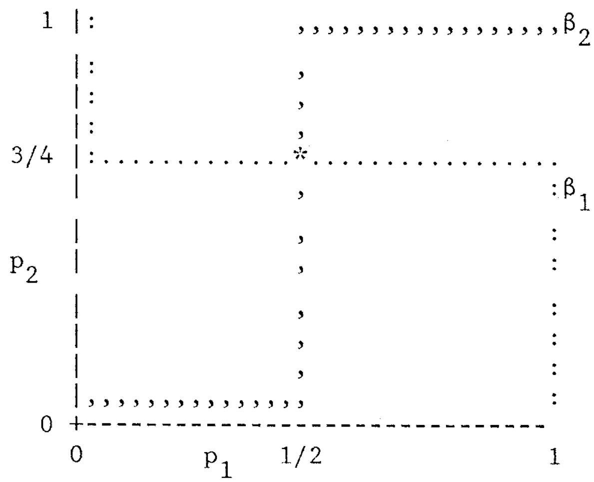

Fig. E -- Best replies for game 5

$p_1 = 0$ (1) is the same as B (T); and
$p_2 = 0$ (1) is the same as R (L).

Figure F shows what player i can get against any strategy of player j if he uses either of his pure strategies. The solid line in Figure F indicates player i's best reply. It is clear that:

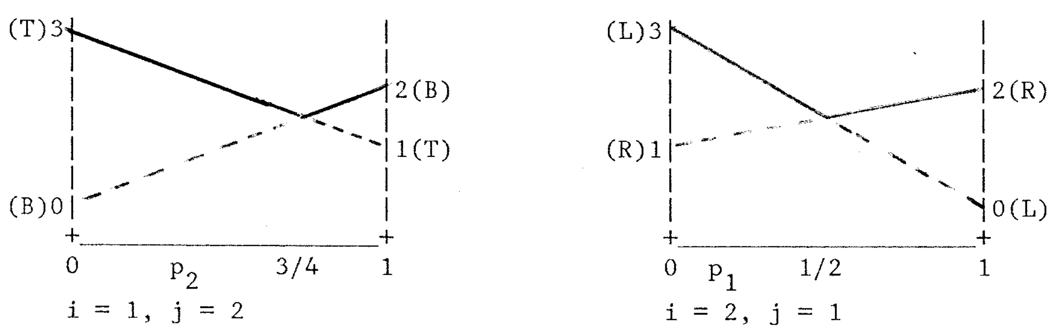

Fig. F -- Player i's payoff versus player j's strategy: best replies

$$\beta_1(p_2) = \begin{cases} 1 & \text{if } p_2 < 3/4 \\ \text{any} & \text{if } p_2 = 3/4 \\ 0 & \text{if } p_2 > 3/4 \end{cases} \quad \text{and} \quad \beta_2(p_1) = \begin{cases} 1 & \text{if } p_1 > 1/2 \\ \text{any} & \text{if } p_1 = 1/2 \\ 0 & \text{if } p_1 < 1/2 \end{cases}$$

Figure G uses a similar diagram to find the maxmin strategies. Player i's payoff is plotted against his own strategy. We again get two lines, one for each pure (not mixed) strategy of the other player, j. The solid line is the "min U_i" function that player i maximizes. Figure G shows that player 1's maxmin strategy is $p_1 = 1/2$, and player 2's maxmin strategy is $p_2 = 3/4$. These are the same as their equilibrium strategies: considering the opponent's capabilities (maxmin) leads to the same prediction as considering his intentions (equilibrium).

This is always true in games of opposed interests. Essentially, it turns on the observation that what player i expects player j to do (minimize U_i) is exactly what player j wants to do (by maximizing U_j). If player i can guarantee at least V_i, player j can prevent player i from getting any more than V_i.

As immediate corollaries to this theorem, the optimal strategies are in equilibrium. There may be other equilibria, but they all have the same payoff. If there are many equilibria, the players never need to coordinate their strategy choices: if player 1 uses any equilibrium strategy, and player 2 uses any equilibrium strategy, the result is

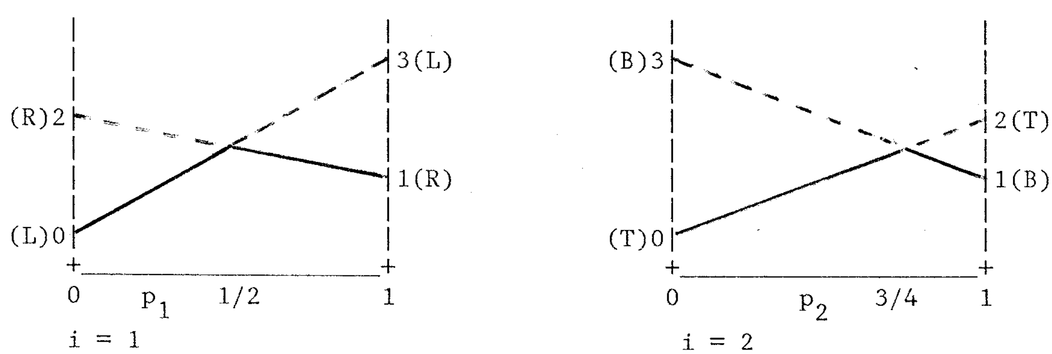

Fig. G -- Player i's payoff against player i's strategy: maxmin always an equilibrium. Finally, compared with an equilibrium, even a coordinated change of strategy will necessarily make one player worse off.

GENERAL NONCOOPERATIVE GAMES

If there are more than two players, or if players' interests are not strictly opposed, these comforting conclusions fall apart. Optimal strategies need not be in equilibrium. Indeed, the assumption that the other player(s) intend to minimize one's own payoff is incorrect.

All one can say is that "active" behavior (best reply) is never worse than the "passive" behavior of optimal strategies. The payoff to each player in equilibrium cannot be less than his security level, and is usually more. Therefore, both players can gain by ignoring each other's capabilities in favor of intentions.

Even in a game with many players, the idea of a best reply is valid: a player maximizes his payoff taking the strategies of other players as fixed. To focus the discussion, consider a **duopoly**: two firms producing a homogeneous good. The quantity produced by firm i is q_i, and the price is:

$$(6) \qquad P = 30 - 2(q_1 + q_2)$$

The two firms have the same cost function: $C(q_i) = 6q_i$ -- there are no fixed costs, and marginal cost is constant (= 6). The payoff (profit) function for each firm is:

$$(7) \qquad \pi_i = Pq_i - C(q_i) = [30 - 2(q_1 + q_2)]q_i - 6q_i = 24q_i - 2q_iq_j - 2q_i^2$$

The best reply of firm i is found by setting the derivative of π_i with respect to q_i equal to 0. This gives:

$$\beta_1(q_2) = 6 - [q_2/2]$$

(8)

$$\beta_2(q_1) = 6 - [q_1/2]$$

These best reply functions are shown in Figure H. Solving them simultaneously gives the (Nash or Cournot) equilibrium: each firm produces 4 units of the good.

To put this result in perspective, let us consider what price-taking ("competitive") firms or a monopolist would do.

In a competitive market, P = MC = $6.00, so the total quantity produced would be 12 units.

The marginal revenue curve corresponding to (6) is MR = 30 - 4Q, where Q is the total supplied. Therefore a monopolist would produce 6 units.

What happens if one firm (say firm 1) realizes that the other firm's output is not really given, but is instead a best reply to firm 1's output? This realization changes firm 1's perception of its profit function to:

(9) $$\pi_1 = [30 - 2(q_1 + \beta_2(q_1))]q_1 - 6q_1$$
$$= [12q_1 - q_1^2]$$

which depends only on q_1. If firm 1 thinks in this way, it is called a **(Stackelberg) leader**, and firm 2 is called a **follower**. Setting the derivative of the expression in equation (9) equal to 0, we get the **Stackelberg equilibrium** quantities: the leader's profit-maximizing output is 6, which induces the follower (firm 2) to produce 3 units.

The selection of firm 1 as the leader was arbitrary; either firm could "figure out" what the other was doing -- once this happens, the "smart" firm need not consider the actual output of the other firm in choosing its own output. What happens if both firms work out that their original (no-reaction) conjectures were wrong?

Well, they are both wrong again, but on a "higher plane" -- each will think of itself as the leader, and will therefore produce 6 units. But this is the competitive level of output, so neither firm will earn any profits. This is called the **Stackelberg disequilibrium**.

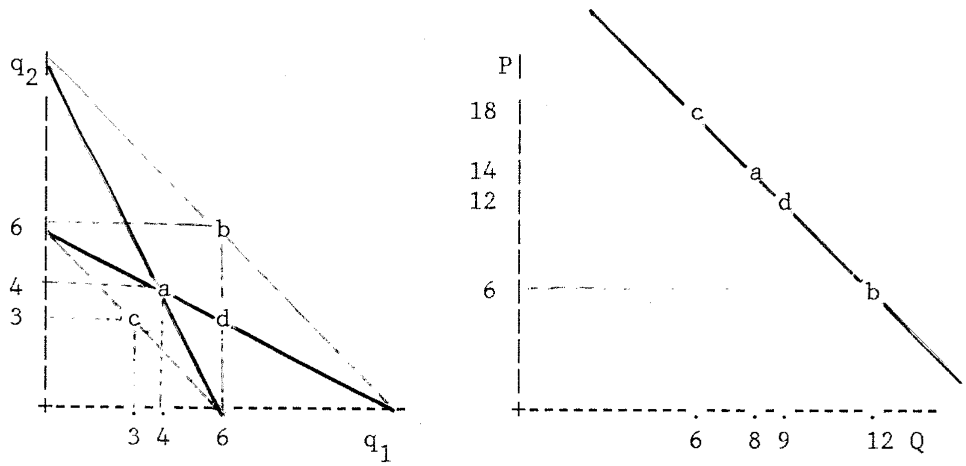

a = Nash equilibrium
b = Pure competition = Stackelberg disequilibrium
c = Monopoly = collusive cartel
d = Stackelberg equilibrium with firm 1 as leader.

Fig. H -- Reaction functions and duopoly outcomes

There is one other possibility we will investigate with this simple model. Oligopolies require entry barriers just as monopolies do, since they earn positive profits (unless they outsmart themselves). Suppose there are n identical firms (instead of just 2). How do the output levels depend on n?

If the total produced by the other (n-1) firms is Q_{n-1}, firm i's profit is:

$$(10) \qquad \pi_i = [30 - 2Q_{n-1} - 2q_i]q_i - 6q_i$$

and i's best reply is $q_i = 6 - [Q_{n-1}/2]$. This equation is the same for all n firms, so we can assume that each produces the same amount, q_n. Then $Q_{n-1} = (n-1)q_n$, and the equilibrium condition is:

$$(11) \qquad q_n = 6 - (n-1)q_n/2 = 12/(n+1)$$

Total output is $12n/(n+1)$, which approaches the competitive output of 12 as $n \to \infty$.

The results are summarized in Table 1, which shows quantities, price, profits, consumer surplus and total welfare for the various possibilities.

There are a number of principles that emerge from these general results. In the first place, Cournot (Nash) equilibrium is better for all firms than pure competition, but worse for all firms than monopoly. If the firms could agree to restrain themselves to 1/n times the monopoly output they would all earn higher profits. This is called collusive, cartel, or price-fixing behavior. It would not be an equilibrium, which means that each firm could "free-ride" on the restraint of the others, and the "agreement" would collapse.

The choice of a solution may be a higher level strategic choice on the part of the players. Figure I shows two possibilities. In the first game, the players choose between collusive play (produce 3 units) and playing the best reply to the other firm. In the second game,

Table 1

DIFFERENT RESULTS FOR THE OLIGOPOLY MODEL

Solution	q_i	Q	P	π_i	π	CS	W
2 firms:							
Nash	4	8	14	32	64	64	128
Competition	6	12	6	0	0	144	144
Cartel	3	6	18	36	72	36	108
Stack. Eq:		9	18		54	81	135
leader	6			36			
follower	3			18			
Stack. Dis.	same as "Competition"						
Nash (n firm)	$\frac{12}{n+1}$	$\frac{12n}{n+1}$	$\frac{30+6n}{n+1}$	$\frac{288}{(n+1)^2}$	$\frac{288n}{(n+1)^2}$	$\frac{144n^2}{(n+1)^2}$	$\frac{144n(n+2)}{(n+1)^2}$

players choose between playing best reply and playing as a Stackelberg leader (produce 6 units).

Strategically, the first game is Prisoners' Dilemma and the second game is Chicken.

In the general {N,S,U} model of a game, a strategy vector is in **equilibrium** if no player can improve his payoff by changing only his strategy. A strategy vector is **Pareto optimal** if there is no other strategy that makes **all** players better off. The "strategic problem" is that few equilibria are Pareto optimal. When people follow their selfish interests, they are likely to miss alternatives that all of them would prefer. Conversely, most Pareto optimal agreements need some kind of enforcement mechanism before they can be carried out. This is very different from the economics of the "invisible hand," which ensures that equilibria are Pareto optimal and *vice versa*.

There are also logical problems with equilibrium. Each player takes the others' strategies as fixed. Equilibrium strategies may involve contingent behavior, or **threats**. Threats only work if players believe that they will be carried out if push comes to shove. Alternatively, a player making a threat that he would not wish to carry out is relying on the rationality of the other players. It is as if players could **precommit** themselves to carrying out their strategies, and do not want to place undue reliance on the rationality of other players. Credibility and reliance are the same thing (for our purposes) and lead to the idea of **perfect equilibrium**: one in which all threats

	cartel	best reply
C	36, 36	27, 40.5
BR	40.5, 27	32, 32

a) Collude or Fight

	best r.	St. leader
BR	32, 32	18, 36
SL	36, 18	0, 0

b) Follow or Lead

Fig. I -- Higher level strategic options for duopolists would be carried out. In an extensive-form game, perfect equilibria can be found through "backwards induction"; starting at the end of the game, find the best (equilibrium) behavior for the last player(s). Then look at the next-to-last move and decide what the equilibrium is in the same way.

How can we find perfect equilibria in the normal form, where threats are implicit in strategies? If a strategy s_i involves a threat, there is another strategy t_i which differs from it only in that the threat is not made. If s is part of an equilibrium in which the threat is not carried out, then both s_i and t_i must be best replies to the strategies of the other players. In fact, s_i and t_i are equivalent except for those strategies of the other players that trigger the threat, in which case s_i is inferior to t_i; this means that s_i is (weakly) **dominated** by t_i. In general, if s_i and t_i are two strategies of player i, and if

(12) $$U_i(s_i, r_{-i}) \le U_i(t_i, r_{-i})$$

for all strategies r_{-i} of the other players, with strict inequality for at least one r_{-i}, we say that s_i is dominated by t_i. s_i can be used in equilibrium, if (12) holds with equality, but i would not use s_i if he was the least bit unsure about what the other players were going to do, since t_i costs him nothing to use, and might help him.

Therefore, in perfect equilibrium all players avoid dominated strategies.[4] Since perfect equilibria are equilibria, they share some of their drawbacks; in particular they might not be Pareto optimal, as the Prisoners' Dilemma in Figure J shows: each player has either a greedy (G) or a helpful (H) strategy. The only equilibrium is (G,G) -- the greedy strategy dominates the helpful strategy. This is a stylized version of our "collude or fight" duopoly game (Ia above). The Prisoners' Dilemma is in many ways the starkest example of the strategic or "free rider" problem alluded to above: individual collective

[4] Strictly, this is true only for games with a finite number of pure strategies.

interests may diverge. We shall use it to illustrate institutions that can reconcile these interests.

Agreements

The {N,S,U} model supports a theory of agreements that goes beyond simple equilibrium notions. Consider the games in Figure J. Agreements should be Pareto optimal: if everyone prefers not to do something, it should not happen (do not look for realism here!). Let us consider two candidate agreements: 1) use a coin flip to pick (B,P) or (P,B) in the Battle of the Sexes; and 2) play (H,H) in the Prisoners' Dilemma. The first agreement is **self-enforcing**: no matter whether the coin comes up Heads or Tails, both parties prefer to go along with the agreement. The second agreement is not self-enforcing: both parties would happily sign the agreement, but each would have his fingers crossed, and neither would carry it out. In fact, no agreement that does not involve choosing equilibria all the time can be self-enforcing.

How can we get around this problem? We have to change the payoffs. One way to do this is to create an exogenous penalty or authority. For example, the players may give hostages to each other to be killed in the event the agreement is not carried out. Suppose that the value of the hostage (to the giver) is 2. The situation is shown in Figure Kb.

With hostages, the only equilibrium is at (H,H) and cooperation is assured. On the other hand, if one party manages to give a hostage they do not like (so the value of the hostage is 0); we get the "Fake hostage" situation of Figure Kc, in which the other player is exploited

	B	P
B	(0,0)	(8,1)
P	(1,8)	(2,2)

9) Battle of the Sexes

	H	G
H	(3,3)	(0,4)
G	(4,0)	(1,1)

10) Prisoners' Dilemma

Fig. J -- Prisoners' Dilemma and Battle of the Sexes

	H	G
H	(3,3)	(0,4)
G	(4,0)	(1,1)

a) No Hostages

	H	G
H	(3,3)	(0 , 2)
G	(2,0)	(-1,-1)

b) Real Hostages

	H	G
H	(3,3)	(0, 2)
G	(4,0)	(1,-1)

c) Fake Hostage

Fig. K -- Hostages in Prisoners' Dilemma

to the maximum extent. This means that we may only have raised the problem to a higher level -- both players benefit if real hostages are used, but each does even better by using a fake hostage, and would use a fake hostage in self-defense if the other player did. This is why **verification** and **signaling** are such an important part of these agreements: people have to be sure that the other side is using "real" hostages.

Examples of "hostage" institutions include: deposits on rentals; earnest money in real estate transactions; engagement rings; prenuptial agreements; and Mutual Assured Destruction. At a more abstract level, penalties can be interpreted as hostages under the rule of law. Property pledged as collateral or other contractual penalties serve to coerce people into good behavior. Indeed, the entire body of criminal law can be seen in the same light: the hostage we give is our absolute right of independent action (life, liberty, and the pursuit of happiness) and we can lose the hostage if due process of law determines that we have violated the agreement implicit in the social compact.

The structure of the game can provide the enforcement mechanism. All social institutions persist and require time for their proper functioning: time to verify that an agreement has been violated; time to identify the culprit; and time to carry out the appropriate punishment. A simple model of such an ongoing interaction is the **repeated game**. To get ourselves to cooperate today, we threaten each other with punishment in the future. One institution that reflects this idea is **reputation**. It is a summary of past behavior used to determine whether the blessings of cooperation will be available to a person in the future.

Consider the Prisoners' Dilemma again. By cheating today, I increase my payoff by \$1.00 (from \$3.00 to \$4.00). If I am punished for this tomorrow, I lose \$2.00 (from the \$3.00 I would have earned to the \$1.00 I get when the agreement breaks down). Therefore, it does not pay me to cheat.

Or does it? If we examine this situation carefully, we see there is something wrong. If the game is to be repeated 10 times, it certainly looks like I should play H. But what about the last period? On the 10^{th} play there is no future, so neither of us can be deterred from cheating. But if this is so, on the 9^{th} play the same is true; I know we will play (G,G) tomorrow regardless of whether I play G or H today. This means that the entire agreement collapses! This accounts for the fact that the prospect of the end of the world (as when people misunderstood Orson Wells' broadcast of "War of the Worlds") is attended with all sorts of bad behavior; nobody cares about tomorrow or its punishments. Eat, Drink, and be Merry....

This cannot be the end of the story. After all, reputations do play an important role in the world. One possibility is that, while no one expects to live forever, no one expects to die tomorrow. An appropriate mathematical model is that I have a probability p of living one more day. This means that I multiply tomorrow's payoff by p, the day after tomorrow's by p^2 and so on. The probability of living forever is therefore $p^{\infty} = 0$. To find out what this does to our agreement, let us do a little computation. First, the "present value" of a stream which pays \$x on each day is:

(13) $$PV(x) = x + px + p^2x + \ldots = x/(1-p)$$

If we have an agreement that pays, say, \$3.00 each day and I plan to cheat, the worst punishment that can be inflicted gives me \$1.00 on each day, starting tomorrow. I will not cheat if:

(14) $$3/(1-p) \geq 4 + p(1) + p^2(1) + \ldots = 4 + p/(1-p), \text{ or}$$
$$p \geq 1/3$$

In general, the agreement must pay me each day at least (1-p) times what I get if I cheat + p times my security level (or the smallest amount I can be forced to accept if I am punished).[5] For people whose p is close to 0 (short-sighted or nearly dead), this is almost the condition for equilibrium; for people with p close to 1 (far-sighted or healthy), almost any individually rational behavior can be supported by an equilibrium of the repeated game. The formula also predicts that circumstances that lower p (high interest rates, greater uncertainty of survival, etc.) will decrease cooperation, since agreements will no longer be viable.

---

[5] This can be written (1-p)(best reply) + p(minmax). Note that the details of the punishment must be known beforehand, so the appropriate security level is the minmax, rather than the more severe maxmin security level appropriate to games of opposed interests.

REFERENCES

Games and Decisions by W. Luce and H. Raiffa, Wiley, 1957.

The Calculus of Consent by J. Buchanan and G. Tullock, University of Michigan Press, 1965.

The Strategy of Conflict by T. Schelling, Harvard University Press, 1980.

Game Theory and Politics by S. Brams, Free Press, 1975.

A Game-Theoretic Approach to Political Economy by M. Shubik, Massachusets Institute of Technology Press, 1984.

RAND/P-7336 INTRODUCTION TO GAME THEORY Jonathan Cave