

MA40128
Final Year Project
A survey of Game Theory

Graeme Taylor

10th May, 2006

Contents

1	Introduction	3
2	Two player zero-sum games	3
2.1	Zero-sum, utility and the assignment of worth	3
2.2	Strategic form for two player zero-sum games	4
2.3	2×2 Payoff matrices	5
2.3.1	A strategy for Player 1	6
2.3.2	A strategy for Player 2	6
2.4	The Minimax Theorem for a two player zero-sum game	7
2.4.1	Games with Saddle points	7
2.4.2	2×2 Games	7
2.4.3	Dominated strategies	9
2.5	Finite strategic form games as a linear programming problem	10
3	Two player general sum games and co-operation	11
3.1	Strategic form for two player general sum games	12
3.2	The prisoner's dilemma	12
3.2.1	Dominated strategies	13
3.2.2	Nash Equilibria	14
3.3	Co-operation in the prisoner's dilemma	14
3.4	The iterated prisoner's dilemma	16
3.4.1	Group interest	17
4	n-player games	18
4.1	Strategic form for n -player general sum games	18
4.2	Nash's Theorem	19
4.2.1	Proof of Nash's Theorem	20
4.3	The 3-player simple majority game	22
4.3.1	Simple Majority	22
5	Co-operative games and coalitions	23
5.1	Coalitional form for n -player games	23
5.1.1	Coalitional form of a strategic form game	26
5.1.2	S -veto games	26
5.2	Shapley Value	27
5.3	Shapley Axioms	27
5.4	Shapley's Theorem	30
6	Concluding remarks	31
7	Figures and sources	32

1 Introduction

This report seeks to give an account of the basic principles of Game Theory, by tracing the progression from the subject's roots in the work of John Von Neumann in the 1920s, to that of John Nash and Lloyd Shapley in the 1950s. Section 2 reconstructs Von Neumann's Minimax theorem after direct calculations with simpler examples of zero-sum games, whilst the Nash equilibrium is introduced in-depth in the context of general sum games during section 4. After a consideration of the limitations of the non-cooperative model, co-operative game theory is examined in section 5, via coalitional games and Shapley's theorem.

Two properties, which it shares with Operational Research, characterise Game Theory- its relative youth, and its connections to other disciplines both within mathematics and beyond. A formal treatment of strategy, rationality and the activity of individuals in groups, whilst initially applied to economic problems, has found applications in topics as diverse as evolutionary biology and computer science. Some examples of such connections have been given where appropriate. One game in particular, the prisoner's dilemma, has proved particularly far-reaching, having been described as "the E.Coli of social psychology". It continues to surprise, and is considered in detail in section 3, with section 3.4.1 describing particularly recent results from this famous problem.

2 Two player zero-sum games

The two player, zero-sum game is probably the barest essence of a game. Two participants are presented with a single choice that determines both victor and spoils, with the winner collecting his payoff directly from the loser. However, some important concerns come to light even within this minimalist framework, indicating their relevance to any work that builds upon such simple examples.

2.1 Zero-sum, utility and the assignment of worth

To make much progress with a mathematical theory of games (as opposed to simply a philosophy), it will be necessary to capture the game in question numerically, in an attempt to quantify the impact of a set of decisions. To do so, however, requires some potentially unrealistic assumptions. Clearly, many games will easily lend themselves to quantification. Any financial transaction will carry with it cash values - profits and losses - for each participant. Other situations may have an ultimately numerical outcome, but the chain of cause and effect is less clear, such as trying to determine the number of votes that will be gained or lost by a political party reforming its platform. Alternatively, the outcome may be difficult to capture numerically at all- just what effect does changing a given note have on the enjoyment offered by a piece of music?

Assuming that a numerical value can be readily assigned to outcomes, there is the added complication of reconciling those values between players. It may be possible for an objective worth to be assigned to an object (or transaction or outcome...) yet for the participants to value that

worth differently. There are many ways in which this may occur. Consider for instance the sale of a business- if a bidder intends to purchase it as a going concern, instead of to simply liquidate the assets, then they will attach additional value to 'goodwill' factors such as the size of the existing customer base, beyond the net worth of existing assets. The average pay-out on a lottery ticket is less than the asking price of the ticket (assuming the organisers have any sense!), yet hundreds of thousands of people purchase a ticket every week; the slim chance at a disproportionately large payout is valued higher than the entry price. Conversely, a relatively small sum of money to a company could be devastating to an individual. If this is the case with legal costs, for instance, then it may be preferable to settle out of court for an affordable amount, even if innocent, rather than run the risk of losing your home.

Such subjective differences are examined in the context of utility values. From here on, it will be generally assumed that the pay-offs discussed are accurate reflections of the worth to the participant, so they may measure expected utility to the player instead of a more objective value. The point is that the pay-off should genuinely capture all of the worth of the outcome; maximising the pay-off may then be assumed to be the only motivation of the player (other motivations having been subsumed into the calculation of that value in the first place). However, this issue is particularly relevant to games which offer a potential for co-operation, and will be revisited then.

The most contentious stipulation is of a zero-sum scenario, where players are diametrically opposed in interest such that the gain of one participant necessitates an equal loss for the other. This assumption is not always tenable - scenarios may offer the potential for mutual gain or loss - and thus will be phased out in due course. For now, it provides a suitable simplifying assumption to make progress in the development of theory, and is not completely irrelevant to 'real world' situations (from gambling to finance).

2.2 Strategic form for two player zero-sum games

Definition 2.2.1. A two player zero-sum game in strategic (or normal) form X, Y, A consists of two strategy sets X and Y , corresponding to the players, and a function $A : X \times Y \rightarrow \mathbb{R}$ representing the 'pay-off'. The game is *finite* if both X and Y are finite sets.

Definition 2.2.2. A play of the game consists of Player 1 choosing a strategy $x \in X$ and Player 2 simultaneously choosing a strategy $y \in Y$. Player 1 is then awarded $A(x, y)$ in winnings, and Player 2 awarded $-A(x, y)$ (i.e, Player 2 loses whatever Player 1 wins; this is the *zero-sum* condition).

Definition 2.2.3. The *payoff matrix* or *game matrix* for the game with $X = \{x_1, \dots, x_m\}$, $Y = \{y_1, \dots, y_n\}$ and payoff function A is given by

$$A = \begin{pmatrix} a_{11} & \cdots & a_{1n} \\ \vdots & & \vdots \\ a_{m1} & \cdots & a_{mn} \end{pmatrix} \text{ where } a_{ij} = A(x_i, y_j)$$

That is, the $(i, j)^{th}$ entry of A determines the winnings for Player 1 and losses for Player 2 when Player 1 chooses strategy $x_i \in X$ and Player 2 chooses strategy $y_j \in Y$. As a shorthand, we may describe Player 1 as choosing the row and Player 2 as choosing the column.

2.3 2×2 Payoff matrices

We consider initially the simplest case - a 2×2 payoff matrix - where each player is presented with a single choice. As a motivational example, consider the game of *two finger Morra*, described by

$$\begin{pmatrix} -2 & 3 \\ 3 & -4 \end{pmatrix}$$

This game is related to an ancient Roman guessing/gambling game, *micare digitis* (to flash with fingers); however, in this simple formulation, it is the parity of the total, rather than a successful guess of its value, which determines the victor. Each player reveals either 1 or 2 fingers, with the winnings being the total number of fingers shown. If the total is odd, Player 1 wins; otherwise Player 2 wins (giving rise to the negative entries in the above payoff matrix).

The purpose of studying the strategic form is to attempt to determine a strategy for each player that is in some sense optimal. The ideal scenario for Player 1 is that there is a strategy that always enables him to win any given play of the game. However, such games are likely to be few and far between, and willing participants for the role of Player 2 even rarer.

However, we can more usefully tackle a refinement of this question, namely determining whether there is a strategy for Player 1 which, *in the long run*, they might expect to profit from. For any play of the game, Player 1 can pick either option 1, or option 2. Their return will then depend on the strategy employed by Player 2 for that particular play. We may suppose that Player 1 chooses option 1 with probability p_1 , either in accordance with some plan or simply by chance; choosing option 2 the rest of the time, i.e., with probability $p_2 = 1 - p_1$. We will refer to this as a *mixed strategy*, although it is worth noting the special case of a *pure strategy*, where $\{p, q\} = \{0, 1\}$ and thus only one option is ever used.

The objective for Player 1, therefore, is to devise a mixed strategy that maximises their payoff. At the same time, Player 2 is trying to minimise the payoff of Player 1, since this maximises their own payoff (by the zero-sum condition). Whilst neither player is aware of the particular option the other intends to take in a given play of the game, their calculations can take into account this motivation on the part of their opponent. Perhaps surprisingly, this does not descend into endless second-guessing, and should Player 1 find an optimal mixed strategy, they can even safely pre-declare the mix (although not a given move) without giving an advantage to Player 2. We shall illustrate how this arises for 2 finger Morra, then consider generalisations to any 2×2 game. Note, however, that we have once again sidestepped some considerations of utility theory in our acceptance of expected payout as a good measure of the worth of a game, especially if the number of plays is to be small.

2.3.1 A strategy for Player 1

Can Player 1 guarantee a certain minimum (and preferably positive) payoff? Note that if she employs the mixed strategy (p_1, p_2) , then her return depends on the strategy of Player 2:

- If Player 2 opts for '1', then the return for Player 1 is -2 (if she played 1) or 3 (if she played 2). Thus on average, she may expect a payoff of $-2p_1 + 3p_2$
- If Player 2 opts for '2', then Player 1's expectation is $3p_1 - 4p_2$

If, therefore, we seek an expected payoff of at least V *regardless of Player 2's strategy*, then we require

- $-2p_1 + 3p_2 \geq V$
- $3p_1 - 4p_2 \geq V$

As a first attempt, consider the case of equality:

$$\begin{aligned} V &= -2p_1 + 3p_2 &= 3p_1 - 4p_2 \\ &7p_2 &= 5p_1 \\ 7(1 - p_1) &= 5p_1 \\ 7 &= 12p_1 \\ \frac{7}{12} &= p_1 \end{aligned}$$

Thus we have a mixed strategy $(\frac{7}{12}, \frac{5}{12})$ where the expected payoff is $-2(\frac{7}{12}) + 3(\frac{5}{12}) = \frac{1}{12} = 3(\frac{7}{12}) - 4(\frac{5}{12})$. So Player 1 can guarantee an expected return of $\frac{1}{12}$ per play (over a large number of plays).

2.3.2 A strategy for Player 2

Analogously to the duality theorem in linear programming, we may determine whether it is possible for Player 1 to ensure a greater expectation by seeing whether Player 2 is able to cap their losses at the $\frac{1}{12}$ per play presented above.

In fact, Player 2 can minimise their losses in this way (and thus Player 1 must be content with the value of $\frac{1}{12}$) by the same strategy. For Player 2 the expected payoffs are $2p_1 - 3p_2$ when Player 1 opts for '1' and $-3p_1 + 4p_2$ when she opts for '2'; so with a mixed strategy of $(\frac{7}{12}, \frac{5}{12})$ Player 2 expects $\frac{-1}{12}$ in either case.

So, on average, Player 1 values the game as being good for at least $\frac{1}{12}$ per play, whilst Player 2 can ensure it is no worse than $\frac{-1}{12}$ per play for them, i.e., it is at best worth $\frac{1}{12}$ to Player 1. This is an example of general behaviour:

2.4 The Minimax Theorem for a two player zero-sum game

Theorem 2.4.1 (Minimax Theorem). *Any finite 2 player zero-sum game in strategic form, has a value V such that Player 1 has a mixed strategy ensuring an average payoff of V , regardless of Player 2's strategy; and Player 2 has a mixed strategy which ensures an average payoff of $-V$, regardless of Player 1's strategy.*

It therefore follows that the optimal strategy for Player 1 is one that ensures a payoff equal to the value. Were they to attempt to secure a larger return, they would fail if Player 2 followed the strategy prescribed by the theorem: a payoff greater than V for Player 1 would force a payoff of less than $-V$ to Player 2, which contradicts the Minimax result.

Von Neumann's proof of the Minimax theorem, advanced in his 1928 paper *Zur Theorie der Gesellschaftspiele*, depended upon topological notions such as Brouwer's fixed point theorem (which we shall make use of in section 4.2.1). With the post-war emergence of Operational Research, a simpler proof of the two player, zero-sum case became accessible. Thus the details of the original proof shall be omitted, and in section 2.5 it shall instead be demonstrated that the Minimax theorem is a linear programming problem and hence solvable through application of (for instance) the Simplex algorithm. First, however, it is worth examining some simple cases where routine algebra can give insight into optimal strategies.

2.4.1 Games with Saddle points

Definition 2.4.2. An element a_{ij} of a matrix A is described as a *Saddle point* if it equals both the minimum of row i and the maximum of column j .

Theorem 2.4.3 (Minimax Theorem for a game with a Saddle point). *Let a_{ij} be a saddle point for a game in strategic form given by X, Y and A . Then the game has value a_{ij} , achieved when Player 1 plays the pure strategy x_i and Player 2 the pure strategy y_j .*

Proof. Player 1 is guaranteed a payoff of at least a_{ij} by using strategy x_i since for any strategy choice $y_{j'}$ by Player 2, $A(x_i, y_{j'}) = a_{ij'} \geq a_{ij}$ since a_{ij} is the minimum of row i . Thus $V \geq a_{ij}$. Player 2 is guaranteed a payoff of at least $-a_{ij}$ by using strategy y_j since for any strategy choice $x_{i'}$ by Player 1, the payoff to Player 2 is $-A(x_{i'}, y_j) = -a_{i'j} \geq -a_{ij}$ as $a_{ij} \geq a_{i'j}$ by virtue of being the maximum of column j . Thus $-V \geq -a_{ij}$ and so $V \leq a_{ij}$. Hence $V = a_{ij}$ and the pure strategies are x_i, y_j . ■

2.4.2 2×2 Games

Clearly, any game described by a 2×2 payoff matrix can be dealt with by theorem 2.4.3. It is easy to check such a small matrix for a Saddle point, and furthermore in the absence of such an element, 2.4.1 can be proven by direct calculation of mixed strategies for each player that yield the same value:

Theorem 2.4.4 (Minimax Theorem for a 2×2 game without a Saddle point). *If the strategic form of a game is given by*

$$\begin{pmatrix} a & b \\ c & d \end{pmatrix}$$

with none of $\{a, b, c, d\}$ a Saddle point, then the game has value

$$V = \frac{ad - bc}{a - b + d - c}$$

Player 1 has a mixed strategy which ensures an average gain of V , regardless of Player 2's strategy. Player 2 has a mixed strategy which ensures an average loss of V , regardless of Player 1's strategy.

Proof. Note first that if $a = b$, then if $c > a = b$ or $d > a = b$ then the smaller of c, d is a Saddle point. Thus both c and d are less than a , so in particular a is a row minimum and a column maximum i.e., a Saddle point. So $a < b$ or $a > b$.

$a < b$ a is a row minimum so it can't be a column maximum. So $c > a$. So c is a column maximum, so it can't be a row minimum. So $d < c$. So d is a row minimum and hence cannot be a column maximum, i.e., $d < b$.

$a > b$ b is a row minimum so it can't be a column maximum. So $d > b$. So d is a column maximum, so it can't be a row minimum. So $c < d$. So c is a row minimum and hence cannot be a column maximum, i.e., $c < a$.

Thus there are two cases

$$\begin{array}{l} \text{Case 1: } \left\{ \begin{array}{l} a < b \\ a < c \\ d < b \\ d < c \end{array} \right. \\ \\ \text{Case 2: } \left\{ \begin{array}{l} a > b \\ a > c \\ d > b \\ d > c \end{array} \right. \end{array}$$

The payoff to Player 1 of a mixed strategy $(p, 1 - p)$ is $ap + c(1 - p)$ or $bp + d(1 - p)$ depending on Player 2's strategy. Proceeding as in 2.3.1, we consider a minimal payoff of

$$V = ap + c(1 - p) = bp + d(1 - p) \tag{1}$$

which arises when

$$(a - c - b + d)p = d - c$$

i.e.,

$$p = \frac{d - c}{(a - b) + (d - c)}$$

If $d - c < 0$ then $d < c$ so we are in Case 1 and $a < b$; otherwise $d - c > 0$, we are in case 2, and $a > b$. Thus in either case $a - b$ and $d - c$ have the same sign, which ensures $0 < p < 1$. Substituting into (1), we obtain

$$\begin{aligned}
 V &= (a - c)p + c \\
 &= \frac{(a - c)(d - c)}{(a - b) + (d - c)} + \frac{c(a - b + d - c)}{(a - b) + (d - c)} \\
 &= \frac{ad - ac - dc + c^2 + ac - bc + dc - c^2}{(a - b) + (d - c)} \\
 &= \frac{(ad - bc)}{(a - b) + (d - c)}
 \end{aligned}$$

As claimed in Theorem 2.4.4.

Similarly, if Player 2 seeks to limit their losses by a mixed strategy $(q, 1 - q)$, Then V, q must satisfy

$$V = ap + c(1 - p) = bp + d(1 - p) \tag{2}$$

Which gives rise to

$$q = \frac{d - b}{(a - b) + (d - c)}$$

Analysis of the two cases confirms $0 < q < 1$ and substituting into (2) gives

$$V = \frac{(ad - bc)}{(a - b) + (d - c)}$$

as before. ■

2.4.3 Dominated strategies

Having given a complete description of strategies for the 2×2 case, it is worth considering whether this technique can be applied to larger matrices. Whilst the analysis offered in the previous sections does not neatly generalise to $m \times n$ matrices, it is occasionally possible to reduce such a matrix to a 2×2 case. To see why this is so, consider the game with payoff matrix given by

$$\begin{pmatrix} -2 & 3 & -1 \\ 3 & -4 & 5 \\ 1 & 4 & 2 \end{pmatrix}$$

In this game, each player has three options. However, Player 1 would be ill-advised to ever play strategy 1 when strategy 3 is available since, regardless of Player 2's choice, the payoff is higher for strategy 3 than strategy 1. Meanwhile, Player 2's third option always offers a larger payoff to Player 1 than column 1, so Player 2 should always prefer their first option. Discarding inferior options, the payoff matrix becomes

$$\begin{pmatrix} 3 & -4 \\ 1 & 4 \end{pmatrix}$$

and we may proceed as for a 2×2 payoff matrix.

We formalise this result as follows:

Definition 2.4.5. For a payoff matrix $A = (a_{ij})$, the i^{th} row *dominates* the k^{th} row if

$$a_{ij} \geq a_{kj} \quad \forall j \in \{1, \dots, n\}$$

The dominance is *strict* if $a_{ij} > a_{kj} \quad \forall j \in \{1, \dots, n\}$.

Definition 2.4.6. For a payoff matrix $A = (a_{ij})$, the j^{th} column *dominates* the k^{th} column if

$$a_{ij} \leq a_{ik} \quad \forall i \in \{1, \dots, m\}$$

The dominance is *strict* if $a_{ij} < a_{ik} \quad \forall i \in \{1, \dots, m\}$.

If the payoff matrix A' is obtained from A by removal of dominated rows and columns, then the value of A' is the value of A , and there is an optimal strategy on A' that remains optimal on A . If all the removed rows and columns were strictly dominated, then the set of optimal strategies is unchanged from A .

Note that after removal of a row or column, a previously undominated row or column may become dominated and thus can itself be removed. Consider, for instance, a slight alteration to the payoff matrix above:

$$\begin{pmatrix} -2 & 3 & 4 \\ 3 & -4 & 5 \\ 1 & 4 & 2 \end{pmatrix}$$

Player 1's strategy 3 no longer dominates strategy 1, since if Player 2 opts for their third strategy, Player 1 is better rewarded by strategy 1. However, this situation will never arise, since strategy 3 is strictly dominated for Player 2 by their first strategy. So we reduce the payoff matrix:

$$\begin{pmatrix} -2 & 3 \\ 3 & -4 \\ 1 & 4 \end{pmatrix}$$

Now it is clear that strategy 3 dominates strategy 1 for Player 1, so we discard the first row:

$$\begin{pmatrix} 3 & -4 \\ 1 & 4 \end{pmatrix}$$

Hence we have arrived at the same game as before- the change made corresponded to an option that was never going to be taken.

2.5 Finite strategic form games as a linear programming problem

The methods of identifying dominated rows or saddle points in payoff matrices are particularly convenient for finding appropriate strategies. There are many more special cases (for instance, matrices where one of the dimensions is 2 can be treated graphically) but these tend to exploit structural features of the matrix in question rather than properties of game-theoretic interest.

Thus we turn our attention to a proof of Theorem 2.4.1 in the general (finite) case of an $m \times n$ payoff matrix. For this it suffices to show that finding an optimal strategy for Player 1 is a linear programming problem. It is clear that if there is any solution, it is finite: the best return Player 1 can guarantee is bounded by $\max a_{ij}$. Further, the dual problem is minimise Player 2's losses, and so by the duality theorem the objectives of each coincide at some value V , which is the value of the game.

To formulate as a linear programming problem, we require an objective function and a set of constraints. If Player 1 is to employ a mixed strategy $\mathbf{P} = (p_1, \dots, p_m)$, to a game with payoff Matrix $A = (a_{ij})$, then the objective is to maximise

$$z = \min_{j \in \{1..n\}} \sum_{i=1}^m p_i a_{ij}$$

However, this is not a linear objective. This can be resolved as follows- introduce a variable x with the constraint $x \leq z$, and maximise x instead. Unravelling the definition of z , this gives n linear conditions on x to be satisfied simultaneously, namely

$$x \leq \sum_{i=1}^m p_i a_{i1}, \quad x \leq \sum_{i=1}^m p_i a_{i2}, \quad \dots, \quad x \leq \sum_{i=1}^m p_i a_{in}$$

i.e., written closer to canonical form,

$$\sum_{i=1}^m p_i a_{ij} - x \geq 0 \quad \forall j \in \{1, \dots, n\}$$

Extra constraints on $\{p_i\}$ also arise. Since \mathbf{P} is a probability distribution, we require

$$p_1 + \dots + p_m = 1$$

Further, for each i , $p_i \geq 0$ and $p_i \leq 1$, since each p_i is to be a probability; this gives another $2m$ conditions.

Thus the set of $2m + n + 1$ constraints on the variables p_1, \dots, p_m, x with linear objective x captures the minimax theorem as a linear programming problem. ■

3 Two player general sum games and co-operation

We now relax the zero-sum condition on games, and consider *general sum* games. This necessitates a move away from Von Neumann's Minimax theorem (and the notion of value) to the more general concept of Nash Equilibria. We will explore the consequences of this new idea in the context of the celebrated prisoner's dilemma.

3.1 Strategic form for two player general sum games

Restricting our attention to two-player games initially, much of the notation of section 2 can be retained, in particular the strategy sets X and Y . However, since Player 2's payoff is no longer immediately determined as the negative of Player 1's, it becomes necessary to consider two payoff matrices.

Definition 3.1.1. The 2 player general sum game where Player 1 has strategy set $X = \{x_1, \dots, x_m\}$ and Player 2 has strategy set $Y = \{y_1, \dots, y_n\}$ such that we can represent Player 1's payoff $A(x, y)$ by a matrix A and Player 2's $B(x, y)$ by a matrix B (both $m \times n$ is described as being a *game in strategic form* X, Y, A, B):

$$\begin{array}{cc} \text{Player 1} & \text{Player 2} \\ \left(\begin{array}{ccc} a_{11} & \cdots & a_{1n} \\ \vdots & & \vdots \\ a_{m1} & \cdots & a_{mn} \end{array} \right) & \left(\begin{array}{ccc} b_{11} & \cdots & b_{1n} \\ \vdots & & \vdots \\ b_{m1} & \cdots & b_{mn} \end{array} \right) \end{array}$$

A more compact notation is the *bimatrix* form:

$$\left[\begin{array}{ccc} (a_{11}, b_{11}) & \cdots & (a_{1n}, b_{1n}) \\ \vdots & & \vdots \\ (a_{m1}, b_{m1}) & \cdots & (a_{mn}, b_{mn}) \end{array} \right]$$

Example 3.1.2. Any zero-sum game in strategic form X, Y, A is a game of strategic form X, Y, A, B Where $B = [-a_{ij}]$ i.e. $B = -A$

Example 3.1.3. In the second form *two finger Morra* becomes

$$\left[\begin{array}{cc} (-2, 2) & (3, -3) \\ (3, -3) & (-4, 4) \end{array} \right]$$

3.2 The prisoner's dilemma

We are now in a position to examine a non-zero sum game. The following is a version of the famous Prisoner's dilemma.

Example 3.2.1 (The prisoner's dilemma). Two prisoners are independently interrogated by the police for a crime they are guilty of, which they may individually either confess to or deny. The authorities lack firm evidence to convict either prisoner of this crime so they can only receive a minor jail sentence for previous misdemeanours if neither confesses. If both prisoners confess to the crime, then there is no doubt over their guilt and each receive the standard sentence. However, as an incentive to confess, the prisoners are told that if they confess whilst the other denies, then they will walk free for their honesty whilst the other is made an example of and receives a still harsher sentence. As a bimatrix, this scenario becomes

$$\left[\begin{array}{cc} (1, 1) & (4, 0) \\ (0, 4) & (3, 3) \end{array} \right]$$

Where row one denotes a confession of guilt by Player 1 (in the standard terminology of the prisoner's dilemma, to *defect*) and row two denotes a denial of guilt by Player 1 (referred to as the strategy of *co-operation*, in the sense that they co-operate with the other prisoner rather than the authorities). Similarly, the first column represents defection by Player 2, and the second co-operation by player 2.

The Prisoner's dilemma is at the heart of many problems in social science. For instance, it can be recast in terms of economics as follows. Two companies may between them control the supply of a particular good, and are able to supply at either a high or low level. Were both to restrict supply by producing at a low level (mutual co-operation), then the price can be kept artificially high. There is insufficient demand in the market to justify high supply (mutual defection) by both companies, and thus if this occurs each will receive a smaller profit. However, restriction of supply by a single company will maintain the higher price point. Hence if one defects (high supply at high price) and the other co-operates (high price, but for a low supply) then the defecting company gains a significant market-share (and profit) advantage over the co-operating one, and so each company is motivated to defect.

3.2.1 Dominated strategies

We can analyse the game from the perspective of Player 1. In the absence of knowledge of Player 2's strategy, it is clear that Player 1 should defect, since row 1 dominates row 2 as follows. Supposing Player 2 defects, then Player 1 is better served by also defecting- both are found guilty of the crime, but Player 1 doesn't incur the penalty for failing to confess. Meanwhile, should Player 2 have co-operated, Player 1 can secure his freedom (and hence the greatest payoff) by defecting.

The situation is exactly the same for Player 2, who must also defect (this is reflected in the payoff matrix by column 1 dominating column 2). Thus we arrive at the scenario of mutual defection, with a payoff to each of 1 (which may be interpreted as years of freedom over the next 5 years, say).

Yet there is a sense in which this behaviour is irrational. Even if motivated entirely by self-interest, the prisoners prefer the outcome of mutual co-operation to mutual defection (a payoff of 3 each, instead of 1 each). This outcome also corresponds to the greater good, the sum of payoffs being greater for mutual co-operation (6) than for exploiting the other prisoner to secure your freedom (4). The problem is that neither may deviate from the defection strategy alone, for to do so simply gets them an extra year in jail whilst their partner dodges a sentence. To secure the benefits of co-operation, the prisoners require some binding arrangement: honour codes like the mafia *omerta* or an external authority such as the state to enforce co-operation. Otherwise a prisoner, acting in accordance with self-interest, maximises their personal gain at the cost of the group: by promising to co-operate to secure the *other* participant's co-operation, then defecting anyway to collect the higher payoff.

To describe mutual defection as the rational strategy, therefore, is to take the *non-cooperative* view of games. In this interpretation, players are either unable to negotiate with each other, or simply unable to trust the promises of the others enough to risk being short-changed by a deceptive partner. Each is therefore best served by a strategy which avoids the potential for exploitation, at a probable cost of being able to co-operate for mutual gain. Such a model of rationality in non-cooperative games is captured by the *Nash equilibrium*.

3.2.2 Nash Equilibria

As before, we denote a mixed strategy for Player 1 by $\mathbf{p} = (p_1, \dots, p_m)$ and for Player 2 by $\mathbf{q} = (q_1, \dots, q_n)$, where

$$\sum_{i=1}^m p_i = 1 = \sum_{j=1}^n q_j$$

to properly define a probability. We can consider these as elements of the vector spaces S_X, S_Y which have the pure strategies as a basis (e.g., $p = p_1x_1 + p_2x_2 + \dots + p_mx_m$ where $x_1, \dots, x_m \in X$ are the pure strategies corresponding to each row). Then the expected return to each player is given by

$$E_1(\mathbf{p}, \mathbf{q}) := \mathbf{p}A\mathbf{q}^T \quad E_2(\mathbf{p}, \mathbf{q}) = \mathbf{q}B^T\mathbf{p}^T$$

Definition 3.2.2. A pair of strategies $(\mathbf{p}, \mathbf{q}) \in S_X \times S_Y$ is a 2-player *Nash equilibrium* for the game given by strategic form X, Y, A, B if neither player gains by unilaterally deviating from the equilibrium. More formally, (\mathbf{p}, \mathbf{q}) is a Nash equilibrium if

$$\begin{aligned} E_1(\mathbf{p}, \mathbf{q}) = \mathbf{p}A\mathbf{q}^T &\geq \tilde{\mathbf{p}}A\mathbf{q}^T = E_1(\tilde{\mathbf{p}}, \mathbf{q}) \quad \forall \tilde{\mathbf{p}} \in S_X \\ E_2(\mathbf{p}, \mathbf{q}) = \mathbf{q}B^T\mathbf{p}^T &\geq \tilde{\mathbf{q}}B^T\mathbf{p}^T = E_2(\mathbf{p}, \tilde{\mathbf{q}}) \quad \forall \tilde{\mathbf{q}} \in S_Y \end{aligned}$$

This is in accordance with our observations on the prisoner's dilemma. Mutual defection is a Nash equilibrium for the game since if a single player moves to co-operation, their expected payoff decreases (from 1 to 0, i.e., they incur another year in jail). Meanwhile, mutual co-operation, whilst preferable to mutual defection, is not a Nash equilibrium since a unilateral decision by one player increases their expected payoff (from 3 to 4). Since both players reach this conclusion, their individually rational changes in strategy from co-operation to defection lead them to mutual defection, a situation that neither prefers.

3.3 Co-operation in the prisoner's dilemma

Surprisingly, it is possible for co-operative behaviour to emerge from a system where it is not *a priori* assumed by the players. There is considerable experimental evidence for people choosing to co-operate more than this game-theoretic analysis would predict- even in the 'one-shot' version of the game. A recent experiment of this kind [2] with university students revealed that defection rates of non-economics majors was under 40%. Whilst economics majors defected 60% of the time in the standard game, when given the opportunity to make (non-binding) deals with the other

participants before play, both categories dropped to a defection rate of around 30%. If one seeks to use game theory to explain behaviour of individuals, these discrepancies between theory and practice must be resolved. Several interpretations are possible.

One way to resolve this problem is to argue that the mutual defection strategy is the correct response to the incorrect problem- that is, the payoff matrix presented is inaccurate in capturing the utility to each player; appropriate modifications will then yield a problem where mutual co-operation is the Nash equilibrium, thus explaining experimental observations.

Typically, such arguments relate to the notion of utility considered in section 2.1 , and introduce an additional cost to defection when the other player co-operates. Mechanisms ranging from guilt, inequality aversion or simple fear of repercussions from external authorities (be it the state or other criminals) have been advanced to explain this psychological difference between the objective payoff and subjective payoff to the defecting player.

A more interesting modification is to preserve the payoff matrix, but consider the relative merits of co-operation. The aversion to co-operation is due to the possibility of receiving a ‘sucker’ payoff of 0 when the other player defects and you co-operate. Obviously for this to occur, one must play against an opponent who opts to defect. Earlier, rational behaviour was considered in terms of average expected outcome over a large number of plays. Since your payoff increases by 2 for mutual co-operation but only decreases by 1 when suckered by a defector, if you expect the probability of another player co-operating to be high enough, it may seem preferable to also co-operate, if only for purely selfish reasons.

We may attempt to formalise this as follows. We introduce a meta-game, whereby a player can assign themselves to one of two categories: category D, members of which always defect; or category C, where the strategy is to defect against players from category D, but to co-operate with others from category C. We can then consider whether it is rational to opt for category C or D.

If the proportion of the population in category C is p , then a player from category D expects a payoff of 1 whereas one from category C expects $3 \times p + 1 \times (1 - p) = 1 + 2p \geq 1$; category C is preferable. This is unsurprising- if one can perfectly identify other co-operators, it pays to co-operate with them. But this analysis has performed a sleight-of-hand: if one can entirely trust other players to co-operate, then there was never any dilemma, since a co-operative game emerges. Without that trust, there is a category superior to C: consider a category C', whose members pretend to be of category C, then defect anyway. This ensures a return of 1 against other defectors, or other players from C', but gains a payoff of 4 rather than 3 from members of category C; who now must factor in the potential of a sucker payoff. In line with the assumptions of non-cooperative game theory, the ability to label oneself as being of a particular category is of no use without a guarantee of honesty.

3.4 The iterated prisoner's dilemma

However, a scenario exists whereby the true nature of another player can be assessed. For this to be the case, the prisoners must play the game multiple times, with knowledge of the previous moves. Thus the ruse of the C' player is no longer of use- with a visible reputation for defection, a player is unlikely to convince another to co-operate. This introduces an incentive to co-operate early (risking a sucker payoff) to benefit from mutual co-operation later. In fact, a strategy along this lines turns out to offer greater expected payoff than persistent defection; co-operation can emerge from a non-cooperative game.

Definition 3.4.1. In an iterated game in which each round consists of a play of the prisoners' dilemma against the same opponent, the *Tit-for-Tat* strategy is:

- In round 1, co-operate.
- In round n for $n \geq 2$, play your opponent's strategy from round $n - 1$.

In 1974, political scientist Robert Axelrod [1] organised a computer simulation of the iterated prisoner's dilemma by inviting game theorists to supply programs implementing their strategy of choice. Each program was run against all the others, itself, and a **random** program which opted for co-operation or defection with equal probability each round. In a preliminary event, Anatol Rapport's Tit-for-Tat program only achieved second place, with victory going to **look ahead**, a program which employed tree-searching techniques similar to those used in computer chess programs. None-the-less, it captured the interest of many of the participants and many of them sought to improve upon tit-for-tat for the main competition. However, it transpired that the most elegant formulation was also the most effective, with Tit-for-Tat scoring higher than any of the other 13 entries to the main event (although, as Axelrod notes, no player submitted the **tit for two tats** program which would have scored higher still against this field).

Theorem 3.4.2. *If the prisoner's dilemma is iterated for an indefinite number of rounds with the probability of a successive round being sufficiently high (i.e., a large number of iterations are carried out) then Tit-for-Tat is a Nash equilibrium.*

Notice the stipulation on an unknown number of rounds: were the precise timing of the final round known, then that round could be played without fear of repercussions and thus reduces to the single iteration case, with defection being the rational behaviour in that round. Further, with too few iterations a strategy of initial co-operation will be unable to catch up with a defection strategy that succeeds in suckering a co-operating player.

Proof. In [4] it is demonstrated that there are only three types of response to a tit-for-tat player t :

- c , co-operation in all future rounds.
- d , defect in all future rounds.
- a , alternate defection and co-operation.

Note that c is equivalent to playing t against t , since mutual co-operation will arise in this scenario. Thus to show that t is the best response to t (that is, a Nash equilibrium), we require that

$$E_1(t, t) \geq E_1(d, t) \text{ and } E_1(t, t) \geq E_1(a, t)$$

In the context of an iterated game, (rather than the strategic forms considered earlier) we can calculate the expected returns E_1 as follows. Suppose a single round is played, and subsequent iterations occur with fixed probability p . Then the expected number of rounds is given by the geometric sum $1 + p + p^2 + p^3 + \dots = \frac{1}{1-p}$ and so $E_1(t, t) = E_1(c, t) = \frac{3}{1-p}$. For a defection strategy d , there is a return of 4 in the first round, then a return of 1 for all subsequent rounds (as tit-for-tat retaliates by continually defecting). Thus

$$E_1(d, t) = 4 \times 1 + 1 \times (p + p^2 + p^3 + \dots) = 4 + p(1 + p + p^2 + \dots) = 4 + p \frac{1}{1-p} = 4 + \frac{p}{1-p}$$

For the alternating strategy a , the participants are perpetually out of sync and thus alternate between sucker payoffs of 4 and 0. This, via a geometric progression in p^2 , gives rise to expectation:

$$E_1(a, t) = 4 \times 1 + 0 \times p + 4 \times p^2 + \dots = 4 \times (1 + p^2 + p^4 + \dots) = 4 \times \frac{1}{1-p^2} = \frac{4}{1-p^2}$$

Hence (see figure 1) if $p > \frac{1}{3}$,

$$E_1(t, t) \geq E_1(d, t) \text{ and } E_1(t, t) \geq E_1(a, t)$$

as desired. ■

3.4.1 Group interest

In 2003-05 a series of iterated prisoner's dilemma competitions was organised by The Congress on Evolutionary Computation Conference (CEC'04) to celebrate the 20th anniversary of Axelrod's original tournament. Nick Jennings and Gopal Ramchurn of Southampton University developed a system to beat tit-for-tat by exploiting the possibility of entering multiple programs. They recognised that if your goal is to achieve the highest score by any one program, then you can do so by introducing helper programs which sacrifice their score for the benefit of the master program. By submitting over 60 programs (223 competed in total), the Southampton team was able to secure the top three scores in the competition (but also many of the lowest).

Simply introducing programs that always co-operated would be insufficient for this task, since any other program that defected would also benefit for a free ride, whilst tit-for-tat would maintain its virtuous circle of mutual co-operation and hence still score well. Thus, each program was designed to use the first few plays to communicate its identity to other Southampton programs by means of a set sequence of co-operation and defection. On achieving a match, the helper programs would proceed to co-operate every round whilst the master program defected, thus collecting the highest possible payoff. When a Southampton program was paired against another team (which would be unable to produce the appropriate sequence of moves), it switched to constant defection to drive down the score of tit-for-tat style strategies. [3]

Thus the goal of the team could be achieved at the expense of some of its players. Often, this is unworkable in real world scenarios- what player would be willing to assume the role of the helper program? However, similar situations can occur in business environments. For instance, the Starbucks corporation found that its strategy of introducing many branches in the same area caused cannibalisation between stores, with each seeing a decline in profits, yet ultimately boosted overall profits. This was because the combined presence of multiple Starbucks forced independent coffee houses out of business. This would allow for Starbucks to monopolise the local demand for coffee, after which it could close the least popular stores. Thus market share is maximised (analogously, master program score), at the expense of both the competition (tit-for-tat's score) and some of their own stores (the helper programs). [6]

In essence, these results show that an optimal strategy for the iterated prisoner's dilemma is not necessarily optimal for an iterated prisoner's dilemma tournament. We can consider such a tournament, with its series of two player games, as being a larger n-player game, where each player fixes a general strategy to play in each sub-game. Then tit-for-tat style programs can be seen as forming a group where each benefits from membership, whilst the members of the Southampton group will not all benefit but can ensure greater success for some of the members. These group dynamics are as relevant as the iterated nature of the sub-game, and we investigate them further in later sections.

4 n-player games

4.1 Strategic form for n-player general sum games

Just as relaxing the zero-sum condition increased the complexity of notation for strategic form (from a payoff matrix to a bimatrix), so increasing the number of players necessitates a more general description of the payoff structure. Recall (from definitions 2.2.1 and 2.2.3) that the entries a_{ij} of the payoff matrix corresponded to the values of a payoff function $A : X \times Y \rightarrow \mathbb{R}$. This can be usefully generalised as follows:

Definition 4.1.1. A *finite n-player game in strategic form* X_1, \dots, X_n consists of strategy sets X_1, \dots, X_n corresponding to the players; and real valued functions $a_1 \dots a_n : X_1 \times \dots \times X_n \rightarrow \mathbb{R}$ such that the payoff to player i when the strategies chosen by each player j is $x_j \in X_j$ is given by $a_i(x_1, \dots, x_n)$.

As in section 3.2.2, each strategy space X_i , consisting of strategies $\{x_1, \dots, x_{m_i}\}$ will give rise to a space of mixed strategies S_{X_i} consisting of m_i -tuples $(p_{i1}, \dots, p_{im_i})$ which describe the probability of player i picking strategy $x_j \in X_i$ as p_{ij} . For these probabilities to be valid, we additionally require that

$$\sum_{j=1}^{m_i} p_{ij} = 1 \quad \forall i$$

Further, we can simplify notation by identifying X_i with the integers $\{1, \dots, m_i\}$ - that is, playing "1" corresponds to a strategy of x_1 and so on.

Then the expected return to player j when each player has mixed strategy \mathbf{p}_j is given by

$$E_j(\mathbf{p}_1, \dots, \mathbf{p}_n) = \sum_{i_1=1}^{m_1} \cdots \sum_{i_n=1}^{m_n} p_{1i_1} \cdots p_{ni_n} a_j(i_1, \dots, i_n)$$

4.2 Nash's Theorem

Characterising expectation as in the previous section, definition 3.2.2 generalises:

Definition 4.2.1. A set of mixed strategies $(\mathbf{p}_1, \dots, \mathbf{p}_n) \in S_{X_1} \times \dots \times S_{X_n}$ is an n -player *Nash equilibrium* for the game given by strategic form X_1, \dots, X_n if no player gains by unilaterally deviating from the equilibrium.

More formally, if $\forall i = 1, \dots, n$ and $\forall \tilde{\mathbf{p}} \in S_{X_i}$,

$$E_j(\mathbf{p}_1, \dots, \mathbf{p}_{i-1}, \mathbf{p}_i, \mathbf{p}_{i+1}, \dots, \mathbf{p}_n) \geq E_j(\mathbf{p}_1, \dots, \mathbf{p}_{i-1}, \tilde{\mathbf{p}}, \mathbf{p}_{i+1}, \dots, \mathbf{p}_n)$$

Then $(\mathbf{p}_1, \dots, \mathbf{p}_n)$ is a Nash equilibrium (since no player j benefits from changing from strategy \mathbf{p}_j to any other strategy $\tilde{\mathbf{p}}$).

Several Nash Equilibria have been encountered so far. Existence of a unique Nash equilibrium in any two player zero-sum game follows from the Minimax theorem, which was demonstrated directly for the 2×2 case and follows by linear programming techniques for games with larger strategy sets. The Prisoner's dilemma gave rise to a Nash equilibrium (mutual defection), as did the iterated version, indicating that the zero-sum condition is not necessary. In his 1950 dissertation, John Nash proved that much more was true:

Theorem 4.2.2. (*Nash's theorem*) Any finite n -player game in strategic form has a Nash equilibrium.

Before attempting a proof, it is worthwhile considering some limitations of the result. Section 3.2.1 considered the rationality of Nash equilibria. Further, the theorem makes no claim of uniqueness, and cannot, as the following example shows.

Example 4.2.3. (The Stag-Hunt) The *Stag-Hunt* game is given in Bimatrix form by

$$\begin{bmatrix} (1, 1) & (2, 0) \\ (0, 2) & (3, 3) \end{bmatrix}$$

Whilst this appears similar to the prisoner's dilemma, the reduced payoffs in the off-diagonal entries remove the incentive to defect for selfish gain. Hence mutual co-operation (a play of $\mathbf{p} = (0, 1)$ by each player) becomes a Nash equilibrium- for player 1, and thus player 2 by symmetry:

$$E_1((0, 1), (0, 1)) = 3 \geq 3 - p = 2p + 3(1 - p) = E_i((p, 1 - p), (0, 1)) \forall p \in [0, 1)$$

That is, unilateral change of strategy from $(0, 1)$ to any other $(p, 1 - p)$ for $p \geq 0$ is disadvantageous.

However, mutual defection (each player using strategy $(1, 0)$) continues to be a Nash equilibrium, since then

$$E_1((1, 0), (1, 0)) = 1 \geq p = E_1((p, 1 - p), (1, 0)) \forall p \in [0, 1]$$

There is also a third Nash equilibrium given by mixed strategies of $\mathbf{p}_1 = \mathbf{p}_2 = (\frac{1}{2}, \frac{1}{2})$ for an expected payoff of $\frac{3}{2}$.

4.2.1 Proof of Nash's Theorem

We first fix notation and introduce some useful results.

Definition 4.2.4. A pure strategy for player k of their i th strategy is denoted by δ_i , i.e., the mixed strategy \mathbf{p}_k consisting of a one in the i th component and zeroes elsewhere. Then, given a set of mixed strategies $\mathbf{p} = (\mathbf{p}_1, \dots, \mathbf{p}_n)$ for all the players, we may define the *expected payoff to Player k by switching to pure strategy i* by

$$E_{k[i]}(\mathbf{p}_1, \dots, \mathbf{p}_n) = E_k(\mathbf{p}_1, \dots, \mathbf{p}_{k-1}, \delta_i, \mathbf{p}_{k+1}, \dots, \mathbf{p}_n)$$

More compactly, we denote this by $E_{k[i]}(\mathbf{p})$.

Lemma 4.2.5. The $E_{k[i]}(\mathbf{p})$ determine $E_k(\mathbf{p})$ by the weighted average $E_k(\mathbf{p}) = \sum_{i=1}^{m_k} p_{ki} E_{k[i]}(\mathbf{p})$

Proof.

$$\mathbf{p}_k = (p_{k1}, \dots, p_{km_k}) \Rightarrow \mathbf{p}_k = p_{k1}\delta_1 + \dots + p_{km_k}\delta_{m_k} = \sum_{i=1}^{m_k} p_{ki}\delta_i$$

So

$$\begin{aligned} E_k(\mathbf{p}) &= E_k(\mathbf{p}_1, \dots, \mathbf{p}_{k-1}, \mathbf{p}_k, \mathbf{p}_{k+1}, \dots, \mathbf{p}_n) \\ &= E_k(\mathbf{p}_1, \dots, \mathbf{p}_{k-1}, \sum_{i=1}^{m_k} p_{ki}\delta_i, \mathbf{p}_{k+1}, \dots, \mathbf{p}_n) \\ &= \sum_{i=1}^{m_k} p_{ki} E_k(\mathbf{p}_1, \dots, \mathbf{p}_{k-1}, \delta_i, \mathbf{p}_{k+1}, \dots, \mathbf{p}_n) \\ &= \sum_{i=1}^{m_k} p_{ki} E_{k[i]}(\mathbf{p}) \end{aligned}$$

■

Lemma 4.2.6. A set of mixed strategies $\mathbf{p} = (\mathbf{p}_1, \dots, \mathbf{p}_n)$ is a Nash equilibrium if

$$E_{k[i]}(\mathbf{p}) \leq E_k(\mathbf{p}) \quad \forall k \in \{1, \dots, n\} \quad \forall i \in X_k$$

Proof. Suppose Player k unilaterally changes from strategy \mathbf{p}_k to an alternative $\tilde{\mathbf{p}}_k = (\tilde{p}_{k1}, \dots, \tilde{p}_{km_k})$. Let $\tilde{\mathbf{p}} = (\mathbf{p}_1, \dots, \mathbf{p}_{k-1}, \tilde{\mathbf{p}}_k, \mathbf{p}_{k+1}, \dots, \mathbf{p}_n)$ be the new strategy set. Then by lemma 4.2.5

$$E_k(\tilde{\mathbf{p}}) = \sum_{i=1}^{m_k} \tilde{p}_{ki} E_{k[i]}(\tilde{\mathbf{p}})$$

But $E_{k[i]}(\tilde{\mathbf{p}}) = E_k(\mathbf{p}_1, \dots, \mathbf{p}_{k-1}, \delta_i, \mathbf{p}_{k+1}, \dots, \mathbf{p}_n) = E_{k[i]}(\mathbf{p})$ so

$$\begin{aligned} E_k(\tilde{\mathbf{p}}) &= \sum_{i=1}^{m_k} \tilde{p}_{ki} E_{k[i]}(\mathbf{p}) \\ &\leq \sum_{i=1}^{m_k} \tilde{p}_{ki} E_k(\mathbf{p}) \quad \text{by hypothesis} \\ &= E_k(\mathbf{p}) \sum_{i=1}^{m_k} \tilde{p}_{ki} \quad \text{since } E_k(\mathbf{p}) \text{ independent of } i \\ &= E_k(\mathbf{p}) \quad \text{since } \mathbf{p}_k \text{ a probability distribution} \end{aligned}$$

Thus Player k does not gain by unilaterally deviating from the strategy \mathbf{p}_k . Since k was arbitrary, this makes \mathbf{p} a Nash equilibrium. \blacksquare

The following result from topology is also needed:

Theorem 4.2.7. (*Brouwer's Fixed point theorem*) *If $\emptyset \neq D \subset \mathbb{R}^n$ (for some n) is compact and convex, then for any continuous function $f : D \rightarrow D$ there exists a fixed point of f , i.e., $\exists z \in D$ s.t. $f(z) = z$.*

We are now in a position to prove Nash's theorem:

Proof. (Nash's theorem)

Let k be arbitrary. Then S_{X_k} is a compact convex subset of \mathbb{R}^{m_k} and hence $D = S_{X_1} \times \dots \times S_{X_n}$ is a compact convex subset of \mathbb{R}^M for $M = \sum_{i=1}^n m_i$. To make use of Brouwer's Theorem, define $f : D \rightarrow D$ by $f(\mathbf{p}) = \mathbf{p}' = (\mathbf{p}'_1, \dots, \mathbf{p}'_n)$ where

$$p'_{ki} = \frac{p_{ki} + \max(0, E_{k[i]}(\mathbf{p}) - E_k(\mathbf{p}))}{1 + \sum_{j=1}^{m_k} \max(0, E_{k[j]}(\mathbf{p}) - E_k(\mathbf{p}))}$$

Notice that \mathbf{p}' is a well defined element of D since each $p'_{ki} \geq 0$ and the denominator ensures that their sum is 1. Further, f is continuous by the continuity of expectation $E_k(\mathbf{p})$. Hence theorem 4.2.7 gives a fixed point $\mathbf{q} = (\mathbf{q}_1, \dots, \mathbf{q}_n) \in D$ such that $f(\mathbf{q}) = \mathbf{q}$. So

$$q_{ki} = \frac{q_{ki} + \max(0, E_{k[i]}(\mathbf{q}) - E_k(\mathbf{q}))}{1 + \sum_{j=1}^{m_k} \max(0, E_{k[j]}(\mathbf{q}) - E_k(\mathbf{q}))} \quad \forall k \in \{1, \dots, n\} \forall i \in \{1, \dots, m_n\} \quad (3)$$

From Lemma 4.2.5 $E_k(\mathbf{q})$ is a weighted average of the $E_{k[i]}(\mathbf{q})$ with each weight at most 1, so the $E_{k[i]}(\mathbf{q})$ cannot all be greater than $E_k(\mathbf{q})$ (Fubini principle). Hence, for at least one i such that $q_{ki} > 0$, $E_{k[i]}(\mathbf{q}) \leq E_k(\mathbf{q})$, which means that $\max(0, E_{k[i]}(\mathbf{q}) - E_k(\mathbf{q})) = 0$ for that i .

But then the numerator of (3) is simply the LHS, q_{ki} , and hence the denominator must be 1. This forces

$$\sum_{j=1}^{m_k} \max(0, E_{k[j]}(\mathbf{q}) - E_k(\mathbf{q})) = 0$$

$$\Rightarrow E_{ki}(\mathbf{q}) \leq E_k(\mathbf{q}) \quad \forall k \text{ and } \forall i$$

Hence by Lemma 4.2.6, \mathbf{q} is a Nash equilibrium. ■

4.3 The 3-player simple majority game

Having discovered that any n -player finite game will offer at least one rational solution by way of a Nash equilibrium, it is tempting to conclude that such games offer no further surprises. However, consideration of even a simple 3-player example will show that (as for the move from zero-sum to general sum games) the introduction of additional players leads to qualitatively different behaviour.

4.3.1 Simple Majority

We will consider the *simple majority* game. This is formulated in section 21 of [5] as follows:

Example 4.3.1. Each player, by a personal move, chooses the number of one of the two other players. Each one makes his choice uninformed about the choices of the two other players. After this payments will be made as follows: If two players have chosen each other's numbers we say that they form a *couple*. Clearly there will be precisely one couple, or none at all. If there is precisely one couple, then the two players who belong to it get one half unit each, while the third (excluded) player correspondingly loses one unit. If there is no couple, then no one gets anything.

Analysed by way of Nash Equilibria, it can be seen that the only irrational situation would be for no couple to form. Denoting a choice of Player a by Player 1, Player b by Player 2 and Player c by Player 3 as the triple (a, b, c) , this would correspond to a set of strategies $(2, 3, 1)$ or $(3, 1, 2)$ with a payoff to each of 0. Any other set would create a couple between two players and leave the third excluded. The six remaining states $(2, 1, 2)$, $(2, 1, 1)$, $(3, 1, 1)$, $(3, 3, 1)$, $(2, 3, 2)$, and $(3, 3, 2)$ are *all* Nash equilibria since they contain couples. This is because the excluded player cannot escape their fate by a unilateral change of strategy, since either strategy fails to match them to a Player who chose them; whilst a change of strategy by a coupled player either places them in a new couple (such as Player 1 moving from $(3, 1, 1)$ to $(2, 1, 1)$) or reduces their payoff from $1/2$ to 0 (such as Player 1 moving from $(3, 3, 1)$ to $(2, 3, 1)$): Hence no player gains an advantage from unilateral changes in strategy. Of course, a couple may still fail to emerge since without binding agreements no player can tell who they should attempt to form a couple with.

As Von Neumann and Morgenstern note, this simple example gives an argument against *laissez faire* Capitalism- despite the absolute, formal fairness of this game due to the symmetry of its rules, there is no reason why the usage of those rules by the participants, nor the outcomes they receive, will be fair. Worse, the players are driven to create unsymmetric outcomes since it is rational and advantageous for them to form couples.

We can view the formation of a couple in the simple majority game as a zero sum payoff by considering the couple as a single entity. Then the couple gains what the excluded player loses; in this case, a payoff of 1 unit. The issue then is how that payoff should be divided between the couple. It was observed that in the above formulation, players are driven to form a couple, but they are indifferent as to which player they form it with. This was because the distribution of the payoff was fair within the couple, regardless of which arises. If this condition is relaxed, then players may gain an incentive to form a particular couple; although which is not immediately obvious.

For instance, suppose the payoff is altered in the event of a couple forming between Player 1 and Player 2 to $1/2 + \epsilon$ and $1/2 - \epsilon$ respectively. Then Player 1 may appear to favour forming a couple with Player 2 to one with Player 3. However, if they expect Player 2 to behave rationally, then they must conclude that Player 2 will opt instead to try and form a couple with Player 3, for a payoff of $1/2$ instead of $1/2 - \epsilon$. Hence Player 1 must also choose Player 3, and settle for a payoff of $1/2$ or risk becoming the excluded player and incurring a cost of 1.

In such a scenario, Player 3 is then a favoured partner, and can form a couple with the player of their choice. Thus it is unreasonable to assume that they will settle for a payoff of $1/2$ since they can safely threaten to form a couple with the other player. In effect, Players 1 or 2 will need to buy their way into a partnership with Player 3; Player 2 will accept any extra price ϵ' that is less than the ϵ it would cost them to form a partnership with Player 1, for instance.

Thus a range of behaviours can arise which depend upon both the payoff to the couple, and the distribution within the couple, which could be subject to complex bargaining between the players. Such scenarios are particularly relevant to economic problems, and to model them more precisely than with Nash equilibria we revisit and formalise the ideas of cooperation first considered in section 3.3. The next section builds upon the notion of couple in the simple majority game to develop a theory of coalitions and the payoff to individuals within them.

5 Co-operative games and coalitions

5.1 Coalitional form for n -player games

We now examine cooperative games, that is, ones in which players may enter into binding arrangements. Each remains motivated by their individual utility payoff, and thus can be expected to only enter an agreement that is personally beneficial to them. However, we will also allow for transferable utility. That is, the participants in a coalition may redistribute the total return to the coalition

amongst themselves, rather than keeping to the individual returns prescribed by the game. In effect, this allows for side payments from one player to another, to give the second an incentive to join a coalition with the first. This, of course, assumes that the payoffs to each player are in equivalent units, and represent a transferable commodity. In a genuine prisoner's dilemma, for instance, neither participant can accept jail time for the other; although mutual cooperation will still arise as their best strategy.

We thus require two things: a rule for determining the return to any coalition; and a means to decide which players will enter into the coalition. Broadly speaking, the payoffs to the coalitions constitute the rules of the game, analogous to the payoff matrices/functions in strategic form games; whilst the formation of coalitions represent the plays (strategies).

Definition 5.1.1. A finite n -player game in coalitional form (X, v) consists of the set of players $X = \{1, 2, \dots, n\}$ and a characteristic function $v : \mathcal{P}(X) \rightarrow \mathbb{R}$ satisfying

- $v(\emptyset) = 0$
- If $S \cap T = \emptyset$ then $v(S) + v(T) \leq v(S \cup T)$

Here $\mathcal{P}(X)$ denotes the *power set* of X , that is, the set of all possible subsets, and thus the characteristic function describes the payoff to any coalition that may form between the players. Note that this is a set of size 2^n . The first condition on v ensures that the *empty coalition* has no value. This is a mathematical formality for later results, since a player cannot enter this coalition: even if Player i shuns all the others, they will be in the coalition $\{i\}$. The second condition, known as *superadditivity*, is considered a natural property: if two distinct coalitions work together, then they should receive at least as much as they did by working independently.

It is an immediate consequence of the superadditivity that the greatest group return is achieved by forming the *grand coalition* X consisting of all n players. However, it is not a given that the grand coalition (or indeed any coalition) will form. Whilst $v(X) \geq v(\{i\})$, Player i would have to be offered at least a return of $v(\{i\})$ from the division of $v(X)$ to be convinced to join. In particular, if $v(\{i\}) < v(X)/n$ then the 'fair' payoff of equal distribution will not interest Player i . Thus the determination of rules for allocation is not without subtlety.

First however, we revisit some earlier examples in the context of coalitional form.

Example 5.1.2. Recall example 3.1.3, the bimatrix form of Two-finger Morra. As a two player game, there are four possible coalitions- $\emptyset, \{1\}, \{2\}$ or $\{1, 2\}$. $v(\emptyset) = 0$ is given, and since Two-finger Morra is zero-sum the value gives that $v(\{1\}) = 1/12$ and $v(\{2\}) = -1/12$. Interpreted directly, $v(\{1, 2\})$ is the total return to the coalition when, working together, the players are able to select any entry in the bimatrix. Since the game is zero-sum, this is always zero (which can be verified by inspection).

Notice that superadditivity holds, and the grand coalition offers a return of 0. Since Player 1 receives a payoff of 1/12 by not entering the grand coalition, Player 2 would have to offer a side payment of at least 1/12 to entice Player 1 into a coalition. But it would be irrational for Player 2 to offer any more than 1/12 to create the coalition, since going it alone only costs him 1/12. Hence the distribution within the grand coalition would be the same as if it did not form: players are indifferent to the formation of a coalition, and the coalitional form precisely mimics the strategic form. This motivates some additional definitions.

Definition 5.1.3. A game in coalitional form is said to be of *constant-sum* if $v(S)+v(X\setminus S) = v(X)$ for all $S \in \mathcal{P}(X)$. If additionally $v(X) = 0$, the game is described instead as *zero-sum*.

Definition 5.1.4. A game in coalitional form is *inessential* if $\sum_{i=1}^n v(\{i\}) = v(X)$. Otherwise, the game is *essential*.

Corollary 5.1.5. Any two person zero-sum game is inessential.

Example 5.1.6. As a coalitional game, the Prisoner's dilemma as formulated in example 3.2.1 is given by $v(\emptyset) = 0$, $v(\{1\}) = v(\{2\}) = 1$, $v(\{1,2\}) = 6$. Hence, the Prisoner's dilemma is an essential coalitional game.

Example 5.1.7. Individual payoffs in the simple majority game with symmetric distribution (example 4.3.1) can be described by the following table, where the rows denote the options for Player 1 and the columns the options for Players 2 and 3 (as an ordered pair):

	(1, 1)	(1, 2)	(3, 1)	(3, 2)
2	(1/2, 1/2, -1)	(1/2, 1/2, -1)	(0, 0, 0)	(-1, 1/2, 1/2)
3	(1/2, -1, 1/2)	(0, 0, 0)	(1/2, -1, 1/2)	(-1, 1/2, 1/2)

The value $v(\{1,2,3\})$ is determined by free choice of any of the 8 strategy combinations, but for any such choice the sum of the payoffs is 0. As always, $v(\emptyset) = 0$. To determine $v(\{1\})$ (and by symmetry $v(\{2\})$ and $v(\{3\})$) we can consider Players 2 and 3 as acting as a single entity against Player 1; the return to their coalition being the sum of the individual returns, which they are motivated to drive as high as possible. It should be clear then that they will form a couple to ensure a return of 1, forcing a payoff of -1 onto Player 1, ensuring that $v(\{1\}) = -1$. However, it is enlightening to see precisely why this occurs.

By treating Players 2 and 3 as a single player with strategy set A, B, C, D the above table reduces to

	A	B	C	D
2	(1/2, -1/2)	(1/2, -1/2)	(0, 0)	(-1, 1)
3	(1/2, -1/2)	(0, 0)	(1/2, -1/2)	(-1, 1)

This is a bimatrix for a 2 player strategic form game; moreover it is the bimatrix of the zero-sum game given by

$$\begin{pmatrix} 1/2 & 1/2 & 0 & -1 \\ 1/2 & 0 & 1/2 & -1 \end{pmatrix}$$

Further, column 4 dominates all other columns, so the strategic form reduces to that column- giving the game a value of -1 for Player 1 (and hence of 1 for the coalition of Players 2 and 3).

So the simple majority game is described in coalitional form by

- $v(\emptyset) = v(X) = 0$
- $v(\{1\}) = v(\{2\}) = v(\{3\}) = -1$
- $v(\{1, 2\}) = v(\{1, 3\}) = v(\{2, 3\}) = 1$

Superadditivity holds, since $v(\{i\}) + v(\{j, k\}) = -1 + 1 = 0 \leq 0 = v(\{i, j, k\})$ for any permutation i, j, k of the players. Further, the game is zero-sum; but it is essential, since

$$v(\{1\}) + v(\{2\}) + v(\{3\}) = -3 < v(X)$$

Hence corollary 5.1.5 does not generalise to any zero-sum game; it is specific to the two-player case.

5.1.1 Coalitional form of a strategic form game

The construction of the characteristic function used in the preceding example can be used, after suitable generalisation, for any strategic form game. Given a coalition $S \in \mathcal{P}(X)$, we consider a two-player zero-sum game between two teams S and $X \setminus S$. The strategy sets for each team consists of the cartesian product of the strategy sets of the individual members of each team. For instance, in example 5.1.7 Players 2 and 3 had strategy sets $\{1, 3\}$ and $\{1, 2\}$ respectively giving rise to the combined strategy set $\{A, B, C, D\} = \{(1, 1), (1, 2), (3, 1), (3, 2)\}$. The payoff to the coalition for any given combination of strategies is then determined by the sum of the payoffs to its members from those strategies. $v(S)$ is then determined by the value of the game, which (due to the Minimax theorem) exists and can be found by the techniques discussed in section 2.

5.1.2 S -veto games

Of particular interest are the class of coalitional games known as S -veto games. In these, a coalition is only effective if some subset S of the players are all members. Thus this gives rise to characteristic functions of the form

$$w_S(T) = \begin{cases} 1 & S \subseteq T \\ 0 & S \not\subseteq T \end{cases}$$

For instance, $S = \{1\}$ gives rise to a dictatorship by Player 1 (an inessential game) whereas $S = X$ forces the grand coalition to form for any player to receive a payoff- or, considered as a voting system, a unanimous verdict. More complicated voting arrangements can be built upon veto systems, such as the United Nations Security Council system of “great power unanimity” which requires the support of all five permanent members (and any four of the ten non-permanent members) to pass major resolutions.

5.2 Shapley Value

Given a coalitional form game (X, v) we seek to construct a value $\Phi(X, v) \in \mathbb{R}^n$ for the game, where the component $\Phi_i(X, v)$ denotes the payoff to player i . This can be interpreted as a measure of the power of player i in the game, since it indicates their contribution to the grand coalition, were one to form.

An example of such a value is the *Shapley value*, constructed as follows. Given a permutation $\pi \in S_n$ (i.e., a bijection from X to X) we can consider the players as forming the coalition one by one, in accordance with the ordering created by π . Thus the coalition is built by considering first the coalition consisting of Player $\pi(1)$, then of players $\pi(1)$ and $\pi(2)$, and so on. By superadditivity, the value at each step either increases or remains constant, so we may assign to each player a non-negative payoff equal to this increase. Let p_π^i denote the set of players who joined the coalition before Player i . Therefore $p_\pi^i = \{j | \pi(j) < \pi(i)\}$. Then the value of Player i to the coalition is given by $v(p_\pi^i \cup \{i\}) - v(p_\pi^i)$.

However, it is unlikely that the payoff constructed in this way will be independent of the ordering π . Thus we consider the value of a player to be their average contribution to the formation of a coalition; where any ordering of players is equally likely. That is,

$$\Phi_i(X, v) = \frac{1}{n!} \sum_{\pi \in S_n} v(p_\pi^i \cup \{i\}) - v(p_\pi^i)$$

since the size of S_n is the number of permutations of the set of players X , namely $n!$, and for any given permutation we may determine Player i 's contribution by the method in the previous paragraph.

5.3 Shapley Axioms

The above construction, determining average contributions to a coalition, is intuitively fair. However, the notion of fairness can be made rigorous by requiring the satisfaction of a number of axioms, some of which have been hinted at in earlier discussion. A number of desirable properties for a value can be advanced.

For group rationality, the total value of the players should be the value of the grand coalition. Thus, to assign a value to each player in X we require that

$$\sum_{i=1}^n \Phi_i(X, v) = v(X)$$

A value should not *a priori* favour any particular player over another. That is, should the return to any coalition featuring Player i and not Player j be the same as the return to the coalition with Player i replaced by Player j , then the values of players i and j should be equal.

Any player whose presence in a coalition does not alter its payoff should receive a value of 0.

Given games (X, v) and (X, w) , then we can define the game $(X, v+w)$ as $(v+w)(S) = v(S) + w(S)$. Logically, we should require that $\Phi(X, v+w) = \Phi(X, v) + \Phi(X, w)$, i.e., that the return of playing the sum of two games is the sum of the returns of each game.

These four properties can be summarised as follows.

Definition 5.3.1. The *Shapley axioms* for games in coalitional form are

- Efficiency: $\sum_{i=1}^n \Phi_i(X, v) = v(X)$
- Symmetry: If $v(S \cup \{i\}) = v(S \cup \{j\}) \forall S \in \mathcal{P}(X) \setminus \{i, j\}$ then $\Phi_i(X, v) = \Phi_j(X, v)$.
- Dummy: If $v(S \cup \{i\}) - v(S) = 0 \forall S \in \mathcal{P}(X)$ then $\Phi_i(X, v) = 0$
- Additivity: $\Phi(X, v+w) = \Phi(X, v) + \Phi(X, w)$ for any games (X, v) and (X, w)

Theorem 5.3.2. A function satisfying the Shapley axioms always exists.

Proof. The function Φ constructed in section 5.2 satisfies the Shapley axioms:

We can recursively define p_π^i via

$$p_\pi^i = \begin{cases} \emptyset & i = 1 \\ p_\pi^{i-1} \cup \{i\} & i \geq 2 \end{cases}$$

Further, we can consider $q_\pi^i = p_\pi^i \cup \{i\}$; in particular, $q_\pi^n = X$. Thus

$$\Phi_i(X, v) = \frac{1}{n!} \sum_{\pi \in S_n} v(q_\pi^i) - v(p_\pi^i)$$

So

$$\sum_{i=1}^n \Phi_i(X, v) = \sum_{i=1}^n \frac{1}{n!} \sum_{\pi \in S_n} v(q_\pi^i) - v(p_\pi^i)$$

$$\begin{aligned}
&= \frac{1}{n!} \sum_{\pi \in S_n} \sum_{i=1}^n \left(v(q_\pi^i) - v(p_\pi^i) \right) \\
&= \frac{1}{n!} \sum_{\pi \in S_n} \left(\sum_{i=2}^n \left(v(q_\pi^i) - v(p_\pi^i) \right) + \left(v(q_\pi^1) - v(p_\pi^1) \right) \right) \\
&= \frac{1}{n!} \sum_{\pi \in S_n} \left(\sum_{i=2}^n v(q_\pi^i) - \sum_{i=2}^n v(q_\pi^{i-1}) + \left(v(q_\pi^1) - v(p_\pi^1) \right) \right) \\
&= \frac{1}{n!} \sum_{\pi \in S_n} \left(v(q_\pi^n) - v(q_\pi^{n-1}) + v(q_\pi^1) - v(\emptyset) \right) \\
&= \frac{1}{n!} \sum_{\pi \in S_n} \left(v(X) - v(q_\pi^1) + v(q_\pi^1) - 0 \right) \\
&= \frac{1}{n!} \sum_{\pi \in S_n} v(X) = \frac{1}{n!} n! v(X) = v(X)
\end{aligned}$$

Thus efficiency holds. Symmetry and the dummy axiom are immediate from the definition of $\Phi_i(X, v)$, whilst additivity follows from the linearity of Σ and the averaging process. ■

Theorem 5.3.3. *The Shapley function is unique.*

Proof. The S -veto games described in section 5.1.2 are a basis for the set of coalitional games: Note first that the value of an S -veto game is completely determined from the Shapley axioms from definition 5.3.1. For a given characteristic function w_S the dummy axiom ensures that $\Phi_i(X, w_S) = 0$ for any $i \notin S$; whilst the symmetry axiom ensures that if $i, j \in S$ then $\Phi_i(X, w_S) = \Phi_j(X, w_S)$, that is, the members of the veto set S have equal value. Since (by the efficiency axiom) the sum of their values is the grand coalition payoff $w_S(X) = 1$, it follows that the individual values for members of S are $\frac{1}{|S|}$. By the same reasoning for $w_S(X) = c$ for an arbitrary constant c , we deduce

$$\Phi_i(X, cw_S) = \begin{cases} \frac{c}{|S|} & i \in S \\ 0 & i \notin S \end{cases}$$

Now for an arbitrary characteristic function v , consider the set of constants c_T for $T \in \mathcal{P}(X)$ constructed inductively on the size of T by $c_\emptyset = 0$ and

$$c_T = v(T) - \sum_{\substack{S \subset T \\ S \neq T}} c_S$$

Then

$$\sum_{S \in \mathcal{P}(X)} c_S w_S(T) = \sum_{S \subset T} c_S = c_T + \sum_{\substack{S \subset T \\ S \neq T}} c_S = v(T)$$

Hence v is uniquely determined by the coefficients c_S and S -veto games w_S as $v = \sum_{S \in \mathcal{P}(X)} c_S w_S(T)$.

By the additivity axiom,

$$\Phi_i(X, v) = \Phi_i(X, \sum_S c_S w_S) = \sum_{S \in \mathcal{P}(X)} \Phi_i(X, c_S w_S) = \sum_{\substack{S \in \mathcal{P}(X) \\ i \in S}} \frac{c_S}{|S|}$$

So $\Phi_i(X, v)$, and thus $\Phi(X, v)$, are uniquely determined by the constants c_S . ■

5.4 Shapley's Theorem

Thus, we have arrived at a strong result-

Theorem 5.4.1. (*Shapley*) *There exists a unique value $\Phi(X, v)$ satisfying the Shapley axioms from definition 5.3.1.*

Proof. By theorem 5.3.3, if such a Φ exists, it is unique. But by theorem 5.3.2, there is a Φ satisfying the Shapley axioms, given by the random arrival formula

$$\Phi_i(X, v) = \frac{1}{n!} \sum_{\pi \in S_n} v(p_\pi^i \cup \{i\}) - v(p_\pi^i)$$

Hence a unique value function, the Shapley value, exists. Further, the coefficients c_S described in 5.3.3 provide an alternative means of calculating the value. ■

Example 5.4.2. Consider three companies A, B and C, which seek to invest in a combined project. The project requires five million pounds in funding to be successful; the three companies have investment budgets of 2, 3 and 4 million pounds respectively. What is the worth of each company in terms of Shapley value?

By considering the coalitions with sufficient funding, we observe that

$$v_{\{A\}} = v_{\{B\}} = v_{\{C\}} = v_\emptyset = 0$$

$$v_{\{A,B\}} = v_{\{A,C\}} = v_{\{B,C\}} = v_{\{A,B,C\}} = 1$$

Following the construction in theorem 5.3.3, we determine the values c_T . c_\emptyset is zero by assumption, and so for each $L \in \{A, B, C\}$,

$$c_{\{L\}} = v(\{L\}) - c_\emptyset = 0 - 0 = 0$$

i.e.,

$$c_{\{A\}} = c_{\{B\}} = c_{\{C\}} = 0$$

Thus

$$c_{\{A,B\}} = v(\{A, B\}) - (c_\emptyset + c_{\{A\}} + c_{\{B\}}) = 1 - (0 + 0 + 0) = 1$$

By symmetry,

$$c_{\{A,B\}} = c_{\{A,C\}} = c_{\{B,C\}} = 1$$

Finally, this gives

$$\begin{aligned} c_{\{A,B,C\}} &= v(\{A, B, C\}) - (c_{\emptyset} + c_{\{A\}} + c_{\{B\}} + c_{\{C\}} + c_{\{A,B\}} + c_{\{A,C\}} + c_{\{B,C\}}) \\ &= 1 - (0 + 0 + 0 + 0 + 1 + 1 + 1) \\ &= -2 \end{aligned}$$

So

$$\begin{aligned} \Phi_1(\{A, B, C\}, v) &= c_{\{A\}}/1 + c_{\{A,B\}}/2 + c_{\{A,C\}}/2 + c_{\{A,B,C\}}/3 \\ &= \frac{0}{1} + \frac{1}{2} + \frac{1}{2} + \frac{-2}{3} \\ &= \frac{1}{3} \end{aligned}$$

By the same argument, $\Phi_1(\{A, B, C\}, v) = \Phi_2(\{A, B, C\}, v) = \Phi_3(\{A, B, C\}, v) = 1/3$. That is, no one company has more influence (or right to the profits) than any of the others. This is intuitively obvious from the problem formulation, since it is the simple majority game of section 4.3: No-one company can afford the project, but *any* two can. This can also be seen from the random arrival formula, as the value of a coalition will only be increased when a second company joins, which, across the set of all permutations, is equally likely for any particular company.

6 Concluding remarks

Thus it can be seen that the techniques of Game Theory provide numerous insights into questions of decision making and strategy. The two more recent theorems described (Nash equilibria and Shapley values) are particularly far reaching, generalising to arbitrary player or strategy sets without modification; yet even the simplest of problems (such as the iterated prisoner's dilemma) can yield complex behaviour. That quantitative predictions are sometimes at odds with intuition or observed behaviour indicates that there is still much to say about the nature of rationality, utility and human behaviour. Attempting to reconcile those differences will strengthen both Game Theory and the social sciences to which it is applied, and as a relatively new mathematical field, Game Theory almost certainly still has many interesting results to offer.

7 Figures and sources




Figure 1: Expected payoff of the three responses

Background reading

- T. S. Ferguson- Lecture notes for *Mathematics 167, Game Theory, Fall 2000*, University of California Los Angeles (UCLA)
<http://www.math.ucla.edu/~tom/GameTheory.html> .
- Yuval Peres- Lecture notes for *STAT 155: Game Theory (Fall 2005)*, Berkeley University of California
<http://www.stat.berkeley.edu/users/peres/155.html> .
- E. Rasmusen- *Games and Information*, Blackwell Publishers.
- E. Winter- *The Shapley Value*
<http://www.ma.huji.ac.il/~Emseyal/shapley4to95.pdf>

References

- [1] R. Axelrod (1980) 'Effective choice in the prisoner's dilemma' *Journal of Conflict Resolution* Vol 24 No. 1 (March).
- [2] R. Frank, T. Gilovich and D. Regan (1993) 'Does studying economics inhibit cooperation?' *Journal of Economic Perspectives* Vol 7 Issue 2 (spring).
- [3] W. Grossman 'New tack wins prisoner's dilemma'
<http://www.wired.com/news/culture/0,1284,65317,00.html> .
- [4] S.P. Hargreaves Heap, Y. Varoufakis 'Game Theory- A critical text' (Second edition, 2004)
- [5] Von Neumann and Morgenstern- *Theory of Games and Economic Behaviour*, Princeton University Press.
- [6] 'Starbucks, monoculture, and American Imperialism'
http://www.everything2.com/index.pl?node_id=1535256 .