Contents

1 Introduction

Games can be categorized into two forms: Strategic Form Games (also called Normal Form Games), and Extensive Form Games. Today we will be focusing on strategic form games.

1.1 Why Do We Care?

- Economy
- Evolutionary Biology
- Large Scale Distributed Systems
- Resource Allocation
- Intelligent Agents

2 Strategic Form Games

1. Let $\mathcal{I} = \mathbb{N}_I$ be a finite set of players, where I is the number of players.
2. For $i \in \mathcal{I}$, let S_i be the (finite) set of pure strategies available to player i.
3. Define the set of pure strategy profiles to be the Cartesian product of the pure strategies sets of the players: $S = S_1 \times S_2 \times \cdots \times S_I$

2.1 Conventions

We write, $s_i \in S_i$ for a pure strategy of player i. We also write $s = (s_1, s_2, \ldots, s_I) \in S$ for a pure strategy profile. Let $\neg i$ denotes the player i's "opponents" (all players other than player i). Thus, we can write, $S_{\neg i} = \prod_{j \in \mathcal{I}; j \neq i} S_j$ Just as before, $s_{\neg i} \in S_{\neg i}$ denotes a pure strategy profile for the opponents of i. We will abuse notation a bit and write a pure strategy profile s as

$$(s_i, s_{\neg i}) \in S \tag{1}$$

When discussion pure strategies, we will write the player i pay-off function as $u_i : S \to R$, where $u_i(s)$ is the von Neumann-Morgenstern utility of player i given the pure strategy profile s.

Definition 2.1 (strategic form game) A strategic form game is is a tuple $(\mathcal{I}, \{S_1, S_2, \ldots, S_I\}, \{u_1, u_2, \ldots, u_I\})$ consisting of a a set of players, their pure strategy spaces, and their pay-off functions.

2.2 Domination and Nash Equilibrium

Definition 2.2 (Dominated Strategy I) A pure strategy s_i is strictly dominated if

$$\exists s_i' \in S_i \text{ s.t. } \forall s_{\neg i} \in S_{\neg i} \quad u_i(s_i', s_{\neg i}) > u_i(s_i, s_{\neg i})$$

A pure strategy s_i is weakly dominated if

$$\exists s_i' \in S_i \text{ s.t. } (\forall s_{\neg i} \in S_{\neg i} \quad u_i(s_i', s_{\neg i}) \geq u_i(s_i, s_{\neg i}) \wedge \exists s_{\neg i} \in S_{\neg i} \quad u_i(s_i', s_{\neg i}) > u_i(s_i, s_{\neg i}))$$

Definition 2.3 (Best Response) The set of best responses for player i to a pure strategy profile $s \in S$ is

$$\mathrm{BR}_i(s) = \{s_i^* \in S_i \mid \forall s_i \in S_i, u_i(s_i^*, s_{\neg i}) \geq u_i(s_i, s_{\neg i})\}$$

Let the joint best response set be $\mathrm{BR}(s) = \prod_i \mathrm{BR}_i(s)$.

Definition 2.4 (Nash Equilibrium)

A pure strategy profile s^* is a Nash equilibrium if for all players i,

$$\forall s_i \in S_i, u_i(s_i^*, s_{\neg i}^*) \geq u_i(s_i, s_{\neg i}^*) \tag{2}$$

Thus a Nash equilibrium is a strategy profile s^* such that $s^* \in \mathrm{BR}(s^*)$.

A Nash equilibrium s^* is strict if each player has a unique best response to his rivals' strategies: $\mathrm{BR}(s^*) = \{s^*\}$.

$$\forall s_i \neq s_i^*, u_i(s_i^*, s_{\neg i}^*) > u_i(s_i, s_{\neg i}^*) \tag{3}$$

2.3 Example

	L	M	R
U	(4,3)	(5,1)	(6,2)
M	(2,1)	(8,4)	(3,6)
D	(3,0)	(9,6)	(2,8)

For column-player, M is dominated by R. Column-player can eliminate M from his strategy space. The pay-off matrix reduces to

	L	R
U	(4,3)	(6,2)
M	(2,1)	(3,6)
D	(3,0)	(2,8)

For row-player, M and D are dominated by U. Row-player can eliminate M and D. The new pay-off matrix is

	L	R
U	(4,3)	(6,2)

Next, column-player eliminates R as it is dominated by U and reduces the pay-off matrix to

	L
U	(4,3)

Note that

$$\begin{aligned} \text{BR}_r(U,L) &= U \\ \text{BR}_c(U,L) &= L \\ \text{BR}(U,L) &= (U,L) \end{aligned} \tag{4}$$

(U; L) is a strict Nash equilibrium.

Remark: If we allow mixed strategies, this is *not* a Nash equilibrium.

3 Two-Player Non-Cooperative Games

3.1 Social Dilemmas

Social dilemmas are social interactions where everyone enjoys the benefits of collective action, but any individual would gain even more without contributing to the common good. (Here we assume others do not imitate her selfish action.)

Key problem: how can we promote cooperation in these situations? Is imposition of a central authority the only solution?

Dawes (1980) observes that the fundamental tensions that generate social dilemmas for humans are present in the three crucial problems of the modern world: resource depletion, pollution, and overpopulation. (Social dilemmas are not exclusive to human interactions, of course.)

Simple social dilemma: two-person game, where each player can either cooperate or defect. Payoffs are as in Table 2.

The mnemonic is: R = reward for cooperation, P = punishment for selfishness, S = sucker, T = temptation (or treachery).

The essence of a social dilemma arises when both players prefer any outcome in which the opponent cooperates to any outcome in which the opponent defects ($\min(Ti, Ri) > \max(Pi, Si)$) but each has a reason to defect.

Table 1: Two Person Games

i / j	C	D
C	(R_i,R_j)	(S_i,T_j)
D	(T_i,S_j)	(P_i,P_j)

Examples of such reasons: the temptation to cheat (if $Ti > Ri$) or the fear of being cheated (if $Si < Pi$). This puts cooperation at risk.

Three well-known social dilemma games:

- Chicken: the problem is greed (temptation) but not fear ($Ti > Ri > Si > Pi$; $i = 1, 2$)
- Stag Hunt: the problem is fear (of being cheated) but not greed ($Ri > Ti \geq Pi > Si; i = 1, 2$)
- Prisoner's Dilemma: both problems coincide in the paradigmatic Prisoner's Dilemma ($Ti > Ri > Pi > Si; i = 1, 2$).

It is common to consider the symmetric versions of these three social dilemma games (e.g., Macy and Flache 2002). Payoffs are as in Table 2.

Table 2: Symmetric Two Person Games (social dilemma)

i / j	C	D
C	(R,R)	(S,T)
D	(T,S)	(P,P)

3.2 Assurance Game (With Two Players)

Voilà comment les hommes purent insensiblement acquérir quelque idée grossière des engagements mutuels, et de l'avantage de les remplir, mais seulement autant que pouvait l'exiger l'intérêt présent et sensible; car la prévoyance n'était rien pour eux, et loin de s'occuper d'un avenir éloigné, ils ne songeaient pas même au lendemain. S'agissait-il de prendre un cerf,

chacun sentait bien qu'il devait pour cela garder fidèlement son poste; mais si un lièvre venait à passer à la portée de l'un d'eux, il ne faut pas douter qu'il ne le poursuivît sans scrupule, et qu'ayant atteint sa proie il ne se souciât fort peu de faire manquer la leur à ses compagnons.
–Rousseau 1754 (Discours sur l'origine et les fondements de l'inégalité parmi les hommes.)

Names for the game: "stag hunt", "assurance game", "coordination game", and "trust dilemma".

Table 3: Stag Hunt Problem

	Stag	Hare
Stag	(2,2)	(0,1)
Hare	(1,0)	(1,1)

Two hunters can catch a stag if they cooperate. On his own, a hunter catches nothing. Each hunter has the alternative of defecting to catch a hare.

- If both row-player and column-player hunt stag, since a stag is worth 4 "utils", they each get 2 "utils."
- If both row-player and column-player hunt hares, since a hare is worth 1 "util", they each get 1 "util."
- If row-player hunts hare, while column-player hunts stag (and hence fails to hunt any thing), then the row-player gets 1 "util" and the column-player gets 0 "util."
- The other case is symmetric.

The strategy "hunt stag" has a minimum utility of 0 if the your opponent hunts hare. The strategy "hunt hare" has a minimum utility of 1 regardless of what the opponent chooses to play. So if you are "uncertainty averse" (maximin), you will choose to hunt hare and thus guarantee that you gets 1 independent of the choice of your opponent. This will maximize the minimum utility under the two possible pure strategies.

Conclusion: two uncertainty averse players will each choose to hunt hares.

3.3 Chicken Game

Rapoport and Chammah (1966) discuss this in detail.

Standard version:

Two drivers head for a single lane bridge from opposite directions. Each can swerve to yield the bridge to the other. If neither swerves, they collide on the bridge, which is the worst outcome for each. The best outcome for each driver is to stay straight while the other swerves. (Then the other is the "chicken" but a crash is avoided).

Each player, in attempting to secure his best outcome, may decide to risk the worst.

Famous variation ("chickie run"): in the movie Rebel Without a Cause, two characters race their cars towards a cliff instead of each other.

Two versions of the payoff matrix for this game: words and numerical payoffs.

Words: each player would prefer to win over tying, prefer to tie over losing, and prefer to lose over crashing. Here C is “swerve” and D is “straight”.

Table 4: Chicken (in words)

ij	C	D
C	(Tie, Tie)	(Lose, Win)
D	(Win, Lose)	(Crash, Crash)

Numerical payoffs: the benefit of winning is 1, the cost of losing is -1, and the cost of crashing is -10.

Table 5: Chicken (with payoffs)

ij	C	D
C	(0,0)	(-1,1)
D	(1,-1)	(-5,-5)

“Chicken” is an example of an “anti-coordination games”: it is mutually beneficial for the players to play different strategies.

Crucial underlying concept: is that players use a shared resource.

(In coordination games, sharing the resource creates a benefit for all: the resource is non-rivalrous and shared usage creates a positive externality. In anti-coordination games the resource is rivalrous but non-excludable: shared usage creates a negative externality.

Because the "loss" of swerving is so trivial compared to the crash that occurs if nobody swerves, the reasonable strategy would seem to be to swerve before a crash is likely. Yet, knowing this, if one believes one's opponent to be reasonable, one may well decide not to swerve at all, in the belief that he will be reasonable and decide to swerve, leaving the other player the winner.

This can be formalized by saying there is more than one Nash equilibrium. The pure strategy equilibria are the two situations wherein one player swerves while the other does not.

Bertrand Russell famously compared the game of Chicken to nuclear brinkmanship:

> Since the nuclear stalemate became apparent, the Governments of East and West have adopted the policy which Mr. Dulles calls 'brinkmanship'. This is a policy adapted from a sport which, I am told, is practised by some youthful degenerates. This sport is called 'Chicken!'. It is played by choosing a long straight road with a white line down the middle and starting two very fast cars towards each other from opposite ends. Each car is expected to keep the wheels of one side on the white line. As they approach each other, mutual destruction becomes more and more imminent. If one of them swerves from the white line before the other, the other, as he passes, shouts 'Chicken!', and the one who has swerved becomes an object of contempt. As played by irresponsible boys, this game is considered decadent and immoral, though only the lives of the players are risked. But when the game is played by eminent statesmen, who risk not only their own lives but those of many hundreds of millions of human beings, it is thought on both sides that the statesmen on one side are displaying a high degree of wisdom and courage, and only the statesmen on the other side are reprehensible. This, of course, is absurd. Both are to blame for playing such an incredibly dangerous game. The game may be played without misfortune a few times, but sooner or later it will come to be felt that loss of face is more dreadful than nuclear annihilation. The moment will come when neither side can face the derisive cry of 'Chicken!' from the other side. When that moment is come, the statesmen of both sides will plunge the world into destruction.
>
> – B. Russell, 1959

3.4 Prisoners' Dilemma

There are two prisoners (row-player and column-player) arrested for a particular crime. The prosecutor has enough evidence to convict them each of a lesser crime and send each to jail for 2 years. He relies one one of them testifying against the other in order to get a full conviction and punish the other prisoner by sending him to for 5 years. If both of them testify against the other (defections: "D, D") then they

Table 6: Prisoners' Dilemma

	C	D
C	(-2,-2)	(-5,0)
D	(0,-5)	(-4,-4)

$\Longrightarrow$

	C	D
C	(3,3)	(0,5)
D	(5,0)	(1,1)

both go to jail for 4 years each, thus getting a payoff of -4. If, on the other hand, both maintain silence (cooperations: "C, C") then they are convicted only of the lesser crime, with payoff of -2 each. If one player testifies (defects, D) and column-player maintains silence (cooperates, C), then defector is rewarded with no jail time and column-player is punished with 5 years in jail.

The pay-offs can be made all non-negative by adding 5 to each payoff and thus getting an apparently equivalent pay-off matrix.

Defection is a dominant strategy:

- Suppose you are the row-player. Strategy C is dominated by the strategy D independent of how column-player plays. So you should defect.
- Similarly, for your opponent (column-player), the strategy C is dominated by the strategy D independent of how you play. Thus row player must defect.

Hence the equilibrium strategy for the players is to defect, even when they could have each gotten better pay-offs by cooperating.