Chapter 1

Introduction

1.1 The Game of Life

In 1970, the mathematician John Conway invented a game called *Life*. This game of Life is played on a grid of square cells-like a chess board but **extending infinitely** in every direction. A cell can either be live or dead. Let us adopt the following convention: a live cell is shown by putting a red marker on its square. A dead cell is shown by leaving the square empty (green square in the pictures below). Each cell in the grid has a neighborhood consisting of the eight cells in every direction including diagonals:

1	2	3
4	□	5
6	7	8

1.1.1 The Rules of the Game

To apply one step of the rules, we count the number of live neighbors for each cell. What happens next depends on this number.

(1) A dead cell with exactly three live neighbors becomes a live cell (birth).

(2) A live cell with two or three live neighbors stays alive (survival).

(3) In all other cases, a cell remains dead or it dies (overcrowding or loneliness).

Be careful that the number of live neighbors is always based on the cells before the rule was applied. This means we must first find all of the cells that change before changing any of them. Conway chose the rules carefully after trying many other possibilities, some of which caused the cells to die too fast and others which caused too many cells to be born. His Game of Life balances these tendencies, making it hard to tell whether a pattern will die out completely, form a stable population, or grow forever. In aother words, in most cases, it is impossible to look at a starting configuration and see what will happen in the future. The only way to find out is to follow the rules of the game, step by step (a job for a computer!).

This game of life has attracted much interest precisely because of the surprising ways the patterns can evolve. From a theoretical point of view it is interesting because it has the power of a universal Turing machine (it is Turing complete), which means that anything that can be computed by a Turing machine can also be computed within Conway's Game of Life. The reason for this is that it is possible to build a pattern which acts like a finite state machine connected to two counters (something known to be equivalent to a universal Turing machine).
Besides, the following problem related to Conway's Game of Life is P-complete: given an initial configuration, a particular cell, and a time t (in unary), is that particular cell be alive after t steps?

However, this Game of Life is not a game in the conventional sense. It looks like a game because it is defined in terms of *rules of the game*. But

(1) there are no players,

(2) there are no winning nor losing; no outcomes, no payoffs. Once the "pieces" are placed in the starting position, the rules determine everything that happens later.

1.2 The Game of Nim

In the movie *Last Year at Marienbad* (1961, dir. A. Resnais) a character called M repeatedly persuades another called X, to play a simple game. He arranges cards on a table in a triangular shape:

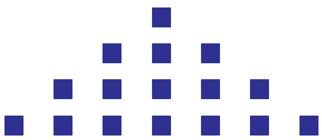

Players take turn, picking as many cards as they like from one row. This goes on, until one of the two players picks the last card and loses.

In the movie, M politely lets X make the first move. Unfortunately, X always loses.

This game is a *Nim Game*, one of the oldest and most engaging of all two-person mathematical games known today. The word *Nim* traces back to the german word *nehmen*. The name and the complete theory of the game were invented a century ago by C. L. Bouton of Harvard University. In the general version, players take turns removing objects (counters, pebbles, coins, pieces of paper) from heaps (piles, rows, boxes), but only from one heap at a time. In the normal convention, the player who removes the last object wins, in the misère convention the player to move last loses.

Nim is really a game since it verifies

Definition 1 *There are three ingredients to a game:*

(a) ***Players:*** *interactive parties who participate in the game. Usually there are two of them, called* Player *and* Opponent, Adam *and* Eve, Abelard *and* Eloïse, He *and* She, 0 *and* 1, *I and II. Or specific names describing their actions in the game (*Verifier *and* Falsifier*).*

(b) ***Rules of the game****, which specify* who *can* when *do* what*. Also called the* game form*.*

(c) ***Winning conditions*** *or even* outcomes *according to what was done in the course of the game.*

The rules of this game can be displayed by describing the *arena* where it takes place. In most cases, this can be viewed as a *graph* where vertices represent the possible configurations of the game, and the edges, the moves that are allowed (also called *legal* moves). For example, below is a graph-representation of a Marienbad game played with only the first two rows:

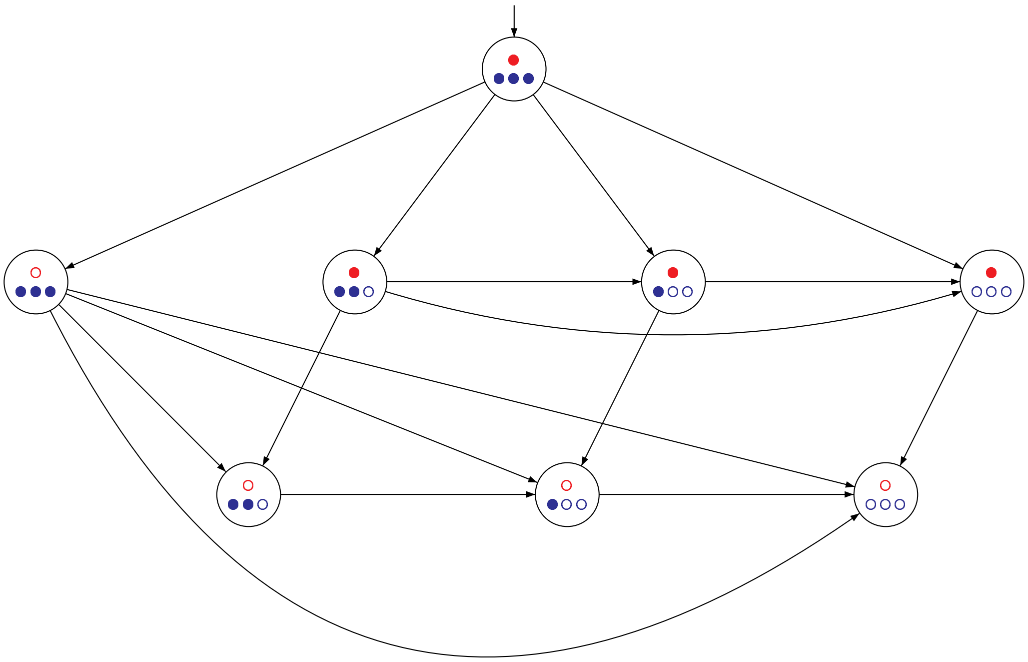

Graph for 2-Marienbad Game

Usually, the representation of the rules of the game in terms of graph, each node is related to a specific player (the one who must play at this vertice). In other words, there is a partition of the set of vertices in the number of players. However, the example above do not mention players'alternations, since players simply take turn.

But this presentation of the game by mean of the description of the arena, do not take into account the *winning conditions.* For this reason, the most detailed representation of the overall rules of the game is the one known as the *extensive form.* It constitutes a way of verifying that the rules of a game are complete and therefore that the game is well-defined. In other words, a game has complete rules if and only if it can be represented in *extensive form*.

This *extensive form* takes the form of a *tree* that amounts to a *flow chart* representation of an outside observer's view of the rules of the game. It is the *tree* of all possible sequences of *legal* moves, i.e., the tree of all authorized sequences of configurations. We call such a legal sequence of configuration a *position*. We recall that a *tree* $\mathcal{T}$ is a set of sequences that is also closed under subsequence ($u \subseteq v \ \wedge \ v \in \mathcal{T} \Longrightarrow u \in \mathcal{T}$).

We write $\mathbb{G}_{\mathcal{T}}$ to denote the game whose extensive form is given by $\mathcal{T}$. In case of the 2-Marienbad Game, this tree takes the form:

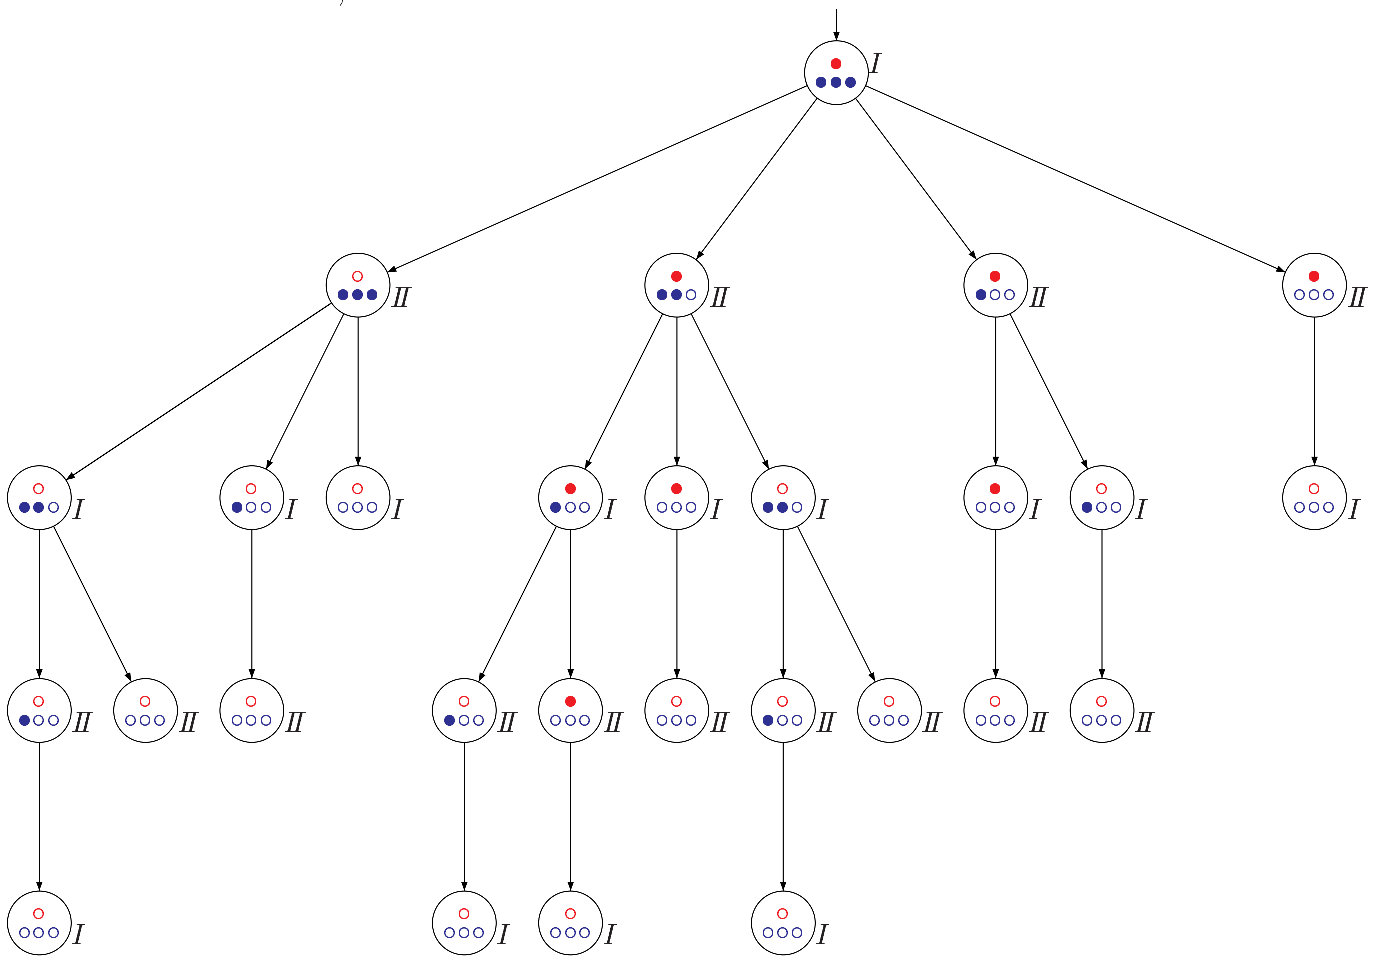

Tree for 2-Marienbad Game

Here also, the winning conditions should be mentionned: reaching a terminal node is losing. If one considers the tree of all possible successive removals of cards in any ways (e.g. in one move removing one card from the first row and also one from the second,etc), the tree for 2-Marienbad Game appears as a subtree of this general tree of all configurations; a subtree that is built by pruning the branches that precisely break the law. In other words, it is the overall tree restricted to legal moves.

An element of the game-tree (i.e. a sub-branch $\langle c_0, \ldots, c_k \rangle \in \mathcal{T}$) is called a *position*.

The innovation that game theory introduces is the *strategic dimension*. One certainly should not think of a player as a timberman who, when chopping wood, does not expect the wood to fight back, but rather as a general. His (her) environment is not neutral. The general has to overcome and anticipate resistance from the enemy. He must expect that the enemy will try to outguess his plans.

A strategy (for a given player P), is a function that tells him (her) what move to make given the previous moves. It is a specification of how to play the game in every contingency It is important to notice that no other players'moves are restricted, only the ones from player P are. Another way to look at a strategy is to consider it from the extensive form of the game.

Definition 2 *A strategy σ for player P is a subtree of the game tree $\mathcal{T}$ that satisfies:*

(a) $\langle c_0 \rangle$ *= the initial position belongs to* σ*,*

(b) if $\langle c_0, \ldots, c_k \rangle \in \sigma$ *is a position for P (meaning it is a position where P has to move next), then there is a unique position* $\langle c_0, \ldots, c_k, c_{k+1} \rangle \in \mathcal{T}$ *that belongs to* σ*.*

(c) if $\langle c_0, \ldots, c_k \rangle \in \sigma$ *is a position for some other player than player P, then* **all** *positions* $\langle c_0, \ldots, c_k, c_{k+1} \rangle \in \mathcal{T}$ *must belong to* σ*.*

In case of the 2-Marienbad Game, a strategy looks like:

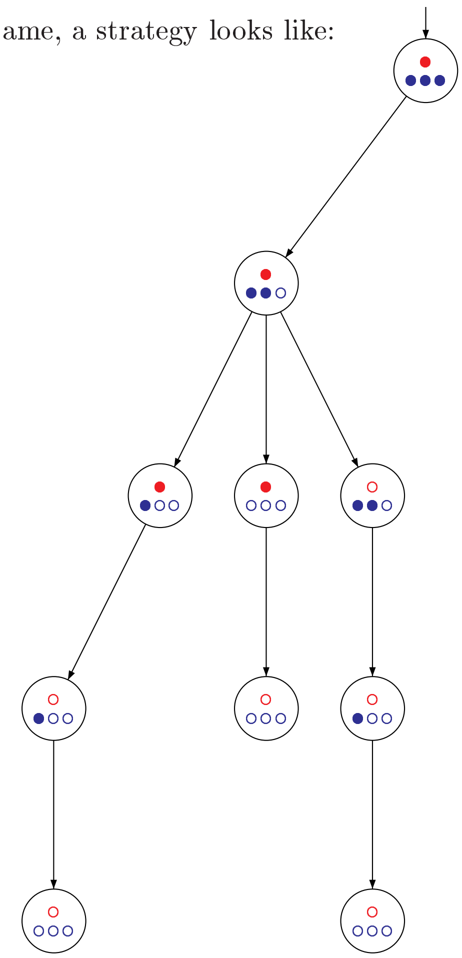

A Strategy for Player I in the 2-Marienbad Game

With this tree-like point of view, a strategy really stands as a sub-game. Since $\sigma \subseteq \mathcal{T}$, one can consider σ as the extensive form of a game derived from the one presented by $\mathcal{T}$: it has the particularity of being a sub-game that leaves no choice to player P's moves.

Remark that if in a game $\mathbb{G}_{\mathcal{T}}$ each player applies a strategy, then this determines a unique *run*; that is a unique branch of $\mathcal{T}$..

A strategy σ (for player P) in a game $\mathbb{G}_{\mathcal{T}}$ is *winning* if player P always wins when applying σ. This means that all branches of σ lead to a winning condition for P.

It is very natural, given any game, to try to find a winning strategy for a player, a one that would defeat the opponent for ever. Games where such a strategy exist are called *determined*.

If we go back to the 2-Marienbad Game, we see that player *I* has a winning strategy, namely *remove all cards from the second row*. More generally, the *Nim Game* is determined, one of the two players always has a winning strategy. It is explained in Bouton's theory based on the binary representation of the cards in each row, and digit counting.

1.3 A Winning Strategy for The Game of Nim

We consider the version of this game where players take matches from heaps, players (*I* and *II*) take turns, player *I* starts and the player who takes the last match wins the game.

By convention, we call a *losing configuration* a configuration such that if one *Player* (*I* or *II*) can secure this configuration by his move, leaving his opponent to move next, then, whatever his opponent may do, the first *Player* can play so as to win the game. Any other configuration is called *winning*. For example the configuration $(2, 2)$ is a losing configuration since if the opponent takes two matches, then the *Player* takes the other two. If the opponent takes one, the *Player* takes one from the other heap. In either case the *Player* wins..

We define a *correct configuration*:

(a) we first express the number of matches in each column in binary. (E.g. the configuration $(3, 9)$ gives $(11, 1001)$).

(b) We then write down these numbers one under the other,

(c) and finaly add up the columns.

Thus the configuration $(2, 3, 6, 7)$ gives rise to the figure

$$\begin{array}{ccc} & 1 & 0 \\ & 1 & 1 \\ 1 & 1 & 0 \\ 1 & 1 & 1 \\ \hline \\ 2 & 4 & 2 \end{array}$$

The configuration is *correct* if the sum of each column is even, in any other cases it is *incorrect*.

Theorem 3 *A configuration in* Nim *is losing if and only if it is correct.*

Proof of theorem 3:

(a) Consider the special case in which no heap contains more than one match. It is straightforward to see that such a configuration is losing if and only if it is correct.

(b) If a player must take from a correct configuration, then he must replace one number defining a row of the figure by a smaller number. Hence he changes at least one of its digits. This shows that when a player takes from a correct configuration, he transforms it into an incorrect one.

(c) If a position is incorrect, then consider the far most left column whose sum is odd. to fix ideas, imagine that the figure is $(even, even, odd, even, odd, even)$, so that it is the third column we're interested in. There is necessarily a 1 in that column, and assume also to fix idea that one row in which it happens is: 011101. Then by replacing this number by 010111 (which is a smaller number!), one restores the correctness of the configuration. Hence, a solution is to remove

$$\begin{array}{rcccccccc} & 0 & 1 & 1 & 1 & 0 & 1 & & \\ - & 0 & 1 & 0 & 1 & 1 & 1 & & \\ \hline \\ & 0 & 0 & 0 & 1 & 1 & 0 & = & 6 \end{array}$$

matches from that heap.

This shows that a player, presented with an incorrect configuration can always convert it into a correct one.

(d) Finally, if a player leaves a correct configuration, the opponent is compelled to convert it into an incorrect one, then the first player can move so as to restore a correct configuration. This process continues until every heap is exhausted or contains one match only. Thus the theorem is reduced to case (2)(a) which is already proved.

$\dashv$ 3

This theorem shows

- if the original configuration is incorrect, then the first player has a winning strategy;
- if the original configuration is correct, then the player who plays second has a winning strategy.

Which

(a) proves that this game is *determined*,

(b) points out the *player* who has a w.s.,

(c) provides a winning strategy.

1.4 Games Economic Experts play

To address problems arising in the context of the Internet, algorithmic thinking is actually combined with game-theoretical and economic concepts.

Games theory originates, to a large extent, from mathematics and economics. It is devided in two branches: *cooperative* and *non-cooperative* game theory..

- The *cooperative* approach to games deals with games in which
  (a) the rules are very broadly defined. The rules of the game are kept implicit.
  (b) Hence, the individual decision problems of players cannot be directly analyzed. Therefore, cooperative theory emphasizes on coalitions of players.
  (c) It implicitly assumes that such coalitions can be formed and maintained. So, commitments are available.
- On the contrary, in *non-cooperative* game theory:
  (a) the rules are complete,
  (b) the ultimate decision units are the (individual) players,
  (c) Commitments are not available, unless allowed for by the rules of the game.

This course definitely focus on *non-cooperative* game theory.

1.5 Real Games

They involve preferences among outcomes, not just winning or losing. Decisions of agents involve expected utility. The aim is to predict rational behavior for agents in various situations, and more generally to explain stable behavior of groups of individuals.
The emphasis is on *strategic games*, presented in *normal form*: that is a description of a game by specifying the strategies and payoffs. As opposed to *extensive form game*: a description of a game by specifying the order of moves and information available to players, as well as the payoffs.
One does not analyze precise successive moves and possible runs, but rather players'over-all options.

1.5.1 Information and Chance

There are two essential differences between games such as chess and games such as poker:

(1) in chess, the current configuration of the game is fully accessible to both players. The only uncertainty is about future moves. While in poker, the players have *imperfect information.* They have only partial knowledge about the current configuration of the game. Chess is a *perfect information* game, poker an *imperfect information* one.

(2) In chess there is no chance move.

The lack of information involves uncertainty about the *current configuration* of the game (particularly situations where different players have access to different information. The possibility of chance moves involves only uncertainty about the future, uncertainty the is resolved as soon as the future materializes.
Both perfect and imperfect information games may involve an element of chance:

	perfect information	imperfect information
no chance	*Chess*	*Battleships*
chance	*Monopoly*	*Poker*

In economics, game theory has focused on decision under uncertainty, hence on games with imperfect information, since it is intended to deal with games derived from *real life* where one rarely has perfect information. Common games with imperfect information are games with simultaneous moves like rock-scissors-paper:

1.5.2 Dominant Strategies

The Prisoners' Dilemma

Imagine two burglars, Jack and John, are captured near the scene of a burglary. Each has to choose whether or not to confess and implicate the other.

- If neither man confesses, then both will serve one year on a charge of carrying a concealed weapon.
- If each confesses and implicates the other, both will go to prison for ten years.
- However, if one burglar confesses and implicates the other, and the other burglar does not confess, the one who has collaborated with the police will go free, while the other burglar will go to prison for twenty years on the maximum charge.

The strategies in this case are: confess or don't confess. The payoffs (penalties, actually) are the sentences served. We can express all this compactly in a "payoff table" of a kind that has become pretty standard in game theory. Here is the payoff table for the Prisoners' Dilemma game:

		John	
		$conf$	$don't$
Jack	$conf$	$(-10,-10)$	$(0,-20)$
	$don't$	$(-20,0)$	$(-1,-1)$

The table is read like this: Each prisoner chooses one of the two strategies. In effect, John chooses a column and Jack chooses a row. The two numbers in each cell tell the outcomes for the two prisoners when the corresponding pair of strategies is chosen. The number to the left of the comma tells the payoff to the person who chooses the rows (Jack) while the number to the right of the column tells the payoff to the person who chooses the columns (John). Thus (reading down the first column) if they both confess, each gets ten years, but if John confesses and Jack does not, Jack gets twenty and John goes free.
So: how to solve this game? What strategies are rational considering that both men have one goal: minimizing the time they spend in jail.

(1) Jack might reason as follows: John can confess or John can keep quiet

  (a) If John confesses.

    (1) Then I get twenty years, if I don't confess,

    (2) and I get ten years if I do so.

  (b) If John does not confess.

    (1) I only get one year if I don't confess,

    (2) but I can go free if I confess.

  So that either way, it is best if I confess. Therefore I'll confess.

(2) But John will certainly reason in the same way.

So that both Jack and John will confess and go to jail for ten years each.

Replace prisonners by countries (here A and B) and to confess by to arm, then you just get another familiar version of the prisonners' dilemma:

		B	
		arm	$disarm$
A	arm	$(1,1)$	$(3,0)$
	$disarm$	$(0,3)$	$(2,2)$

What has happened here is that the two prisoners have fallen into something called a *dominant strategy equilibrium.*

Definition 4

Dominant Strategy: *Let an individual player in a game evaluate separately each of the strategy combinations he may face, and, for each combination, choose from his own strategies the one that gives the best payoff. If the same strategy is chosen for each of the different combinations of strategies the player might face, that strategy is called a "dominant strategy" for that player in that game.*

Dominant Strategy Equilibrium: *If each player has a dominant strategy and plays it, then that combination of (dominant) strategies and the corresponding payoffs are said to constitute the dominant strategy equilibrium for that game.*

In the Prisoners' Dilemma game, to confess is a dominant strategy, and when both prisoners confess, that is a dominant strategy equilibrium.

1.5.3 Nash Equilibrium

One of game theory's basic premises is that a rational player, in order to guess opponents' actions, puts himself in the opponents' shoes, and simultaneously with his own he also solves the opponents' optimization problems. In a fully rational world this assumption leads to a Nash equilibrium.

The Battle of the sexes

A husband and his wife must agree on how to spend the evening. The husband prefers to go to the ballgame and the wife to the opera. However, they prefer agreement to disagreement:

		She	
		$ballgame$	$opera$
He	$ballgame$	$(2,1)$	$(0,0)$
	$opera$	$(0,0)$	$(1,2)$

This game has *no dominant strategy equilibrium.* If the husband thinks the wife's strategy is to choose the opera, his best response is to choose opera rather than ballgame. Conversely, if he thinks the wife's strategy is to choose the ballgame, his best response is ballgame. While in the prisoner's dilemma, the best response does not depend on what the other player is thought to be doing, in the battle of the sexes, the best response depends entirely on what the other player is thought to be doing.

For such games without dominant strategies another notion of equilibrium named after Nobel Memorial Laureate John Nash was introduced:

Definition 5 *Two strategies are in a* **Nash Equilibrium** *if each participant chooses the best strategy, given the strategy chosen by the other participant.*

In other words, if σ and τ are strategies for respectively player I and $I\!I$, the couple (σ, τ) forms a Nash Equilibrium if the following holds:

- if τ is fixed, then σ is the best strategy for player I *against* τ
- if σ is fixed, then τ is the best strategy for player $I\!I$ *against* σ.

The battle of the sexes game has two Nash equilibria: both go to the opera, or both go to the ball game: if each expects the other to go to the opera (ballgame) the best response is to go to the opera (ballgame).
A Nash equilibrium does not attempt to explain why players play the way they do. It merely proposes a way of playing so that no player would have a motivation to play differently.

Matching Pennies

In Conan Doyle's Sherlock Holmes story *The Last Problem*, Moriarity is pursuing Holmes by train in order to kill Holmes and save himself. The train stops at Canterbury on the way to Paris. If both stop at Canterbury, Moriarity catches Holmes and wins the game. Similarly if both stop at Paris. Conversely, if they stop at different places, Holmes escapes.

		Moriarty	
		Canterbury	*Paris*
Sherlock Holmes	*Canterbury*	$(1, -1)$	$(-1, 1)$
	Paris	$(-1, 1)$	$(1, -1)$

This is an example of a *zero sum* game: one player's loss is another player's gain. In the story, Holmes stops at Canterbury, while Moriarity continues on to Paris. It is a matching *pennies game*:

		$I\!I$	
		head	*tails*
I	*head*	$(1, -1)$	$(-1, 1)$
	tails	$(-1, 1)$	$(1, -1)$

But it is easy to see that this is not a Nash equilibrium: Moriarity should have anticipated that Holmes would get off at Canterbury, and so his best response was to get off also at Canterbury. This game has *no Nash equilibrium*: whichever player loses should anticipate losing, and so choose different strategy.

However, instead of simply choosing Canterbury or Paris, a player can flip a fair mental coin to decide what to do. This is an example of a *random strategy* (also called *mixed strategy*), which simply means a particular way of choosing randomly among the different strategies.

Theorem 6 (Von Neuman-Morgenstern Minimax Theorem) *Every game with finitely many players and strategies admits a mixed strategy Nash equilibrium.*

In our matching pennies game, each player should randomize 50/50 between Canterbury and Paris. If Moriarity randomizes 50/50 between the two alternatives, then Holmes has a 50% chance of winning and 50% chance of losing regardless of whether he chooses to stop at Canterbury or Paris. Since he is indifferent between the two choices, he does not mind flipping a coin to decide between the two, and so there is no better choice than for him to randomize 50/50 himself. Similarly when Holmes is randomizing 50/50, there is no better choice for Moriarity to do the same. Each player, correctly anticipating that his opponent will randomize 50/50 can do no better than to do the same.
Mixed strategy Nash equilibrium does not require a positive reason for playing the equilibrium strategy. In matching pennies, Holmes and Moriarity are indifferent: they have no positive reason to randomize 50/50 rather than doing something else. However, it is only an equilibrium if they both happen to randomize 50/50.

However, the problem that arises with Nash Equilibria is that very often there are too many of them...