Introduction to Solution Concepts in Game Theory

Uriel Feige

April, 9, 2024

1 Introduction

1.1 What is a game?

A game has *players*. We shall assume here that the number of players in a game is at least two and finite. (For simplicity, we shall assume players of male gender, though the gender of the players is irrelevant, and the players may also be genderless software programs.)

Every player i has a set S_i of *strategies* available to him. We shall assume here that the sets S_i are finite, though game theory also addresses games with infinite strategy sets. In a game, every player selects one strategy from his respective set of strategies.

The *outcome* of the game is the profile of strategies selected by the players, one strategy for each player. Hence it is a vector of strategies.

Every player is assumed to have his own *preference relation* over outcomes. This preference relation induces a partial order over outcomes. For simplicity, let us assume here that the partial order is given in form a numerical value for each outcome, and the player prefers those outcomes with higher numerical value over those with lower numerical value. These values are often referred to as *payoffs* for the players. It is often the case that the value of the payoff is meant to have quantitative significance beyond the preference of order over outcomes. For example, a payoff 2 would be considered twice as good as a payoff of 1. In this case, the payoff function would typically be called a *utility* function for the player.

A game is represented in *normal* form (a.k.a. strategic form, matrix form, or bi-matrix for two players) if it explicitly lists for each player and every outcome the value of the outcome for that player. For two player games, a game in normal form is often depicted as a pair of payoff matrices, one for the row player and the other for the column player, or even as a single matrix with both payoffs written in each entry.

Often, normal form representation of games is prohibitively large. There are many other more succinct representations of games. Many games (such as the game of chess) are represented implicitly. Two common forms of succinct representations for general classes of games are *extensive form* (typically, a game tree for two person games), and *graphical games* (a graphical representation of multiplayer games in which the payoff of a player is affected only by actions of players in his immediate neighborhood).

1.2 Solution concepts

A description of a game does not say how the players actually choose their strategies. Game theory tries to answer this question. There are two different aspects to this question.

- *Descriptive.* Given a game, how will people actually play it? What strategies will they choose? These questions involve considerations from social science and psychology, and their study may well require experimentation.
- *Prescriptive.* Given a game, how should players play it? Recommend strategies for the players. This is the more theoretical part of game theory, and here mathematics and computer science have a more prominent role.

Though the descriptive and prescriptive aspects are related (e.g., if a strategy is recommended, will the players actually play it?), the emphasize here will be on the prescriptive aspect of game theory.

We will try to associate *solution concepts* with a game. The solution concept will offer to players a recommendation of which strategy to play. The recommended strategy will satisfy some optimality conditions (that depends on the solution concept).

Before presenting the solution concepts that will be discussed here, let us explicitly list some of the assumptions that we shall make here.

1. *The game is given.* A game is sometimes an abstraction of a real life situation. Reaching the right abstraction (who the players are, what actions are available to them, what are their preferences) might be a very difficult task and will not be addressed here. We assume that the game is given.
2. *Self awareness.* The player is aware of those aspects of the game that are under his control. He knows which strategies are available to him, and his own preference relation over outcomes. (Indirectly, this requires awareness of strategies available to the other players, as they define the possible outcomes.)
3. *Full information.* For games in normal form, this is essentially the same as self awareness. However, for games given in some succinct representation, full information is different from self awareness. For example, for games in extensive form, in means that after every move the player knows exactly at what node of the game tree the game is. Chess is an example of a game of full information. Poker is an example of a game of partial information.

4. *No computational restrictions.* When suggesting a solution concept, we shall ignore the question of whether finding a recommended strategy under this concept is computationally tractable. Of course, in many games (chess being one example) the issue of computational limitations (a.k.a. *bounded rationality*) is an important aspect of the game. We remark that even though the solution concepts themselves will not a-priori be required to be computationally efficient, we shall eventually be interested in their computational complexity.

We now present some of the solution concepts that will be discussed more thoroughly later. The underlying assumption is that players are *rational*. Here the word rational is meant to convey that a player attempts to maximize his own payoff. It is worth pointing out that a player may appear to behave irrationally from this respect. This can often be attributed to the failure of one or more of the assumptions listed above.

1.3 Solutions in pure strategies

Dominant strategies. A strategy s is *dominant* for player i if regardless of the strategies played by other players, the payoff of player i is strictly maximized by playing s_i.

Formally, for every set of strategies s_{-i} for all players but i and every strategy $s' \neq s$,

$$u_i(s, s_{-i}) \geq u_i(s', s_{-i})$$

Dominant strategies do not always exist (e.g., for a payoff matrix that is the identity matrix). However, when they do exist, they are a very favorable solution concept. For games in normal form, they can be computed efficiently.

An interesting example is the game *prisoner's dilemma*, which has the following game matrix.

	Cooperate	Defect
Cooperate	-2; -2	-9; -1
Defect	-1; -9	-8; -8

Both players have dominant strategies, but the outcome of playing them is inferior for both players than the outcome if alternative strategies are played. The source of the problem is that a player can slightly improve his own payoff at the cost of other players loosing a lot. If each player selfishly maximizes his payoff, everyone looses.

Examples such as prisoner's dilemma point to a shortcoming of the concept of rational behavior - it might lead to undesirable outcomes.

It is worth mentioning in this context an area of game theory that will be addressed in other parts of the course, that of *mechanism design*. At a high level, the purpose of this area is to set up games in such a way that rational behavior will always lead to desirable outcomes, avoiding situations such as the prisoner's dilemma.

A well known example for mechanism design is that of a Vickery auction. Consider an auctioneer who has one item for sale, and k bidders (players), where bidder i has value u_i for the item (known only to bidder i). The process of the auction is a sealed bid auction, in which first all players submit their bids in sealed envelopes, and then the envelopes are opened and the highest bidder wins. We assume that the *payoff* of an agent is 0 if he does not win, and his value for the item minus his payment if he wins. In a *first price auction*, the winner pays his bid. In general, there are no dominant strategies in first price auctions, and the bids of the bidders depend on their beliefs regarding what other bidders will bid. In a *second price (Vickery) auction*, the winner pays the second highest bid. This has the desirable outcome that players have dominant strategies - to bid their true value. Moreover, in this case the outcome is that the item is allocated to the player who desires it most, which promotes economic efficiency. Hence the second price auction is an example of a game that is designed in such a way that its solution under a standard solution concept optimizes some economic goal: maximizing total economic welfare. (Note that here money paid by the bidder to the seller is assumed to have no effect on the total economic welfare, because the total sum of money remains constant.)

Nash equilibrium. In this solution concept, one recommends strategies to all players in the game, with the property that given that all other players stick to their recommendation, the strategy recommended to a player is strictly better than (or at least as good as) any other strategy.

Formally, a Nash equilibrium in a k player game is a vector $\bar{s} = (s_1, \ldots, s_k)$ such that for every player i and any strategy $s_i' \neq s_i$

$$u_i(s_i, \bar{s}_{-i}) > u_i(s_i', \bar{s}_{-i})$$

For weak Nash equilibrium the inequality need not be strict.

An example of a two player game with a Nash equilibrium is the *battle of sexes*, with male and female players who want to go some movie together, but have different tastes in movies.

	Action	Romantic
Action	5; 4	2; 2
Romantic	1; 1	4; 5

```
|----------|----------|
```

There are no dominant strategies, but two Nash equilibria: the top left corner and the bottom right corner. Each player prefers a different Nash equilibrium.

A famous multi-player example is the stable marriage problem, which we will address later in the course.

An example of a game that does not have a Nash Equilibrium is that of *matching pennies*.

```
|----------|----------|
|   1; 0   |   0; 1   |
|----------|----------|
|   0; 1   |   1; 0   |
|----------|----------|
```

An issue that comes up with Nash equilibrium and not in dominant strategies is that games may have multiple Nash equilibria. Moreover, different players may prefer different Nash equilibria. This might make Nash equilibria unstable in practice. (A player may deviate from the recommended strategy, suffering some loss, but also inflicting loss to others, in the hope that other players deviate as well, and that a new Nash equilibrium that is more favorable to the player is reached. A strike by a worker's union may be explained as an attempt by the workers to switch to a different equilibrium point between the workers and the employers.)

Subgame optimality. This notion addresses to some extent the issue of multiple solutions of a game under a given solution concept. It applies to games that take multiple rounds. As players make moves and the game progresses, the portion of the game that remains is called a *subgame*. A profile of strategies is a Nash equilibrium that is subgame perfect if for every subgame (including those subgames that are only reachable by "irrational" moves) the strategies restricted to this subgame form a Nash equilibrium.

There are two different motivations for this notion, that we illustrate by examples.

Chess. There is a difference between having a strategy that plays optimally only from the initial position, and a strategy that plays optimally from any position. For example, assume that it is true that white has a winning strategy in chess. Then the strategy of playing arbitrarily is optimal for black (in the sense that no other strategy guarantees a higher payoff), but not subgame optimal (it does not take advantage of situations in which white has previously blundered).

Ultimatum game. This example illustrates well several aspects of solution concepts. It is a game for splitting 10 dollars among two players. The column player offers how to split the money, and the row player may either accept the split, or reject it, in which case neither player gets anything. For simplicity of the presentation, assume that the first player is allowed to suggest only one of the following three options (1,9),(5,5),(9,1). In extensive form, the game tree then has six leaves, whereas in normal form it has three columns and eight rows. The following matrix illustrates the payoffs for the row player. The payoffs for the column player are 0 or $10 - x$, depending on whether the payoff x for the row player is 0 or more.

0	0	0
0	0	9
0	5	0
0	5	9
1	0	0
1	0	9
1	5	0
1	5	9

The row player has exactly one dominant strategy - to always accept the offer (the last row). The column player does not have any dominant strategy.

The game has several Nash equilibria. For example, one of them is that the row player accepts only (5,5) splits, and the column player offers a (5,5) split. Clearly, the row player prefers this Nash equilibrium to playing his dominant strategy (which leads to the Nash equilibrium of the column player offering a (1,9) split). Hence the row player may appreciate a possibility to manoeuver the game towards one of those Nash equilibria that is better from his perspective than the Nash equilibrium that involves him playing his dominant strategy. However, the dynamics of the game do not allow this. The column player plays first, and then can play no more. If the column player plays (1,9), this results in a subgame in which the row player has only two possible strategies, either to accept or reject. The only subgame optimal decision is to accept. Hence the only *subgame perfect equilibrium* is the one in which the row player always accepts. In the case of the ultimatum game, the notion of subgame perfect equilibrium allows us to select one out of the many Nash equilibria.

It turns out that when the ultimatum game is played in real life (experimental economists have actually experimented with this game), the row players often do not play their dominant strategy. Likewise, the column players do not always offer the (1,9) split. An explanation for this is that in real life scenarios, the payoff matrix does not really represent the true payoffs for the players. Besides the monitory payoffs represented in the payoff matrix, there are other forms payoffs (feeling of pride, feeling of fairness, feeling of creating a reputation) that if properly represented in the description of the game would explain this "irrational" sort of behavior.

1.4 Mixed strategies

In certain games, it is desirable for one of the players to play in a nondeterministic way. Matching pennies is one such example. Given its zero sum game payoff functions, it is clear that if a player's move can be predicted by the other player, he will lose. Another example would be a game of partial information such as poker. If a player plays deterministically (his moves are dictated only by the cards that he holds and previous moves that he has seen), then other players may be able to infer (something about) his cards, and use this information to improve their chance of winning (and hence the player suffers, the game being a zero sum game). Experienced poker players base their play not only on information directly available from the play of the hands, but also on other factors, and hence from the point of view of game theory (that fails to completely model all factors involved), their play is nondeterministic.

A way game theory models the issue of playing in a nondeterministic way is through the concept of a *mixed strategy*, which is a probability distribution over strategies. For example, one mixed strategy for matching pennies is to play each option with probability 1/2. The point is that a-priori, other players may be aware of the probability distribution that a player is using in his mixed strategy (it might be announced, or inferred), but they do not know the actual strategy that is selected until after they select their own strategies. (There are obvious modifications to this last statement when one deals with multi-round games.)

It turns out that the notion of mixed strategies opens up the possibility for more solution concepts. To realize this potential, one assumes that payoff functions are actually utility functions (numerical values are a linear scale on which preferences have exact values), and moreover that players are concerned only with expected payoff (and not the distribution of payoffs). For example, getting a payoff of either 3 or 5, each with probability 1/2, is assumed to be equivalent to getting a payoff of 4 with probability 1. This assumption is often referred to as the players being *risk neutral* (with other options being *risk seeking* or *risk averse*). Risk neutrality is a natural assumption in certain situations. For example, if we assume that a player is involved in many independent games throughout his lifetime each involving a relatively small payoff, and the payoffs from these games add up, then the law of large numbers shows that the total payoff converges to the sum of expectations, and hence the function to optimize per game is indeed the expected payoff.

Another situation in which maximizing the expected payoff is natural (regardless of risk neutrality) is if a game has only two possible payoffs for a player (say, either *win* or *lose*). In this case, maximizing the expected payoff is equivalent to maximizing the probability of winning.

Mixed strategies are sometimes criticized as not being realistic (though they become more realistic when the players are computer programs). The argument is that human players do not really choose randomly among strategies. This relates more to the descriptive aspects of game theory than the prescriptive aspects, and hence is not of major concern to us. However, let us point out that this issue is discussed extensively in game theory literature, and various justifications are offered for mixed strategy. An interesting one is to consider a two player game in which each player is really a "super-player" composed of a population of individuals. Each individual plays a pure strategy. Every individual has random encounters with individuals of the other population. Hence here the randomness is not in the choice of strategies for each individual, but in the choice of encounters, which effectively corresponds to the super-player corresponding to the other population having a mixed strategy. In a Nash equilibrium (over the populations), every individual is playing optimally (in expectation) against the other population. One may say that its strategy is "fitted" to the environment (the environment for an individual is the other population). If the distribution of strategies over populations is not at a Nash equilibrium, then there may be some strategy not currently used by any individual, which is better fitted to the environment. Such a strategy may then "invade" the space of strategies, and be adopted by many individuals ("survival of the fittest"), changing the mixed strategy of the respective super-player.

Two player zero sum games and the minimax theorem.

A game is referred to as *constant sum* if there is some constant c such that regardless of the outcome of the game, the sum of payoffs for the players is c. The special case of $c = 0$ is referred to as a zero sum game. Most results for zero sum games carry over to constant sum games with $c \neq 0$, by scaling the payoffs of each of the n players by an additive factor of c/n (thus making the game 0-sum).

A solution concept for two player zero sum games is offered by the minimax theorem. It assumes risk neutral players that have the conservative goal of securing at least some minimum payoff. For the maximizing player (player 1), this amounts to finding the highest possible value t^+ and a mixed strategy s_1 such that $\min_{s_2} E[u(s_1, s_2)] \geq t^+$. Here u is the payoff function for the maximizing player, and $-u$ is the payoff function for the minimizing player. The strategy s_2 in the above expression ranges over pure strategies. The value t^+ is a max-min value $\max_{s_1} \min_{s_2} E[u(s_1, s_2)]$.

Likewise, the minimizing player seeks a smallest possible value t^- and a strategy s_2 satisfying $\max_{s_1} E[u(s_1, s_2)] \leq t^-$. The value t^- is a min-max value $\min_{s_2} \max_{s_1} E[u(s_m, s_2)]$, where s_2 ranges over mixed strategies, and s_1 ranges over pure strategies.

The famous minimax theory of Von-Neumann says that $t^+ = t^-$. Namely, for every finite constant sum two player game, for mixed strategies the following equality holds:

$$\max_{s_1} \min_{s_2} E[u(s_1, s_2)] = \min_{s_2} \max_{s_1} E[u(s_1, s_2)]$$

The expected payoff at this mutually optimal point is referred to as the *value* of the game. It is essential that strategies are allowed to be mixed in the minimax theorem, as the game of matching pennies illustrates.

Observe that the minimax theorem implies that a player may announce his mixed strategy upfront, and still get the expected payoff guaranteed by the minimax theorem.

A modern proof of the minimax theorem is based on linear programming duality (we will review this proof later), and implies also a polynomial time algorithm for computing the minimax value. von-Neumann's original proof (from 1928) predated the concept of linear programming duality, and applies also to some classes of games with infinite strategy space.

For matching pennies, the value of the game is 0. The unique optimal mixed strategies for the players are to play the two options with equal probability. Note however that in this case, if one of the players plays his optimal mixed strategy, then the other player can play arbitrarily without changing the value of the game. In fact, this is a rather typical situation in many games: when one player chooses an optimal mixed strategy, the other player has a choice of several pure strategies, each of which is optimal (in expectation).

Mixed Nash. A mixed Nash equilibrium (defined and proven to exist by John Nash) is a profile of mixed strategies, one mixed strategy for each player. It has the property that given that the other players follow this profile, no player has an incentive for deviating from his own mixed strategy. That is, every strategy in the support of his mixed strategy is a best response to the mixed strategies of the other players (gives maximum expected payoff among all pure strategies).

In what follows $\bar{s}$ denotes a profile of mixed strategies, s_i denotes the mixed strategy that it associates with player i, and $\bar{s}_{-i}$ denotes the profile of mixed strategies that it associates with the other players. For $\bar{s}$ to be a mixed Nash equilibrium it is required that for every player i and for every pure strategy s'_i for player i,

$$E[u_i(s_i, \bar{s}_{-i})] \geq E[u_i(s'_i, \bar{s}_{-i})]$$

Mixed Nash equilibria exist for every finite game, with any number of players. Nash proved this as a consequence of certain nonconstructive fixed point theorems. Later we shall present an algorithmic proof for the case of two players (though the algorithm is not a polynomial time algorithm). The question of whether there exists a polynomial time algorithm for computing a mixed Nash equilibrium will be discussed as well.

It is common to use the terms *mixed Nash* and *pure Nash* for short, and in the absence of the prefix maixed/pure, assume that the Nash equilibrium is mixed. We shall follow this terminology from now on.

Correlated equilibrium. A mixed Nash is a product distribution over profiles. To realize a mixed Nash, it suffices that each player has a private source of randomness. In contrast, a correlated equilibrium will be an arbitrary distribution over strategy profiles. To realize it, some coordination mechanism involving a common source of randomness is needed. We shall not discuss here in detail what may serve as a coordination mechanism, but rather mention one common example - that of the traffic light. In the absence of any prior agreed upon traffic rules, the game played by two drivers who arrive from perpendicular directions to a traffic junction at the same time resembles the game of *chicken*.

	Stop	Go
Stop	0; 0	0; 2
Go	2; 0	-9; -9

If both players observe a common signal (the traffic light), they can use it to reach a correlated equilibrium (if you see green, go, and if you see red, stop). Once the correlated equilibrium strategy is announced, then whenever two drivers arrive at a junction with a traffic light, it is in the interest of the players to follow its recommendation (if they trust that the other player also follows it).

To be a correlated equilibrium, a distribution over a profile of strategies has to obey the following condition. Observe that for every player, given a recommendation by the coordinating device, there is some marginal probability distribution over the strategy profile of the other players. It is required that the strategy recommended to the player is a best response strategy with respect to this marginal profile.

The notion of a correlated equilibrium was first suggested by Aumann. Like Nash equilibrium, it exists for every finite game (simply because Nash is a special case). It answers two concerns of game theory. One is that it often can offer an equilibrium of higher expected payoff than any Nash equilibrium. The other is that there are polynomial time algorithms for computing it (for games in normal form). We shall discuss the algorithmic aspects later, and here we shall just present an example showing a correlated equilibrium better than any Nash.

2; 1	1; 2	0; 0
0; 0	2; 1	1; 2
1; 2	0; 0	2; 1

A correlated equilibrium that picks each nonzero cell with probability 1/6 has expected payoff 3/2 for each player. Given a recommendation, a player cannot improve the expected payoff by deviating from it. For example, given the recommendation to play the first row, the row player knows that the column player must have received a recommendation to play one of the first two columns, and assuming that the column player follows the recommendation, the first row indeed gives the highest expected payoff. The game has a unique Nash equilibrium, namely, each player chooses a strategy uniformly at random, and if gives an expected payoff of only 1 to each player.

1.5 Summary of introduction

We defined the notion of a game, and presented some solution concepts for games. These include the notions of dominant strategies, pure Nash, subgame perfect equilibrium, minimax value, mixed Nash and correlated equilibrium. We discussed some of the assumptions that are used in making these definitions, and some of the shortcoming of these definitions.

In later parts, we shall accept the definitions of the solution concepts as given, and discuss algorithms for computing them.