Olivier Gossner

Game Theory

Lecture notes. Solution concepts for Normal Form games: $\min\max$**, correlated equilibrium, rationalisable strategies**

Through the M1, you have learned about the Nash equilibrium concept in normal form games. According to this concept, all players randomise independently over their own sets of actions in such a way that each player best-responds to other player's strategies. In this chapter, we study three alternative concepts of importance for normal form games.

Our first concept pertains to games between two players where the gain of one is the loss of the other, called zero-sum games. In such a game, the objectives of both players are purely antagonistic. We define the $\max\min$ of the game as the maximum payoff that player 1 can guarantee, and compare it with the $\min\max$ which is the minimal payoff that player II can defend. The celebrated $\min\max$ theorem tells us that in finite zero-sum games with mixed strategies, the two are equal.

The second concept is correlated equilibrium. In a correlated equilibrium, all players receive information that is not pertinent to the payoffs in the game. Yet, this information may be relevant to how players act in the game. This is the case because all player's informations are correlated. Player i's information is telling about other player's informations, hence about their choices, so that it finally influences i's actions. The concept of correlated equilbrium distribution captures this idea, and extends the Nash equilibrium notion.

Finally we study notion of rationalisable strategies. A strategy for player i is rationalisable when it is a best-response for player i over a conjecture of other player's strategies. Since we will impose that player i conjectures that other players are themselves playing rationalisable strategies, the definition is somehow circular. We will show how to get out of this circularity, how the concept is related to iterated deletion of dominated strategies, and to common knowledge of rationality.

1 Zero-sum games

A *zero-sum* game is a game between two players: 1 and 2. Player i has a finite action A_i, i's set of mixed strategies is $S_i = \Delta(A_i)$. The payoff function of the game is $g\colon A_1 \times A_2 \to \mathbb{R}$, where $g(a_1, a_2)$ is interpreted as a gain for player 1 and a loss to player 2. Since the gain of one player is the loss of the other one, the sum of gains is always zero, and the game is called *zero-sum*.

The question we study are the following:

- Say that player 1 can *guarantee* a payoff of x when player 1 has a strategy that ensures his payoff is at least x, no matter what player 1 chooses. How much can player 1 guarantee in a zero-sum game?
- Say that player 1 can *defend* a payoff of x when, no matter what player 2 plays, player 1 has a strategy that gives him at least x against this particular strategy of 2. How much can player 2 defend in a zero-sum game?
- Can player 1 defend more than he can guarantee? Less? The same?

1.1 Penalty shootouts

This subsection is devoted to the study of an example, penalty shootouts. A good reference on Game Theory applied to penalty shootouts, together with an empirical analysis, is **?**.

Example 1.1 (Penalty shootouts). In the game of penalty shootouts, player 1 kicks the ball, and player 2 tries to stop it.

As a simplification, we allow only two strategies for player 1, shooting left (L) or right (L). Since player 1's natural foot is her right foot, she misses the goal more often if she shoots right than if the shoots left, let α_L be the probability that she sends the ball to the goal while shooting left, and $\alpha_R < \alpha_L$ is the probability that the shoots at the goal while aiming right.

The goalee also has two strategies, as he decides where to try to stop the ball, left (l) or right (r). If she jumps on the side where the ball is shot, she catches it and the penalty is stopped. So, player 1 wins whenever he shoots properly at the goal and player 2 doesn't jump on the correct side. Otherwise, player 2 wins.

What will the shooter and goalee do?

Questions:

1. Will the shooter aim more often to the left (easier) or to the right?
2. Will the goalee try to stop the ball more often by jumping to the left or to the right?

We write the game as a zero-sum game between players 1 and 2. The payoff matrix to player 1, taking into account the probability of shooting properly, is shown in figure 1. The payoff to player 2 is not shown, it is the opposite of the payoff to player 1.

	l	r
L	0	α_L
R	α_R	0

Figure 1: Payoff matrix in the penalty shootout game

1.1.1 How much can player 1 guarantee?

In this part, we consider that player 1 chooses a strategy, i.e., a probability p of shooting right (and $1-p$ of shooting left). Then player 2, plays a best-response to player 1. For instance, we can assume that player 1 has a long record of shooting, and that player 2's coach has studied all past shoutouts of player 1. In this scenario, what is the best probability p that player 1 can choose and what is the corresponding payoff? We represent the situation by drawing, for each probability p of player 1, player 1's payoff if player 2 chooses l or r. Given p, player 1 receives the minimum of the two, since player 2 can choose the action that minimses 1's payoff. This leads to a concave function of p. Finally, 1 chooses the value of p that maximises this minimum of the two payoffs. In our case, this is achieved at the value $p^* = \frac{\alpha_L}{\alpha_L+\alpha_R}$. And the maximum that player 1 can guarantee is $\underline{v} = \frac{\alpha_L \alpha_R}{\alpha_L+\alpha_R}$, where the max min $\underline{v}$ is defined as

$$\underline{v} = \max_p \min_q g(p,q),$$

where $g(p,q)$ is player 1's payoff when mixed strategies p,q are used. See figure 2 for a graphical analysis.

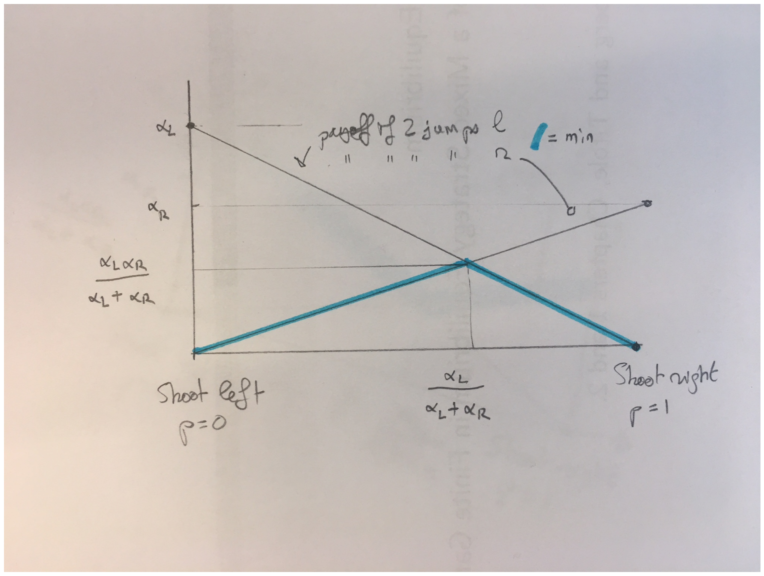

Figure 2: Computation of the max min in the game of penalty shootouts.

1.1.2 How much can player 1 defend?

Now we go the opposite direction. We assume that player 2 chooses a strategy represented by a probability q of jumping right (and $1-q$ on the left). Based on the knowledge of player 2's strategy, player 1 can choose whether to shoot left or right. In order to find the best possible strategy for player 2, we represent, for each value of q, player 1's payoff if 1 choose L, and R. For any value of q, player 1 chooses the maximum between the two, so that we consider the maximum of the two graphs, which corresponds to a convex function, see figure 3. Player 2's objective is to minimise this convex function, and this is done at the value $q^* = \frac{\alpha_R}{\alpha_L+\alpha_R}$. The corresponding maximum value that player 1 can defend is $\overline{v} = \frac{\alpha_L\alpha_R}{\alpha_L+\alpha_R}$. It is called the min max and is defined as:

$$\underline{v} = \min_q \max_p g(p,q)$$

1.1.3 Comments

We remark that, in the example $\underline{v} = \overline{v}$. This is not a coincidence, but a consequence of the min max theorem that we will see below.

We also remark that $p^* > 1/2$, this means that the shooter kicks more often on his weak side than on his strong side. This comes from the fact that, if it were the opposite, the goalee would

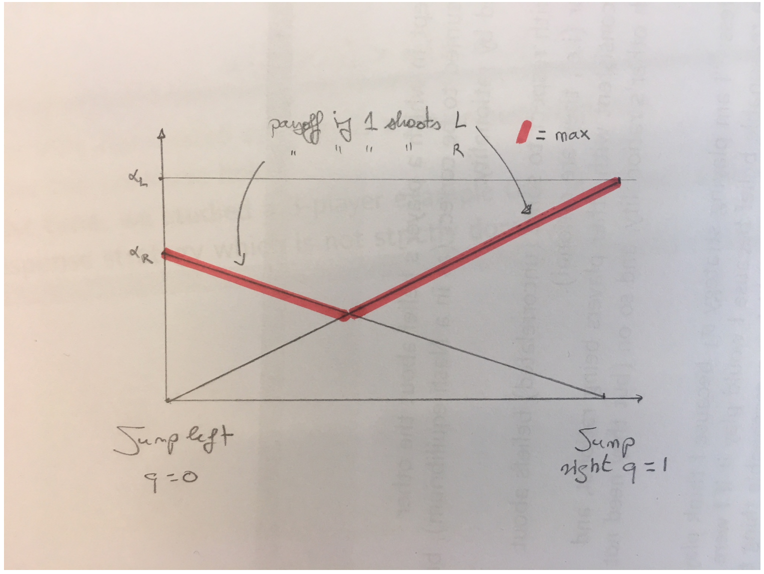

Figure 3: Computation of the min max in the game of penalty shootouts.

try to stop the ball always on the right side, and hence, the shooter would have incentives to augment the probability with with he shoots on the left side.

Similarly, $q^* < 1/2$, which means that the goalee tries to stop the ball more often on the strong side of the kicker than on his weak side. If she did try to stop the ball more often on his weak side, the goalee would want to shoot always on the right, and the goalee would have incentives to increase the probability with with she tries to stop the ball on the right side.

1.2 The min max theorem

Consider a zero-sum game (A_1, A_2, g) between player 1 and 2, where A_1, A_2 are finite. We let the spaces of mixed strategies be $S_i = \Delta(A_i)$, and extend the payoff function g to $S_1 \times S_2$ by letting:

$$g(s_1, s_2) = \sum_{a_1, a_2} s_1(a_1) s_2(a_2) g(a_1, a_2) = \mathbb{E}_{s_1, s_2} g(a_1, a_2)$$

Theorem 1.2 (min max theorem)**.** *For any zero-sum game G with finite action spaces:*

1.

$$\sup_{s_1 \in S_1} \inf_{s_2 \in S_2} g(s_1, s_2) = \inf_{s_2 \in S_2} \sup_{s_1 \in S_1} g(s_1, s_2)$$

and their common value is called the value *of the game,* $v(G)$.

2. *there exist* optimal strategies s_1^*, s_2^* *such that:*

$$\inf_{s_2 \in S_2} g(s_1^*, s_2) = \sup_{s_1 \in S_1} g(s_1, s_2^*) = v(G)$$

Remark 1.3. *Without making any assumptions on the game, we have the inequality*

$$\sup_{s_1 \in S_1} \inf_{s_2 \in S_2} g(s_1, s_2) \leq \inf_{s_2 \in S_2} \sup_{s_1 \in S_1} g(s_1, s_2),$$

thus, the maximum that player 1 can guarantee is always less or equal than what this player defends. This reflects the fact that the situation in which player 1 has to choose first a strategy, and player 2 can best-respond, is less advantageous to player 1 than the situation in which player has to choose a strategy, and player 1 best-responds.

To prove the inequality, see that for any $t_1 \in S_1$ *and* $t_2 \in S_2$*, we have*

$$\inf_{s_2 \in S_2} g(t_1, s_2) \leq g(t_1, t_2) \leq \sup_{s_1 \in S_1} g(s_1, t_2)$$

so that for fixed t_1*, by taking the inf on the right side:*

$$\inf_{s_2 \in S_2} g(t_1, s_2) \leq \inf_{s_2 \in S_2} \sup_{s_1 \in S_1} g(s_1, s_2)$$

and the inequality obtains by taking the supremum of the left side over all t_1*.*

Remark 1.4 (A proof of the min max theorem through the Nash existence theorem)**.** *The* min max *theorem can be seen as a consequence of the Nash existence theorem. Indeed, if we consider a Nash equilibrium* (s_1^*, s_1^*)*, the Nash equilibrium property imply precisely that* s_1^* *maximises 1's payoff over all* $s_1 \in S_1$ *against* s_2^**, and that* s_2^* *minimises 1's payoff over all* $s_2 \in S_2$ *against* S_1*. We thus have:*

$$\inf_{s_2 \in S_2} g(s_1^*, s_2) = \sup_{s_1 \in S_1} g(s_1, s_2^*).$$

This means that player 1 *can guarantee as much as he can defend, taking first a sup on the left hand side, then the inf on the right-hand side, shows:*

$$\sup_{s_1 \in S_1} \inf_{s_2 \in S_2} g(s_1, s_2) \geq \inf_{s_2 \in S_2} \sup_{s_1 \in S_1} g(s_1, s_2).$$

If we combine this with the inequality in remark 1.3, we obtain the stated properties in the min max *theorem.*

Remark 1.5. *The second point of the* min max *theorem tells us that the* sup *is a* max*, i.e. is achieved at some point, and the same for the* inf*. Given the assumptions of our theorem, this is not complicated to see. Indeed, for every* s_1 *we have:*

$$\inf_{s_2 \in S_2} g(s_1, s_2) = \inf_{a_2 \in A_2} g(s_1, s_2) = \min_{a_2 \in A_2} g(s_1, a_2).$$

and given that A_2 is finite, the function $s_1 \mapsto \min_{a_2 \in A_2} g(s_1, a_2)$ is continuous, hence attains its minimum on S_1. A similar argument applies for player 2. For more complex games, the existence of optimal strategies, x‘ i.e. achieving the min *and the* max*, is a lot more complicated problem, and sometimes an open one.*

The proof of the min max theorem through the Nash existence theorem is complete and rigorous, nothing is missing. Yet, we are going to present an alternative proof of the min max theorem. The reason is that the Nash existence theorem is a complex one, derived through Kakutani or Brouwer's fixed point theorems, which assert the existence of fixed points for general functions or correspondences. In a mathematical sense, the min max theorem is much more basic and simple, as it relies only on duality/separation theory. Hence we are going to show a proof that uses only the required tools.

We present a visual proof of the min max theorem when the set of strategies S_1 for player 1 is an arbitrary compact and convex set, and player 2 has a finite number of actions $A_2 = \{1, \dots, J\}$. We define the set of achievable payoffs by 1 defined as the following subset of $\mathbb{R}^J$

$$X = \{(g(s_1, j))_{j \in A_2}, s_1 \in S_1\}$$

When considering the max min, player 1 chooses a point $x \in X$ that maximises $\min\{x_j, j \in A_2\}$. This point is found by looking the furthest point in X in the NE direction along the first diagonal. Let x^* be a corresponding such maximiser. We need to show that the min max is less or equal than $v = \min\{x^*_j, j \in A_2\}$, i.e that there exists a mixed strategy of player 2 against which player 1 cannot achieve more than v.

We first separate coordinates for which x^* achieves more than v from other ones. Let $\tilde{J} = \{j \in A_2, x^*_j = v\}$. For coordinates $j \notin \tilde{J}$, $x^*_j > v$. Let $\tilde{X}$ be the projection of $\mathbb{R}^{A_2}$ on $\mathbb{R}^{\tilde{J}}$, and consider the subset of $\mathbb{R}^{\tilde{J}}$: $\tilde{Y} = \{x \in \mathbb{R}^{\tilde{J}}, \forall j \in \tilde{J}\, x_j > v\}$. $\tilde{Y}$ is an open, convex set and it doesn't intersect $\tilde{X}$ since otherwise player 1 could guarantee more than v.

By the separating hyperplane theorem, there exists a linear mapping $f : \mathbb{R}^{\tilde{J}} \to \mathbb{R}, x \mapsto \sum_{j \in \tilde{J}} \alpha_j x_j$ and a number $c \in \mathbb{R}$ such that

1. $f(x) > c$ for every $x \in \tilde{Y}$
2. $f(x) \leq c$ for every $x \in \tilde{X}$.

Since f has to be increasing in all arguments $\alpha_j \geq 0$ for all $j \in \tilde{J}$, and at least one of α_j is not 0. We also deduce by continuity that $f((v, \dots, v)) = \sum_{j \in \tilde{J}} \alpha_j v = c$, so that $v = \frac{c}{\sum_{j \in \tilde{J}} \alpha_j}$.

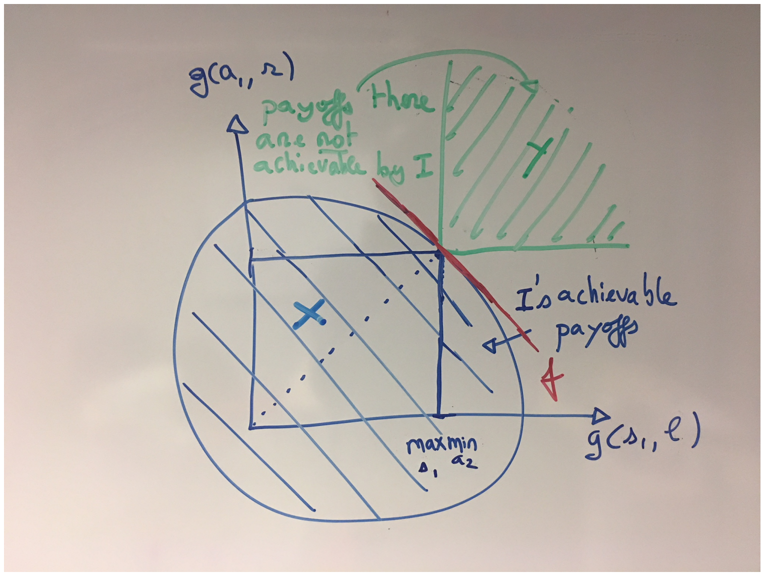

Figure 4: Illustration of the proof of the min max theorem when player 1's action set is compact convex, player 2's action set is $\{l, r\}$, and g is continuous.

Our next step is to reinterpret the coefficients α_j as a mixed strategy for player 2. Let $s_2^*(j) = \frac{\alpha_j}{\sum_{j' \in J'} \alpha'_j}$ for $j \in \tilde{J}$, and $s_2^*(j) = 0$ for $j \notin \tilde{J}$. Consider any strategy a_1 of player 1, and the corresponding point x in X. We have

$$\begin{aligned} g(a_1, s_2^*) &= \sum_{j \in \tilde{J}} s_2^*(j) x_j \\ &= \frac{1}{\sum_{j' \in J'} \alpha'_j} \sum_{j \in \tilde{J}} \alpha_j x_j \\ &\leq \frac{1}{\sum_{j' \in J'} \alpha'_j} c \\ &= v \end{aligned}$$

We have thus shown that, against the strategy s_2^*, player 1 cannot obtain more than v and thus:

$$\min_{s_2 \in S_2} \max_{s_1 \in S_1} g(s_1, s_2) \leq v$$

Now we can conclude. Since we also know that

$$v = \max_{s_1 \in S_1} \min_{s_1 \in S_1} g(s_1, s_2) \leq \min_{s_2 \in S_2} \max_{s_1 \in S_1} g(s_1, s_2),$$

we have proven that:

$$\max_{s_1 \in S_1} \min_{s_2 \in S_2} g(s_1, s_2) = \min_{s_2 \in S_2} \max_{s_1 \in S_1} g(s_1, s_2) = v,$$

which is the min max theorem.

2 Correlated equilibrium

In a Nash equilibrium, all players choose actions independently, and possibly randomly. The Nash condition says that no player wants to change her own strategy, given the strategies of the other ones. There are several reasons why players may not randomize independently. First, they may have access to information based on which they make their own decisions. Since the information received by different agents is correlated, so are their actions. Second, we may think as a distribution of actions as a subjective belief in a player's mind. Assume for instance a game with several Nash equilibria, and one player who believes that others are playing one of these, but not knowing which; in such a player's mind, other player's actions are not independent, they are correlated.

Example 2.1 (Public randomisation in the battle of sexes). Consider the game of battle of sexes. Both players have a choice between which concert to go to, either B(ach), or S(stravinsky). Player 1 has a preference for Bach, and player 2 for Stravinsky. Players' preferences include two components, according to the first component, players like to be with each other, and according to the second one, they like to go to their preferred concert.

	B	S
B	2.1	0,0
S	0,0	1.2

Figure 5: Game of the battle of sexes

This game has two Nash equilibria in pure strategies: (B,B) and (S,S). It also has an equilibrium in mixed strategies: $(2/3B+1/3S, 1/3B+2/3S)$. The mixed strategies equilibrium yields an expected payoff of $2/3$ to each player, which is Pareto dominated by the payoff of each pure equilibrium. Pure Nash equilibria are better than the mixed one, but they are unfair in the sense that each provides more advantage to one player than to the other. A simple way that players can choose between the two pure Nash equilibria is by flipping a fair coin, then play (B,B) if the coin shows Heads, and (S,S) if the coin shows Tails.

Convince yourself that, in the game in which players first see the (common) outcome of the coin, then decide to which concert they go, these strategies do form a Nash equilibrium. Assuming that players follow these strategies, what is their distribution of actions? I.e. with what probability do they end up playing each of the 4 action profiles (B,B), (B,S), (S,B), and (S,S)?

The previous example shows how a correlation device can help randomize between equilibria. In the next example, we show how the use of a randomization device can be useful beyond randomization over Nash equilibria.

Example 2.2 (Rebel without a cause). In the movie "Rebel without a cause" (1955), two cars race towards towards a cliff. See what happens here: `https://www.youtube.com/watch?v=u7hZ9jKrwvo`. The rule of the game is ""We are both heading for the cliff, who jumps first, is the Chicken". We analyse this game by assuming each player has two strategies, "tough" (T), or "chicken" (C). If both players play tough, they both jump off the cliff. If one plays tough and the other chicken, the tough one wins, and the chicken loses the face. If both play chicken, they stop at the same time, so that no-one loses the face and both players can be good friends. The payoff matrix given by figure 6.

	C	T
C	6.6	2,7
T	7,2	0.0

Figure 6: Game of chicken

This game has two Nash equilibria in pure strategies: (T,C) and (C,T). It also has a symmetric mixed strategies equilibrium: $(2/3C+1/3T, 2/3T+1/3T)$. The symmetric equilibrium yields an expected payoff of $14/3$ to each player, which is Pareto dominated by a randomisation over the two pure Nash equilibria with equal probabilities.

How can players generate a payoff which is equal to both of them (fairness) and even higher than randomising between the two pure Nash equilibria? Consider the following system of *traffic lights*. Assume that each player sees a light which may either be red (R) or green (G). No player can see the light of the other one. There are three possibilities: the light of player 1 can be red while the light of player 2 is green, or the other way round, or both lights may be red. The three possibilities have the same probability. We summarize the probabilities of signal pairs in figure 7.

	R	G
R	$\frac{1}{3}$	$\frac{1}{3}$
G	$\frac{1}{3}$	0

Figure 7: Probabilities of signal pairs from traffic light: player 1's signals are in rows, and player 2's are in columnds.

Now consider the following recommendations to the players: play T if your light is green, and C if your light is red.

We check that these recommendations form a Nash equilibrium of the game with signals. If a player receives a green signal, this player knows that the other must have received a red signal,

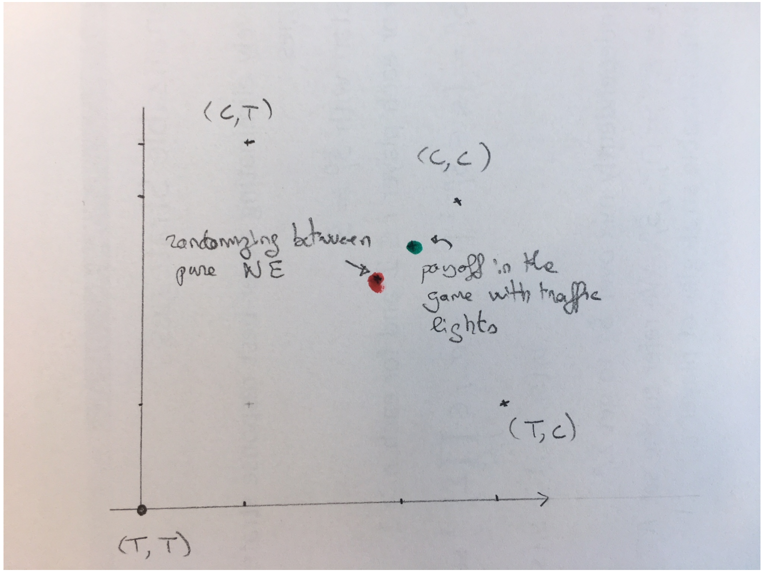

Figure 8: Payoffs in the game of chicken

hence will play C. The only best-response is to play T in this case, which corresponds to the recommended action. If a player receives a red signal, this player puts probability 1/2, 1/2 on the other one receiving a red or a green signal, hence probability 1/2, 1/2 on his opponent playing C or T. Against the strategy $1/2C + 1/2T$, playing C yields an expected payoff of 4, which is more than the expected payoff of $7/2$ obtained by playing T, hence the recommended action C is the only best-response. We have thus shown that following the recommandations form a Nash equilibrium.

When players follow the recommendations, they receive an expected payoff of 5 each, which is strictly larger than $9/2$ that obtained by randomising over the Nash equilibria.

2.1 Correlated Equilibrium Distributions

We consider a game G with finite action sets A_i. We let $A = \Pi_i A_i$. Player i's payoff function is $g_i \colon A \to \mathbb{R}$.

Definition 2.3. A *correlated equilibrium distribution* of G is a distribution $\mu \in \Delta(A)$ such that, for player i, and actions $a_i, b_i \in A_i$:

$$\sum_{a_{-i} \in A_{-i}} \mu(a_i, a_{-i}) g_i(a_i, a_{-i}) \geq \sum_{a_{-i} \in A_{-i}} \mu(a_i, a_{-i}) g_i(b_i, a_{-i}) \quad (1)$$

We let $C(G)$ be the set of correlated equilibrium distributions of G.

In words. Consider the situation in which a *mediator* picks an action profile $a \in A$ according to the distribution μ, and informs each player i of the component a_i of A. The distribution μ is a correlated equilibrium distribution of G whenever the strategies that consist in playing the recommended actions form a Nash equilibrium. The inequalities (1) precisely state that no player i has incentives to deviate to action $b_i \neq a_i$ when action a_i is recommended.

What can be said about the set of correlated equilibrium distributions?

Remark 2.4 (CED form a convex compact polyhedron)**.** *If we look at the system of equations (1), we see that they are a finite set of weak inequalities on* μ*. This means that the set of correlated equilibrium distributions* $C(G)$ *is a* closed, convex polyhedron*. It is also bounded as it is included in the set of all distributions, hence it is a* compact *set. In geometrical terms, this is a much simpler type of set than the set of Nash equilibria, which are defined by* algebraic *equations, i.e. ones that involve polynomials.*

Remark 2.5 (Contains the Nash equilibrium distributions)**.** *Consider any Nash equilibrium in mixed strategies* $s = (s_i)_i$ *of G, and the corresponding Nash equilibrium distribution* μ_s *given by* $\mu_s(a) = \Pi_i s_i(a_i)$*. One can check easily that* μ_s *is a correlated equilibrium distribution of G. Indeed, once knowing what action she has chosen, no player has incentives to change this action to another one. This implies that every Nash equilibrium distribution is a correlated equilibrium distribution. Furthermore the set of correlated equilibrium distributions contains the convex hull of all the Nash equilibrium distributions, i.e. the set of all convex distributions of Nash equilibrium distributions.*

Remark 2.6 (Is not empty by the Nash existence theorem, but this is cheating!)**.** *Since any finite game has a Nash equilibrium in mixed strategies, it also has a correlated equilibrium distribution. This shows that the set of correlated equilibrium distributions is non-empty for any finite game. This proof of existence is correct, but requires the use of a very sophisticated instrument (fixed point theorems) in order to show to study a simple problem, which is a linear one (defined by a finite set of linear weak inequalities).* **?** *present a basic proof of the existence of correlated equilibrium distributions that relies only on the* $\min\max$ *theorem.*

2.2 Correlated Equilibria

So far, we have studied games in which players receive signals that correspond to action recommendations in G, and we looked for the joint distributions of signals for which every player has incentives to follow these recommendations, should other players do the same. More generally, we can allow the set of signals to be any set, so that signals do not necessarily correspond to actions in G.

For instance, look back at the game of "Battle of Sexes", we can ask what happens if players do not have a coin to flip, but may instead roll a die. The roll of the die, with its value $1,\ldots,6$, is then observed by both players. Players can agree on the following type of rule: if the result is odd, coordinate on B, and if it even coordinate on S. We see that the strategies that correspond to following the recommendations form a Nash equilibrium of the game in which players observe the result of the die, then choose an action in the game.

I we break this down a little bit, this means that, for each player i, we have a set X_i of signals, and there is a distribution $\mu \in \Delta(X)$ over profiles of signals. We are considering the game which:

1. A joint signal x is drawn in X according to μ,
2. Each player i observes the signal x_i
3. Each player i then chooses an action a_i, and we let $a = (a_i)$ the action profile
4. The payoff to each player i is $g_i(a)$.

A behavioral strategy for player i in this game is a mapping $f_i \colon X_i \to S_i$, and if $f \colon X \to S$ is the profile of strategies of the players, the corresponding vector payoff is:

$$\gamma(f) = \sum_x \mu(x) g(f(x)).$$

We note by $\Gamma(\mu, G)$ the game above in which each player i chooses f_i and the vector payoff function is γ. Note that each strategy profile also induces a distribution of action profiles $\nu(f) \in \Delta(A)$ given by:

$$\nu(f)(a) = \sum_x \mu(x) f(x)(a).$$

When f is a Nash equilibrium of $\Gamma(\mu, G)$, we say that the corresponding distribution $\nu(f)$ is a *correlated equilibrium distribution of G induced by μ*. We let $CED(\mu, G)$ the set of such correlated equilibrium distributions induced by μ.

The question that we are asking is whether 1) allowing for any set of signals (instead of imposing $X_i = A_i$) and 2) allowing players to choose any strategy (instead of just following the recommendation) increases the set of equilibrium distributions generated. In fact, it doesn't, as shown by the following Theorem:

Theorem 2.7 (Revelation principle)**.** *For every set of signal profiles X and every distribution μ on X,*

$$CED(\mu, G) \subseteq C(G).$$

Note also that, by the definitions, whenever μ is a correlated equilibrium distribution, then μ is a correlated equilibrium distribution induced by itself:

$$\mu \in C(G) \Rightarrow \mu \in CED(\mu, G).$$

The revelation principle together with the above remark imply that the set of correlated equilibria is the set of all correlated equilibria induced by all possible distributions over all possible sets of signals.

The proof of the revelation principle is relatively straightforward. Consider a game G, a set of signal profiles X, a distribution $\mu \in \Delta(X)$ and an equilibrium f of $\Gamma(\mu, G)$ that induces a distribution ν on A. Let P be the joint probability distribution on X and A induced by ν and f.

Consider the following "thought experiment": a signal profile x is drawn according to μ, then a in A is drawn according to $f(x)$. Player i is informed of x_i and of a_i. Assuming players $j \neq i$ play action a_j, and player i is allowed to change from a_i to any other action, should player i do so? Of course not, since a_i is an action chosen with positive probability after signal x_i under the equilibrium strategies f, it is optimal to play a_i after receiving signal x_i.

Consider the similar thought experiment, but in which player i is informed of a_i only (and not of x_i). Can it be optimal now to change from a_i to any other action? Intuitively not: since a_i would be optimal after any possible x_i, it should also be optimal when x_i is not known! To formalize this, let $B(a_i)$ be the set of beliefs on A_{-i} against which the action a_i is optimal. This is the set of beliefs for which no other action a'_i gives a higher payoff than a_i:

$$B(a_i) = \cap_{a'_i \in A_i} \{q \in \Delta(A_{-i}), \mathbb{E}_q g(a_i, a_{-i}) \geq \mathbb{E}_q g(a'_i, a_{-i})\}.$$

Now, note that $B(a_i)$ is a convex polyhedron, as it is defined by a finite family of linear inequalities. The "convex" part is what we are interested for the following. We know that a_i is an optimal action after each pair (x_i, a_i) that happens with positive probability under P (by the point above). Therefore we know that for every (x_i, a_i) such that $P(x_i, a_i) > 0$, the conditional probability $P(a_{-i}|x_i, a_i)$ on A_{-i} satisfies

$$P(a_{-i}|a_i) \in B(a_i).$$

We want to show that, after receiving a_i, it is optimal to choose action a_i, thus we need to prove that the conditional probability $P(a_{-i}|a_i)$ also satisfies

$$P(a_{-i}|x_i, a_i) \in B(a_i).$$

This is in fact a consequence of the convexity of $B(a_i)$, and of the following implication of Bayes's rule;

$$P(a_{-i}|x_i, a_i) = \sum_{x_i} P(x_i|a_i) P(a_{-i}|x_i, a_i).$$