Chapter 2: Two-person zero-sum games

February 24, 2010

In this section we study games with only two players. We also restrict attention to the case where the interests of the players are completely antagonistic: at the end of the game, one player gains some amount, while the other loses the same amount. These games are called "two person zero sum games".

Military games such as pursuit-evasion problems, are a rich source of two-person zero-sum games. While in most economics situations the interests of the players are neither in strong conflict nor in complete identity, this specific class of games provides important insights into the notion of "optimal play". In some 2-person zero-sum games,each player has a well defined "optimal" strategy, which does not depend on her adversary decision (strategy choice). In other games, no such optimal strategy exists. Finally, the founding result of Game Theory, known as the *minimax theorem,* says that optimal strategies exist when our players can randomize over a finite set of deterministic strategies.

1 Games in strategic form

A two-person zero-sum game <u>in strategic form</u> is a triple $G = (S, T, u)$, where S is a set of strategies available to the player 1, T is a set of strategies available to the player 2, and $u : S \times T \to \mathbf{R}$ is the payoff function of the game G; i.e., $u(s, t)$ is the resulting gain for player 1 and the resulting loss for player 2, if they choose to play s and t respectively. Thus, player 1 tries to maximize u, while player 2 tries to minimize it. We call any strategy choice (s, t) an *outcome* of the game G.

When the strategy sets S and T are finite, the game G can be represented by an n by m matrix A, where $n = |S|$, $m = |T|$, and $a_{ij} = u(s_i, t_j)$.

The <u>secure utility level</u> for player 1 (the minimal gain he can guarantee himself, no matter what player 2 does) is given by

$$\underline{m} = \max_{s \in S} \min_{t \in T} u(s, t) = \max_i \min_j a_{ij}.$$

A strategy s^* for player 1 is called *prudent*, if it realizes this secure max-min gain, i.e., if $\min_{t \in T} u(s^*, t) = \underline{m}$.

The secure utility level for player 2 (the maximal loss she can guarantee herself, no matter what player 1 does) is given by

$$\overline{m} = \min_{t \in T} \max_{s \in S} u(s,t) = \min_{j} \max_{i} a_{ij}.$$

A strategy t^* for player 2 is called *prudent*, if it realizes this secure min-max loss, i.e., if $\max_{s \in S} u(s,t^*) = \overline{m}$.

The secure utility level is what a player can get for sure, even if the other player behaves in the worst possible way. For each strategy of a player we calculate what could be his or her worst payoff, resulting from using this strategy (depending on the strategy choice of another player). A prudent strategy is one for which this worst possible result is the best. Thus, by a prudent choice of strategies, player 1 can guarantee that he will gain at least $\underline{m}$, while player 2 can guarantee that she will loose at most $\overline{m}$. Given this, we should expect that $\underline{m} \leq \overline{m}$. Indeed:

Lemma 1 *For all two-person zero-sum games, $\underline{m} \leq \overline{m}$.*

Proof: $\underline{m} = \max_{s \in S} \min_{t \in T} u(s,t) = \min_{t \in T} u(s^*,t) \leq u(s^*,t^*) \leq \max_{s \in S} u(s,t^*) = \min_{t \in T} \max_{s \in S} u(s,t) = \overline{m}$.

Definition 2 *If $\underline{m} = \overline{m}$, then $m = \underline{m} = \overline{m}$ is called* the value *of the game G. If $\underline{m} < \overline{m}$, we say that G has no value.*

An outcome $(s^*,t^*) \in S \times T$ is called a saddle point *of the payoff function u, if $u(s,t^*) \leq u(s^*,t^*) \leq u(s^*,t)$ for all $s \in S$ and for all $t \in T$.*

Remark 3 *Equivalently, we can write that $(s^*,t^*) \in S \times T$ is a saddle point if $\max_{s \in S} u(s,t^*) \leq u(s^*,t^*) \leq \min_{t \in T} u(s^*,t)$*

When the game is represented by a matrix A, (s^*,t^*) will be a saddle point, if and only if $a_{s^*t^*}$ is the largest entry in its column and the smallest entry in its row.

A game has a value if and only if it has a saddle point:

Theorem 4 *If the game G has a value m, then an outcome (s^*,t^*) is a saddle point if and only if s^* and t^* are prudent. In this case, $u(s^*,t^*) = m$. If G has no value, then it has no saddle point either.*

Proof. Suppose that $m = \underline{m} = \overline{m}$, and s^* and t^* are prudent strategies of players 1 and 2 respectively. Then by the definition of prudent strategies

$$\max_{s \in S} u(s,t^*) = \overline{m} = m = \underline{m} = \min_{t \in T} u(s^*,t).$$

In particular, $u(s^*,t^*) \leq m \leq u(s^*,t^*)$; hence, $u(s^*,t^*) = m$. Thus, $\max_{s \in S} u(s,t^*) = u(s^*,t^*) = \min_{t \in T} u(s^*,t)$, and so (s^*,t^*) is a saddle point. Conversely, suppose that (s^*, t^*) is a saddle point of the game, i.e., $\max_{s \in S} u(s, t^*) \leq u(s^*, t^*) \leq \min_{t \in T} u(s^*, t)$. Then, in particular, $\max_{s \in S} u(s, t^*) \leq \min_{t \in T} u(s^*, t)$. But by the definition of $\underline{m}$ as $\max_{s \in S} \min_{t \in T} u(s, t)$ we have $\min_{t \in T} u(s^*, t) \leq \underline{m}$, and by the definition of $\overline{m}$ as $\min_{t \in T} \max_{s \in S} u(s, t)$ we have $\max_{s \in S} u(s, t^*) \geq \overline{m}$. Hence, using Lemma 1 above, we obtain that $\min_{t \in T} u(s^*, t) \leq \underline{m} \leq \overline{m} \leq \max_{s \in S} u(s, t^*)$. It follows that $\overline{m} = \max_{s \in S} u(s, t^*) = u(s^*, t^*) = \min_{t \in T} u(s^*, t) = \underline{m}$. Thus, G has a value $m = \underline{m} = \overline{m}$, and s^* and t^* are prudent strategies. $\blacksquare$

Example 1 *Matching pennies* is the simplest game with no value: each player chooses Left or Right; player 1 wins +1 if their choices coincide, loses 1 otherwise.

Example 2 The *noisy gunfight* is a simple game with a value. The two players walk toward each other, with a single bullet in their gun. Let $a_i(t), i = 1, 2$, be the probability that player i hits player j if he shoots at thime t. At $t = 0$, they are far apart so $a_i(0) = 0$; at time $t = 1$, they are so close that $a_i(1) = 1$; finally a_i is a continuous and increasing function of t. When player i shoots, one of 2 things happens: if j is hit, , player $i$wins \$1 from j and the game stops (j cannot shoot any more); if i misses, j hears the shot, and realizes that i cannot shoot any more so j waits until $t = 1$, hits i for sure and collects \$1from him. Note that the *silent* version of the *gunfight* model (in the problem set below) has no value.

In a game with a value, prudent strategies are optimal—using them, player 1 can guarantee to get at least m, while player 2 can guarantee to loose at most m.

In order to find a prudent strategy:

– player 1 solves the program $\max_{s \in S} m_1(s)$, where $m_1(s) = \min_{t \in T} u(s, t)$ (maximize the minimal possible gain);

– player 2 solves the program $\min_{t \in T} m_2(t)$, where $m_2(t) = \max_{s \in S} u(s, t)$ (minimize the maximal possible loss).

We can always find such strategies when the sets S and T are finite.

Remark 5 *(Infinite strategy sets) When S and T are compact (i.e. closed and bounded) subsets of $\mathbf{R}^k$, and u is a continuous function, prudent strategies always exist, due to the fact that any continuous function, defined on a compact set, reaches on it its maximum and its minimum.*

In a game without a value, we cannot deterministically predict the outcome of the game, played by rational players. Each player will try to guess his/her opponent's strategy choice. Recall matching pennies.

Here are several facts about two-person zero-sum games in normal form.

Lemma 6 *(rectangularity property) A two-person zero-sum games in normal form has at most one value, but it can have several saddle points, and each player can have several prudent (and even several optimal) strategies. Moreover,* *if (s_1, t_1) and (s_2, t_2) are saddle points of the game, then (s_1, t_2) and (s_1, t_2) are also saddle points.*

A two-person zero-sum games in normal form is called *symmetric* if $S = T$, and $u(s,t) = -u(t,s)$ for all s, t. When S, T are finite, symmetric games are those which can be represented by a square matrix A, for which $a_{ij} = -a_{ji}$ for all i, j (in particular, $a_{ii} = 0$ for all i).

Lemma 7 *If a symmetric game has a value then this value is zero. Moreover, if s is an optimal strategy for one player, then it is also optimal for another one.*

Proof. Say the game (S, T, u) has a value v, then we have

$$v = \max_s \min_t u(s,t) = \max_s \{-\max_t u(t,s)\} = -\min_s \max_t u(t,s) = -v$$

so $v = 0$. The proof of the 2d statement is equally easy. ■

2 Games in extensive form

A *game in extensive form* models a situation where the outcome depends on the consecutive actions of several involved agents ("players"). There is a precise sequence of individual moves, at each of which one of the players chooses an action from a set of potential possibilities. Among those, there could be chance, or random moves, where the choice is made by some mechanical random device rather than a player (sometimes referred to as "nature" moves).

When a player is to make the move, she is often unaware of the actual choices of other players (including nature), even if they were made earlier. Thus, a player has to choose an action, keeping in mind that she is at one of the several possible actual positions in the game, and she cannot distinguish which one is realized: an example is bridge, or any other card game.

At the end of the game, all players get some payoffs (which we will measure in monetary terms). The payoff to each player depends on the whole vector of individual choices, made by all game participants.

The most convenient representation of such a situation is by a *game tree*, where to non terminal nodes are attached the name of the player who has the move, and to terminal nodes are attached payoffs for each player. We must also specify what information is available of a player at each node of the tree where she has to move.

A **strategy** is a full plan to play a game (for a particular player), prepared in advance. It is a *complete specification* of what move to choose in any potential situation which could arise in the game. One could think about a strategy as a set of instructions that a player who cannot physically participate in the game (but who still wants to be the one who makes all the decisions) gives to her "agent". When the game is actually played, each time the agent is to choose a move, he looks at the instruction and chooses according to it. The representative, thus, does not make any decision himself!

Note that the reduction operator just described does not work equally well for games with n -players with multiple stages of decisions.

Each player only cares about her final payoff in the game. When the set of all available strategies for each player is well defined, the only relevant information is the profile of final payoffs for each profile of strategies chosen by the players. Thus to each game in extensive form is attached a *reduced* game in strategic form. In two-person zero sum games, this reduction is not conceptually problematic, however for more general n-person games, it does not capture the dynamic character of a game in extensive form, and for this we need to develop new equilibrium concepts: see Chapter 5.

In this section we discuss games in extensive form *with perfect information.*

Example 3 *Gale's chomp game*: the player take turns to destroy a $n \times m$ rectangular grid, with the convention that if player i kills entry (p, q), all entries (p', q') such that $(p', q') \geq (p, q)$ are destroyed as well. When a player moves, he must destroy one of the remaining entries.The player who kills entry $(1, 1)$ loses. In this game player 1 who moves first has an optimal strategy that guarantees he wins. This strategy is easy to compute if $n = m$, not so if $n \neq m$.

Example 4 *Chess and Zermelo's theorem.* The game of Chess has three payoffs, $+1, -1, 0$. Although we do not know which one, one of these 3 numbers is the value of the game, i.e., either White can guarantee a win, or Black can, or both can secure a draw.

Definition 8 *A finite game in extensive form with perfect information is given by*

1) a tree, with a particular node taken as the origin;

2) for each non-terminal node, a specification of who has the move;

3) for each terminal node, a payoff attached to it.

Formally, a tree is a pair $\Gamma = (N, \sigma)$ where N is the finite set of nodes, and $\sigma : N \to N \cup \emptyset$ associates to each node its predecessor. A (unique) node n_0 with no predecessors (i.e., $\sigma(n_0) = \emptyset$) is the origin of the tree. Terminal nodes are those which are not predecessors of any node. Denote by $T(N)$ the set of terminal nodes. For any non-terminal node r, the set $\{n \in N : \sigma(n) = r\}$ is the set of successors of r. The maximal possible number of edges in a path from the origin to some terminal node is called the length of the tree Γ.

Given a tree Γ, a two-person zero-sum game with perfect information is defined by a partition of N as $N = T(N) \cup N_1 \cup N_2$ into three disjoint sets and a payoff function defined over the set of terminal nodes $u : T(N) \to \mathbf{R}$.

For each non-terminal node n, $n \in N_i$ $(i = 1, 2)$ means that player i has the move at this node. A move consists of picking a successor to this node. The game starts at the origin n_0 of the tree and continues until some terminal node n_t is reached. Then the payoff $u(n_t)$ attached to this node is realized (i.e., player 1 gains $u(n_t)$ and player 2 looses $u(n_t)$).

We do not necessary assume that $n_0 \in N_1$. We even do not assume that if a player i has a move at a node n, then it is his or her opponent who moves at its successor nodes (if the same player has a move at a node and some of its successors, we can *reduce* the game and eliminate this anomaly).

The term "perfect information" refers to the fact that, when a player has to move, he or she is perfectly informed about his or her position in the tree. If chance moves occur later or before this move, their outccome is revealed to every player.

Recall that a *strategy* for player i is a complete specification of what move to choose at each and every node from N_i. We denote their set as S, or T, as above.

Theorem 9 *(Kuhn) Every finite two-person zero-sum game in extensive form with perfect information has a value. Each player has at least one optimal (prudent) strategy in such a game.*

Proof. The proof is by induction in the length l of the tree Γ. For $l = 1$ the theorem holds trivially, since it is a one-person one-move game (say, player 1 is to choose a move at n_0, and any of his moves leads to a terminal node). Thus, a prudent strategy for the player 1 is a move which gives him the highest payoff, and this payoff is the value of the game. Assume now that the theorem holds for all games of length at most $l - 1$, and consider a game G of length l. Without loss of generality, $n_0 \in N_1$, i.e., player 1 has a move at the origin. Let $\{n_1, ..., n_k\}$ be the set of successors of the origin n_0. Each subtree Γ_i, with the origin n_i, is of length $l - 1$ at most. Hence, by the induction hypothesis, any subgame G_i associated with a Γ_i has a value, say, m_i. We claim that the value of the original game G is $m = \max_{1 \leq i \leq k} m_i$. Indeed, by moving first to n_i and then playing optimally at G_i, player 1 can guarantee himself at least m_i. Thus, player 1 can guarantee that he will gain at least m in our game G. But, by playing optimally in each game G_i, player 2 can guarantee herself the loss of not more than m_i. Hence, player 2 can guarantee that she will lose at most m in our game G. Thus max-min and min-max payoffs coincide and m is the value of the game G. ■

The value of a finite two-person zero-sum game in extensive form, as well as optimal strategies for the players, are easily found by solving the game backward. We start by any non-terminal node n, such that all its successors are terminal. An optimal choice for the player i who has a move at n is clearly one which leads to a terminal node with the best payoff for him/her (the max payoff if $i = 1$, or the min payoff if $i = 2$). We can write down this optimal move for the player i at the node n, then delete all subtree which originates at n, except the node n itself, and finally assign to n the best payoff player i can get. Thus, the node n becomes the terminal node of so reduced game tree. After a finite number of such steps, the original game will reduce to one node n_0, and the payoff assigned to it will be the value of the initial game. The optimal strategies of the players are given by their optimal moves at each node, which we wrote down when reducing the game.

Remark 10 *Consider the simple case, where all payoffs are either* $+1$ *or* -1 *(a player either "wins" or "looses"), and where whenever a player has a move* *at some node, his/her opponent is the one who has a move at all its successors. An example is Gale's chomp game above. When we solve this game backward, all payoffs which we attach to non-terminal nodes in this process are* $+1$ *or* -1 *(we can simply write "+" or "−"). Now look at the original game tree with "+" or "−" attached to each its node according to this procedure. A "+" sign at a node* n *means that this node (or "this position") is "winning" <for player 1>, in a sense that if the player 1 would have a move at this node he would surely win, if he would play optimally. A "−" sign at a node* n *means that this node (or "this position") is "loosing" <for player 1>, in a sense that if the player 1 would have a move at this node he would surely lose, if his opponent would play optimally. It is easy to see that "winning" nodes are those which have at least one "loosing" successor, while "loosing" nodes are those whose all successors are "winning". A number of the problems below are about computing the set of winning and losing positions.*

3 Mixed strategies

Penalty kicks in soccer, serves in tennis: in each case the receiver must anticipate the move of the sender to increase her chances of a winning move. So the sender must use an appropriate mixture of shots.

Bluffing in Poker When optimal play involves some bluffing, the bluffing behavior needs to be unpredictable. This can be guaranteed by delegating a choice of when to bluff to some (carefully chosen!) random device. Then even the player herself would not be able to predict in advance when she will be bluffing. So the opponents will certainly not be able to guess whether she is bluffing. See the bluffing game (problem 17) below.

Matching pennies: the matrix $\begin{pmatrix} 1 & -1 \\ -1 & 1 \end{pmatrix}$, has no saddle point. Moreover, for this game $\underline{m} = -1$ and $\overline{m} = 1$ (the worst possible outcomes), i.e., a prudent strategy does not provide any of two players with any minimal guarantee. Here a player's payoff depends completely on how well he or she can predict the choice of the other player. Thus, the best way to play is to be unpredictable, i.e. to choose a strategy (one of the two available) completely *random*. It is easy to see that if each player chooses either strategy with probability 1/2 according to the realization of some random device (and so without any predictable pattern), then "on average" (after playing this game many times) they both will get zero. In other words, under such strategy choice the "expected payoff" for each player will be zero. Moreover, we show below that this randomized strategy is also optimal in the mixed extension of the deterministic game.

Schelling's toy safe. Ann has 2 safes, one at her office which is hard to crack, another "toy" fake at home which any thief can open with a coat-hanger (as in the movies). She must keep her necklace, worth $10,000, eithe at home or at the office. Bob must decide which safe to visit (he has only one visit at only one safe). If he chooses to visit the office, he has a 20% chance of opening the safe. If he goes to ann's home, he is sure to be able to open the safe. The point of this example is that the presence of the toy safe helps Ann, who should actually use it to hide the necklace with a positive probability.

Even when using mixed strategies is clearly warranted, it remains to determine which mixed strategy to choose (how often to bluff, and on what hands?). The player should choose the probabilities of each deterministic choice (i.e. on how she would like to program the random device she uses). Since the player herself cannot predict the actual move she will make during the game, the payoff she will get is uncertain. For example, a player may decide that she will use one strategy with probability 1/3, another one with probability 1/6, and yet another one with probability 1/2. When the time to make her move in the game comes, this player would need some random device to determine her final strategy choice, according to the pre-selected probabilities. In our example, such device should have three outcomes, corresponding to three potential choices, relative chances of these outcomes being 2 : 1 : 3. If this game is played many times, the player should expect that she will play 1-st strategy roughly 1/3 of the time, 2-nd one roughly 1/6 of the time, and 3-d one roughly 1/2 of the time. She will then get "on average" 1/3 (of payoff if using 1-st strategy) +1/6 (of payoff if using 2-nd strategy) +1/2 (of payoff if using 3-d strategy).

Note that, though this player's opponent cannot predict what her actual move would be, he can still evaluate relative chances of each choice, and this will affect his decision. Thus a rational opponent will, in general, react differently to different mixed strategies.

What is the rational behavior of our players when payoffs become uncertain? The simplest and most common hypothesis is that they try to maximize their expected (or average) payoff in the game, i.e., they evaluate random payoffs simply by their expected value. Thus the **cardinal** values of the deterministic payoffs now matter very much, unlike in the previous sections where the **ordinal** ranking of the outcomes is all that matters to the equilibrium analysis. We give in Chapter 2 some axiomatic justifications for this crucial assumption.

The expected payoff is defined as the weighted sum of all possible payoffs in the game, each payoff being multiplied by the probability that this payoff is realized. In matching pennies, when each player chooses a "mixed strategy" $(0.5, 0.5)$ (meaning that 1-st strategy is chosen with probability 0.5, and 2-nd strategy is chosen with probability 0.5), the chances that the game will end up in each particular square (i, j), i.e., the chances that the 1-st player will play his i-th strategy and the 2-nd player will play her j-th strategy, are $0.5 \times 0.5 = 0.25$. So the expected payoff for this game under such strategies is $1 \times 0.25 + (-1) \times 0.25 + 1 \times 0.25 + (-1) \times 0.25 = 0$.

Definition 11 *Consider a general finite game $G = (S, T, u)$, represented by an n by m matrix A, where $n = |S|$, $m = |T|$. The elements of the strategy sets S and T ("sure" strategy choices, which do not involve randomization) are called pure or deterministic strategies. A mixed strategy for the player is a probability distribution over his or her deterministic strategies, i.e. a vector of probabilities for each deterministic strategy which can be chosen during the actual game playing. Thus, the set of all mixed strategies for player 1 is $X =$*

$\{(x_1,...,x_n) : \sum_{i=1}^n x_i = 1, x_i \geq 0\}$, *while for player 2 it is* $Y = \{(y_1,...,y_m) : \sum_{i=1}^m y_j = 1, y_j \geq 0\}$.

Note that when player 1 chooses $x \in X$ and player 2 chooses $y \in Y$, the expected payoff of the game is equal to the matrix product $x^T Ay$:

$$x^T Ay = (x_1,...,x_n) \begin{pmatrix} a_{11} & ... & a_{1m} \\ ... & ... & ... \\ a_{n1} & ... & a_{nm} \end{pmatrix} \begin{pmatrix} y_1 \\ ... \\ y_m \end{pmatrix} = \sum_{i=1}^n \sum_{j=1}^m x_i a_{ij} y_j,$$

and each element of this double sum is $x_i a_{ij} y_j = a_{ij} x_i y_j =$
$a_{ij} \times \text{Pro}[1 \text{ chooses } i] \times \text{Pro}[2 \text{ chooses } j] = a_{ij} \times \text{Pro}[1 \text{ chooses } i \text{ and } 2 \text{ chooses } j]$.

The number $x^T Ay$ is a weighted average of the expected payoffs for player 1 when he uses x against player's 2 pure strategies (where weights are probabilities that player 2 will use these pure strategies).

$$\begin{aligned} x^T Ay &= x^T \begin{pmatrix} a_{11} & ... & a_{1m} \\ ... & ... & ... \\ a_{n1} & ... & a_{nm} \end{pmatrix} \begin{pmatrix} y_1 \\ ... \\ y_m \end{pmatrix} = x^T \left[y_1 A_{\cdot 1} + ... + y_m A_{\cdot m} \right] = \\ &= y_1 \left[x^T A_{\cdot 1} \right] + ... + y_m \left[x^T A_{\cdot m} \right] = y_1 \left[x^T A e^1 \right] + ... + y_m \left[x^T A e^m \right]. \end{aligned}$$

Here $A_{\cdot j}$ is j-th column of the matrix A, and $e^j = (0,...,0,1,0,...,0)$ is the (m-dimensional) vector, whose all coordinates are zero, except that its j-th coordinate is 1, which represents the pure strategy j of player 2. Recall $A_{\cdot j} = Ae^j$.

We define the secure utility level for player 1<2> (the minimal gain he can guarantee himself, no matter what player 2<1> does) in the same spirit as before. The only change is that it is now the "expected" utility level, and that the strategy sets available to the players are much bigger now: X and Y, instead of S and T.

Let $v_1(x) = \min_{y \in Y} x^T Ay$ be the minimum payoff player 1 can get if he chooses to play x. Then $v_1 = \max_{x \in X} v_1(x) = \max_{x \in X} \min_{y \in Y} x^T Ay$ is the secure utility level for player 1. Similarly, we define $v_2(y) = \max_{x \in X} x^T Ay$, and $v_2 = \min_{y \in Y} v_2(y) = \min_{y \in Y} \max_{x \in X} x^T Ay$, the secure utility level for player 2.

Given the above decomposition of $x^T Ay$, and $v_1(x) = \min_{y \in Y} x^T Ay$, the minimum of $x^T Ay$, will be attained at some pure strategy j (i.e., at some $e^j \in Y$). Indeed, if $x^T Ae^j > v_1(x)$ for all j, then we would have $x^T Ay = \sum y_j \left[x^T Ae^j \right] > v_1(x)$ for all $y \in Y$. Hence, $v_1(x) = \min_j x^T A_{\cdot j}$, and $v_1 = \max_{x \in X} \min_j x^T A_{\cdot j}$. Similarly, $v_2(y) = \max_i A_{i \cdot} y$, where $A_{i \cdot}$ is the i-th row of the matrix A, and $v_2 = \min_{y \in Y} \max_i A_{i \cdot} y$.

As with pure strategies, the secure utility level player 1 can guarantee himself (minimal amount he could gain) cannot exceed the secure utility level payer 2 can guarantee herself (maximal amount she could lose): $v_1 \leq v_2$. This follows from Lemma 1.

Such prudent mixed strategies $\overline{x}$ and $\overline{y}$ are called maximin strategy (for player 1) and minimax strategy (for player 2) respectively.

Theorem 12 *(The Minimax Theorem) $v_1 = v_2 = v$. Thus, if players can use mixed strategies, any game with finite strategy sets has a value.*

Proof. Let $n \times m$ matrix A be the matrix of a two person zero sum game. The set of all mixed strategies for player 1 is $X = \{(x_1, ..., x_n) : \sum_{i=1}^{n} x_i = 1, x_i \geq 0\}$, while for player 2 it is $Y = \{(y_1, ..., y_m) : \sum_{i=1}^{m} y_j = 1, y_j \geq 0\}$. Let $v_1(x) = \min\limits_{y \in Y} x \cdot Ay$ be the smallest payoff player 1 can get if he chooses to play x. Then $v_1 = \max\limits_{x \in X} v_1(x) = \max\limits_{x \in X} \min\limits_{y \in Y} x \cdot Ay$ is the secure utility level for player 1. Similarly, we define $v_2(y) = \max\limits_{x \in X} x \cdot Ay$, and $v_2 = \min\limits_{y \in Y} v_2(y) = \min\limits_{y \in Y} \max\limits_{x \in X} x \cdot Ay$ is the secure utility level for player 2. We know that $v_1 \leq v_2$.

Consider the following closed convex sets in $\mathbb{R}^n$:

$L = \{z \in \mathbb{R}^n : z = Ay \text{ for some } y \in Y\}$ is a convex set, since $Ay = y_1 A_{\cdot 1} + ... + y_m A_{\cdot m}$, where $A_{\cdot j}$ is j-th column of the matrix A, and hence L is the set of all convex combinations of columns of A, i.e., the convex hull of the columns of A. Moreover, since it is a convex hull of m points, L is a convex polytope in $\mathbb{R}^n$ with m vertices (extreme points), and thus it is also closed and bounded.

Cones $K_v = \{z \in \mathbb{R}^n : z_i \leq v \text{ for all } i = 1, ..., n\}$ are obviously convex and closed for any $v \in \mathbb{R}$. Further, it is easy to see that $K_v = \{z \in \mathbb{R}^n : x \cdot z \leq v \text{ for all } x \in X\}$.

Geometrically, when v is very small, the cone K_v lies far from the bounded set L, and they do not intersect. Thus, they can be separated by a hyperplane. When v increases, the cone K_v enlarges in the direction $(1, ..., 1)$, being "below" the set L, until the moment when K_v will "touch" the set L for the first time. Hence, $\overline{v}$, the maximal value of v for which K_v still can be separated from L, is reached when the cone $K_{\overline{v}}$ first "touches" the set L. Moreover, $K_{\overline{v}}$ and L have at least one common point $\overline{z}$, at which they "touch". Let $\overline{y} \in Y$ be such that $A\overline{y} = \overline{z} \in L \cap K_{\overline{v}}$.

Assume that $K_{\overline{v}}$ and L are separated by a hyperplane $H = \{z \in \mathbb{R}^n : \overline{x} \cdot z = c\}$, where $\sum_{i=1}^{n} \overline{x}_i = 1$. It means that $\overline{x} \cdot z \leq c$ for all $z \in K_{\overline{v}}$, $\overline{x} \cdot z \geq c$ for all $z \in L$, and hence $\overline{x} \cdot \overline{z} = c$. Geometrically, since $K_{\overline{v}}$ lies "below" the hyperplane H, all coordinates $\overline{x}_i$ of the vector $\overline{x}$ must be nonnegative, and thus $\overline{x} \in X$. Moreover, since $K_{\overline{v}} = \{z \in \mathbb{R}^n : x \cdot z \leq \overline{v} \text{ for all } x \in X\}$, $\overline{x} \in X$ and $\overline{z} \in K_{\overline{v}}$, we obtain that $c = \overline{x} \cdot \overline{z} \leq \overline{v}$. But since vector $(\overline{v}, ..., \overline{v}) \in K_{\overline{v}}$ we also obtain that $c \geq \overline{x} \cdot (\overline{v}, ..., \overline{v}) = \overline{v} \sum_{i=1}^{n} \overline{x}_i = \overline{v}$. It follows that $c = \overline{v}$.

Now, $v_1 = \max\limits_{x \in X} \min\limits_{y \in Y} x \cdot Ay \geq \min\limits_{y \in Y} \overline{x} \cdot Ay \geq \overline{v}$ (since $\overline{x} \cdot z \geq c = \overline{v}$ for all $z \in L$, i.e. for all $z = Ay$, where $y \in Y$). Next, $v_2 = \min\limits_{y \in Y} \max\limits_{x \in X} x \cdot Ay \leq \max\limits_{x \in X} x \cdot A\overline{y} = \max\limits_{x \in X} x \cdot \overline{z} = \max\limits_{i=1,\cdots,n} \overline{z_i} \leq \overline{v}$ (since $\overline{z} \in K_{\overline{v}}$).

We obtain that $v_2 \leq \overline{v} \leq v_1$. Together with the fact that $v_1 \leq v_2$, it gives us $v_2 = \overline{v} = v_1$, the desired statement. Note also, that the maximal value of $v_1(x)$ is reached at $\overline{x}$, while the minimal value of $v_2(y)$ is reached at $\overline{y}$. Thus, $\overline{x}$ and $\overline{y}$

constructed in the proof are optimal strategies for players 1 and 2 respectively. ■

4 Computation of optimal mixed strategies

How can we find the (a) maximin strategy $\overline{x}$, the (a) minimax strategy $\overline{y}$, and the value v of a given game?

If the game with deterministic strategies (the original game) has a saddle point, then $v = m$, and the maximin and minimax strategies are deterministic. Finding them amounts to find an entry a_{ij} of the matrix A which is both the maximum entry in its column and the minimum entry in its row.

When the original game has no value, the key to computing optimal mixed strategies is to know their *supports,* namely the set of strategies used with strictly positive probability. Let $\overline{x}, \overline{y}$ be a pair of optimal strategies, and $v = \overline{x}^T A \overline{y}$. Since for all j we have that $\overline{x}^T A e^j \geq \min_{y \in Y} \overline{x}^T A y = v_1(\overline{x}) = v_1 = v$, it follows that $v = \overline{x}^T A \overline{y} = \overline{y}_1 \left[\overline{x}^T A e^1\right] + ... + \overline{y}_m \left[\overline{x}^T A e^m\right] \geq \overline{y}_1 v + ... + \overline{y}_m v = v\left(\overline{y}_1 + ... + \overline{y}_m\right) = v$, and the equality implies $\overline{x}^T A_{\cdot j} = \overline{x}^T A e^j = v$ for all j such that $\overline{y}_j \neq 0$. Thus, player 2 receives her minimax value $v_2 = v$ by playing against $\overline{x}$ any pure strategy j which is used with a positive probability in her minimax strategy $\overline{y}$ (i.e. any strategy j, such that $\overline{y}_j \neq 0$).

Similarly, player 1 receives his maximin value $v_1 = v$ by playing against $\overline{y}$ any pure strategy i which is used with a positive probability in his maximin strategy $\overline{x}$ (i.e. any strategy i, such that $\overline{x}_i \neq 0$). Setting $x^* = \{i | \overline{x}_i > 0\}$ and $T^* = \{j | \overline{y}_j > 0\}$, we see that $\overline{x}, \overline{y}$ solve the following system with unknown x, y

$$x^T A_{\cdot j} = v \text{ for all } j \in T^*; A_{i \cdot} y = v \text{ for all } i \in S^*$$

$$\sum_{i=1}^{n} x_i = 1, x_i \geq 0, \sum_{i=1}^{m} y_j = 1, y_j \geq 0$$

Finding a solution $\overline{x}, \overline{y}$ to this system for a given pair $S^* T^*$ is not enough to ensure that $(\overline{x}, \overline{y})$ is actually a mixed strategy equilibrium ofthe initial game. We must also check that $\overline{x}$ (resp. $\overline{y}$) is also a best reply against the strategy $\overline{y}$ (resp. $\overline{x}$) of the opponent, i.e., that the following system of inequalities hold

$$A_{i \cdot} \overline{y} \leq v \text{ for all } i \in S \diagdown S^*; \quad \overline{x}^T A_{\cdot j} \geq v \text{ for all } j \in T \diagdown T^*$$

The main difficulty in this approach is that there are 2^{n+m} possible choices for S^*, T^*, and no systematic way to guess!

One way to limit the search for S^*, T^* is to assume that they are of the same size, because if they are not, one of the two systems $x^T A_{\cdot j} = v$ for all $j \in T^*$, or $A_{i \cdot} y = v$ for all $i \in S^*$, is *over-determined,* so that the corresponding payoff sub-matrix is not of maximal rank. But this is not going to always work, as shown in the following example:

$$\begin{bmatrix} 0 & 0 \\ 1 & -2 \\ -2 & 1 \end{bmatrix}$$

where the unique saddle point has the row player using Top with probability 1, while the Column player randomizes with equal probabilities over the two pure strategies.

Here is another useful but not systematic approach. In many $n \times n$ games (each player has n pure strategies), one can get an idea about the support of an optimal pair by assuming a full support and solving the corresponding system of equalities (as above, except for $x_i \geq 0$ and $y_j \geq 0$). If its solution is non negative, it is a pair of optimal strategies. If not, the set of pure strategies i, j where $x_i \geq 0$ and $y_j \geq 0$ gives plausible bounds of the support of an optimal strategy. But this trick is not always going to work. Consider the 3×3 game with payoffs

$$\begin{bmatrix} 8 & 0 & -1 \\ -3 & 4 & 7 \\ 0 & -2 & 0 \end{bmatrix}$$

where the trick suggests to give zero weight to the middle column, when in fact the optimal strategy puts weight on the left and middle columns (and on the top and middle rows).

A more rigorous approach to simplify the search for the supports of optimal mixed strategies uses the successively elimination of *dominated* rows and columns.

Definition 13 *We say that the i-th row of a matrix A dominates (resp. strictly dominates) its k-th row, if $a_{ij} \geq a_{kj}$ for all j and $a_{ij} > a_{kj}$ for at least one j (resp. $a_{ij} > a_{kj}$ for all j). Similarly, we say that the j-th column of a matrix A dominates (resp. strictly dominates) its l-th column, if $a_{ij} \geq a_{il}$ for all i and $a_{ij} > a_{il}$ for at least one i (resp. $a_{ij} > a_{il}$ for all i).*

In other words, a pure strategy (represented by a row or a column of A) dominates another pure strategy if, no matter what the other player does, the choice of the first (dominating) strategy is at least as good as the choice of the second (dominated) strategy, and for some strategies of the other player, it is strictly better.

Strict domination requires more: the first strategy yields a strictly better payoff than the second one, irrespective of the strategic choices of other players.

Proposition 14 *a) If the row i of a matrix A is strictly dominated, then any optimal strategy $\overline{x}$ of player 1 has $\overline{x}_i = 0$. Thus removing a strictly dominated strategy does not lose any optimal strategy for either player.*

b) If the row i of a matrix A is dominated, then player 1 has an optimal strategy $\overline{x}$ such that $\overline{x}_i = 0$. Moreover, any optimal strategy, for any player, in the game obtained by removing dominated rows from A will also be an optimal strategy in the original game. Thus a player can always find an optimal mixed strategy using only undominated strategies, but some of his optimal strategies may be dominated.

c) Two symmetrical statements hold for strictly dominated and dominated columns of player 2.

Removing dominated rows of A gives a smaller matrix A_1. Removing dominated columns of A_1 leaves us with a yet smaller matrix A_2. We can continue by removing dominated rows of A_2, etc., until we obtain a matrix which does not contain dominated rows or columns. The optimal strategies and the value for the game with this reduced matrix will still be the optimal strategies and the value for the initial game represented by A. This process is called "iterative elimination of dominated strategies". See the problems for examples of application of this technique.

An example where each player has one dominated pure optimal strategy (and infinitely many mixed optimal strategies) follows

$$\begin{bmatrix} 0 & 0 & 1 \\ 0 & 0 & 0 \\ -1 & 0 & \cdot \end{bmatrix}$$

4.1 2×2 games

Suppose that $A = \begin{pmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{pmatrix}$. This game does not have saddle point if and only if $[a_{11}, a_{22}] \cap [a_{12}, a_{21}] = \varnothing$. In this case, a pure strategy cannot be optimal for either player (check it!). It follows that optimal strategies (x_1, x_2) and (y_1, y_2) must have all components positive. Let us repeat the argument above for the 2×2 case. We have $v = x^T Ay = a_{11}x_1y_1 + a_{12}x_1y_2 + a_{21}x_2y_1 + a_{22}x_2y_2$, or

$$x_1(a_{11}y_1 + a_{12}y_2) + x_2(a_{21}y_1 + a_{22}y_2) = v.$$

But $a_{11}y_1 + a_{12}y_2 \leq v$ and $a_{21}y_1 + a_{22}y_2 \leq v$ (these are the losses of player 2 against 1-st and 2-nd pure strategies of player 1; but since y is player's 2 optimal strategy, she cannot lose more then v in any case). Hence, $x_1(a_{11}y_1 + a_{12}y_2) + x_2(a_{21}y_1 + a_{22}y_2) \leq x_1v + x_2v = v$. Since $x_1 > 0$ and $x_2 > 0$, the equality is only possible when $a_{11}y_1 + a_{12}y_2 = v$ and $a_{21}y_1 + a_{22}y_2 = v$. Similarly $a_{11}x_1 + a_{21}x_2 = v$ and $a_{12}x_1 + a_{22}x_2 = v$. We also know that $x_1 + x_2 = 1$ and $y_1 + y_2 = 1$.

We have a linear system with 6 equations and 5 variables x_1, x_2, y_1, y_2 and v. The minimax theorem guarantees us that this system has a solution with $x_1, x_2, y_1, y_2 \geq 0$. One of these 6 equations is actually redundant. The system has a unique solution provided the original game has no saddle point. In particular

$$v = \frac{a_{11}a_{22} - a_{12}a_{21}}{a_{11} + a_{22} - a_{12} - a_{21}}$$

Note that the denominatior is non zero because $[a_{11}, a_{22}] \cap [a_{12}, a_{21}] = \varnothing$.

4.2 2×n games

By focusing on the player who has two strategies, one computes the value as the solution of a tractable linear program. Say player 1 has two strategies and the game

$A = \begin{pmatrix} a & b & c & \cdots & g \\ a' & b' & c' & \cdots & g' \end{pmatrix}$ has no pure strategy equilibrium (no value). Then the optimal mixed strategy of player 1 is $x = (p, 1-p)$, and it solves

$$\max_{0 \le p \le 1} \min\{ap + a'(1-p), bp + b'(1-p), \cdots, gp + g'(1-p)\}$$

because the best reply of player 2 can be chosen a pure strategy. It is easy to represent the function $\min\{ap + a'(1-p), bp + b'(1-p), \cdots, gp + g'(1-p)\}$ as a minimum of straight lines, and to find its maximum.

See the examples in Problem 9.

4.3 Symmetric games

The game with matrix A is symmetric if $A = -A^T$ (Exercise: check this). Recall that the value of a symmetric game is zero (Lemma 7). Moreover, if x is an optimal strategy for player 1, then it is also optimal for player 2.

5 Von Neumann's Theorem

It generalizes the minimax theorem. And it is a special case of the more general Nash Theorem in Chapter 4.

Theorem 15 *The game (S, T, u) has a value and optimal strategies if S, T are convex compact subsets of some euclidian spaces, the payoff function u is continuous on $S \times T$, and for all $s \in S$, all $t \in T$*

$$t' \to u(s, t') \text{ is quasi-convex in } t'; \; s' \to u(s', t) \text{ is quasi-concave in } s'$$

Example 8 *Borel's model of poker.*
Each player bids \$1, then receives a hand $m_i \in [0,1]$. Hands are independently and uniformly distributed on $[0,1]$.Each player observes only his hand.Player 1 moves first, by either folding or bidding an additional \$5. If 1 folds, the game is over and player 2 collects the pot. If 1 bids, player 2 can either fold (in which case 1 collects the pot) or bid \$5 more to see: then the hands are revealed and the highest one wins the pot.

A strategy of player i can be any mapping from $[0,1]$ into $\{F, B\}$, however it is enough to consider the following simple *threshold* strategies s_i : fold whenever $m_i \le s_i$, bid whenever $m_i > s_i$. Notice that for player 2, actual bidding only occur if player 1 bids before him. Compute the probability $\pi(s_1, s_2)$ that $m_1 > m_2$ given that $s_i \le m_i \le 1$:

$$\pi(s_1, s_2) = \frac{1 + s_1 - 2s_2}{2(1 - s_2)} \text{ if } s_2 \le s_1$$

$$= \frac{1 - s_2}{2(1 - s_1)} \text{ if } s_1 \le s_2$$

from which the payoff function is easily derived:

$$u(s_1, s_2) = -6s_1^2 + 5s_1s_2 + 5s_1 - 5s_2 \text{ if } s_2 \leq s_1$$
$$= 6s_2^2 - 7s_1s_2 + 5s_1 - 5s_2 \text{ if } s_1 \leq s_2$$

The Von Neumann theorem applies, and the utility function is continuously differentiable. Thus the saddle point can be found by solving the system $\frac{\partial u}{\partial s_i}(s) = 0, i = 1, 2$. This leads to

$$s_1^* = (\frac{5}{7})^2 = 0.51; s_2^* = \frac{5}{7} = 0.71$$

and the value -0.51: player 2 earns on average 51 cents.

Two more simplistic models of poker are in the problems below.

6 infinite games

When the sets of pure strategies are infinite, mixed strategies can still be defined as probability distributions over these sets, but the existence of a value for the game in mixed strategies is no longer guaranteed.

Example 5: *a silly game*
Each player chooses an integer in $\{1, 2, \cdots, n, \cdots\}$. The one who chooooses the largest integer wins \$1 from the other, unless they choose the same number, in which case no money changes hands. A mixed strategy is a probability distribution $x = (x_1, x_2, \cdots, x_n, \cdots), x_i \geq 0, \sum_1^\infty x_i = 1$. Given any such strategy chosen by the opponent, and any positive ε, there exists n such that $\sum_n^\infty x_i \leq \varepsilon$, therefore playing n guarantees a win with probability no less than $1 - \varepsilon$. It follows that in the game in mixed strategies, $\max_{x \in X} \min_{y \in Y} u(x, y) = -1 < +1 = \min_{y \in Y} \max_{x \in X} u(x, y)$.

Theorem 16 *(Glicksberg Theorem). If the sets of pure strategies S, T are convex compact subsets of some euclidian space, and the payoff function u is continuous on $S \times T$, then the game in mixed strategies (where each player uses a probability distribution over pure strategies) has a value.*

However, knowing that a value exists does not help much to identify optimal mixed strategies, because the support of these mixed strategies can now vary in a very large set!

An example where Glicksberg Theorem applies is the subject of Problem 13.2.

A typical case where Glicksberg Theorem does *not* apply is when S, T are convex compacts, yet the payoff function u is discontinuous. Below are two such examples: in the first one the game nevertheless has a value and optimal strategies, in the second it does not.

Example 6 *Mixed strategies in the silent gunfight*

In the silent gunfight (Problem 5; see also the noisy version Example 2 in section 1.2), we assume $a(t) = b(t) = t$, so that the game is symmetric, and its value (if it exists) is 0. The payoff function is

$$\begin{aligned} u(s,t) &= s - t(1-s) \text{ if } s < t \\ u(s,t) &= -t + s(1-t) \text{ if } t < s \\ u(s,t) &= 0 \text{ if } s = t \end{aligned}$$

It is enough to look for a symmetric equilibrium. Note that shooting near $s = 0$ makes no sense, as it guarantees a negative payoff to player 1. In fact the best reply of player 1 to the strategy t by player 2 is $s = 1$ if $t < \sqrt{2} - 1$, $s = t - \varepsilon$ if $t > \sqrt{2} - 1$.

This suggests that the support of an optimal mixed strategy will be $[a, 1]$, for some $a \geq 0$, and that the optimal strategy has a density $f(t)$ over $[a, 1]$. We compute player 1's expected payoff from the *pure* strategy $s, a \leq s \leq 1$, against the strategy f by player 2

$$\overline{u}(s,f) = \int_a^s (s(1-t) - t)f(t)dt + \int_s^1 (s(1+t) - t)f(t)dt$$

The equilibrium condition is that $\overline{u}(s,f) = 0$ for all $s \in [a, 1]$. This equality is rearranged as

$$s - (1+s)\{\int_a^s tf(t)dt\} - (1-s)\{\int_s^1 tf(t)dt\} = 0$$

Setting $H(s) = \int_s^1 tf(t)dt$, this writes

$$s = (1+s)(H(a) - H(s)) + (1-s)H(s) \Leftrightarrow H(s) = H(a)\frac{1+s}{2s} - \frac{1}{2}$$

Taking $H(1) = 0$ into account gives $H(a) = \frac{1}{2}$, then

$$H(s) = \frac{1-s}{4s} \Rightarrow f(s) = \frac{1}{4s^3}$$

Finally we find a from

$$1 = \int_a^1 f(t)dt \Rightarrow a = \frac{1}{3}$$

Example 7 *Campaign funding*

Each player divides his \$1 campaign budget between two states A and B. The challenger (player 1) wins the overall game (for a payoff \$1) if he wins (strictly) in one state, where the winner in state A is whomever spends the most money, but in state B the incumbent (player 2) has an advantage of \$0.5 so the challenger only wins if his budget there exceeds that of the incumbent by more than \$0.5. Here is the normal form of the game:

$S = T = [0, 1]$ s (resp. t) is spent by player 1 (resp. 2) in state A

$$
\begin{aligned}
u(s,t) &= +1 \text{ if } t < s \text{ or } s + \frac{1}{2} < t \\
u(s,t) &= -1 \text{ if } s < t < s + \frac{1}{2} \\
u(s,t) &= 0 \text{ if } s = t \text{ or } s + \frac{1}{2} = t
\end{aligned}
$$

Clearly in the pure strategy game $\max_s \min_t u(s,t) = -1 < +1 = \min_t \max_s u(s,t)$. We claim that in the mixed strategy game we have

$$
\max_{x \in X} \min_{y \in Y} u(x,y) = \frac{1}{3} < \frac{3}{7} = \min_{y \in Y} \max_{x \in X} u(x,y) \tag{1}
$$

Suppose first that player 2's mixed strategy y guarantees

$$
\sup_{s \in [0,1]} \overline{u}(s,y) < \frac{3}{7} \tag{2}
$$

Applying (2) at $s = 1$ gives $y(1) > \frac{4}{7}$, and at $s = 0$

$$
y(]\frac{1}{2}, 1]) - y(]0, \frac{1}{2}[) < \frac{3}{7} \tag{3}
$$

Applying (2) at $s = \frac{1}{2} - \varepsilon$, and letting ε go to zero, gives

$$
y([0, \frac{1}{2}[) + y(1) - y([\frac{1}{2}, 1[) \leq \frac{3}{7}
$$

Summing the latter two inequalities yields

$$
2y(1) + y(0) - y(\frac{1}{2}) \leq \frac{6}{7}
$$

Combined with $y(1) > \frac{4}{7}$, this implies $y(\frac{1}{2}) \geq \frac{2}{7}$, and (3) gives similarly $y(]0, \frac{1}{2}[) > \frac{1}{7}$. This is a contradiction as $y(1) + y(\frac{1}{2}) + y(]0, \frac{1}{2}[) \leq 1$, hence inequality (2) is after all impossible.

Next one checks easily that player 2's strategy

$$
y^* = \frac{1}{7}\delta_{\frac{1}{4}} + \frac{2}{7}\delta_{\frac{1}{2}} + \frac{4}{7}\delta_1
$$

guarantees $\sup_{[0,1]} \overline{u}(s, y^*) = \frac{3}{7}$.

To prove the other half of property (1), we assume the mixed strategy x is such that

$$
\inf_{t \in [0,1]} \overline{u}(x,t) > \frac{1}{3}
$$

and apply this successively to $t = 1$ and $t = \frac{1}{2} - \varepsilon$, letting ε go to zero. We get

$$
x([0, \frac{1}{2}[) - x(]\frac{1}{2}, 1[) > \frac{1}{3} \text{ and } -x([0, \frac{1}{2}[) + x([\frac{1}{2}, 1]) \geq \frac{1}{3}
$$

Summming these two inequalities $x(\frac{1}{2})+x(1) > \frac{2}{3}$, a contradiction of $x([0,\frac{1}{2}[) > \frac{1}{3}$. Finally player 1's strategy

$$x^* = \frac{1}{3}\delta_0 + \frac{1}{3}\delta_{\frac{1}{2}} + \frac{1}{3}\delta_1$$

guarantees $\inf_{[0,1]} \overline{u}(x^*,t) = \frac{1}{3}$.

7 Problems on Chapter 2

7.1 Pure strategies

Problem 1

Ten thousands students formed a square. In each row, the tallest student is chosen and Mary is the shortest one among those. In each column, a shortest student is chosen, and John is the tallest one among those. Who is taller—John or Mary?

Problem 2

Compute $\overline{m} = \min\max$ and $\underline{m} = \max\min$ values for the following matrices:

$$\begin{array}{cccc} 2 & 4 & 6 & 3 \\ 6 & 2 & 4 & 3 \\ 4 & 6 & 2 & 3 \end{array} \qquad \begin{array}{cccc} 3 & 2 & 2 & 1 \\ 2 & 3 & 2 & 1 \\ 2 & 2 & 3 & 1 \end{array}$$

Find all saddle points.

Problem 3. *Gale's roulette*

a)Each wheel has an equal probability to stop on any of its numbers. Player 1 chooses a wheel and spins it. Player 2 chooses one of the 2 remaining wheels (while the wheel chosen by 1 is still spinning), and spins it. The winner is the player whose wheel stops on the higher score. He gets \$1 from the loser.

Numbers on wheel #1: 2,4,9; on wheel #2: 3,5,7; on wheel #3: 1,6,8

Find the value and optimal strategies of this game

b) Variant: the winner with a score of s gets \$$s$ from the loser.

Problem 4 *Land division game.*

The land consists of 3 contiguous pieces: the unit square with corners $(0,0),(1,0),(0,1),(1,1)$, the triangle with corners $(0,1),(1,1),(0,2)$, the triangle with corners $(1,0),(1,1),(2,1)$. Player 1 chooses a vertical line L with 1st coordinate in $[0,1]$. Player 2 chooses an horizontal line M with 2d coordinate in $[0,1]$. Then player 1 gets all the land above M and to the left of L, as well as the land below M and to the right of L. Player 2 gets the rest. Both players want to maximize the area of their land. Find the value and optimal strategies.

Problem 5 *Silent gunfight*

Now the duellists cannot hear when the other player shoots. Payoffs are computed in the same way. If v is the value of the *noisy* gunfight, show that in the silent version, the values $\overline{m} = \min\max$ and $\underline{m} = \max\min$ are such that $\underline{m} < v < \overline{m}$.

Problem 6.1

Two players move in turn and the one who cannot move loses. Find the winner (1-st or 2-nd player) and the winning strategy.

In questions a) and b), both players move the same piece.

a) A castle stays on the square a1 of the 8×8 chess board. A move consists in moving the castle according to the chess rules, but only in the directions up or to the right.

b) The same game, but with a knight instead of a castle.

In questions c) and d), a move consists of adding a new piece on the board.

c) A move consists in placing a castle on the 8 by 8 chess board in such a way, that it does not threatens any of the castles already present.

d) The same game, but bishops are to be placed instead of castles.

Problem 6.2

Dominos can be placed on a $m \times n$ board so as to cover two squares exactly. Two players alternate placing dominos. The first one who is unable to place a domino is the loser.

a) Show that one of the two players, First or Second Mover, can guarantee a win.

b) Who wins in the following cases:

$n = 3, m = 3$

$n = 4, m = 4$

c) Who wins in the following cases:

n and m even

n even, m odd

d) *(much harder)* Who wins if $n = 1$? If n and m are odd?

Problem 6.3

Two players move in turn until one of them cannot move. In the *standard* version, that player loses; in the *miser* version, whoever was the last mover loses. Find the winner (1-st or 2-nd moverer) and the winning strategy in both standard and miser versions for the following games.

a) From a pile of n coins, the players take turns to remove *one* or *two* coins. Show that n is a losing position *iff* $n = 0(3)$ in the standard version, *iff* $n = 1(3)$ in the miser version.

b) Same as in a), but now the players can remove *one* or *four* coins?

c) Same as in a), but now the players can remove *one, three* or *five* coins?

d) We now have two piles, of size n and m, and the players take turns to remove *one* or *two* coins from one of the piles. Show that n, m is losing in the standard version *iff* $n = m(3)$, *iff* $n \neq m(3)$ in the miser version.

e) From one of the two piles as in d), the players can remove *one* or *four* coins.

f) We still have two piles of size n, m, but now the players can remove *any number of coins (and at least one)* from one of the piles.

g) Marienbad game: we have p piles of sizes $n_1, \cdots, n_p$. A player can remove any number of coins (and at least one) from one of the (non empty) piles. Show that in the standard version, a position $n_1, \cdots, n_p$ is winning *iff*

$$\text{for all } t, 1 \leq t \leq T : \sum_{k=1}^{p} a_k^t \text{ is even; and } \sum_{k=1}^{p} a_k^t > 0 \text{ for at least one t}$$

when $n_k = a_k^T a_k^{T-1} \cdots a_k^t \cdots a_k^1$ is the diadic representation of n_k, augmented by enough zeros on the left so that all n_k have the same number of digits. What is the solution of the miser version of this game?

Problem 6.4

a) The game starts with two piles, of respectively n and m coins. A move consists in taking one pile away and dividing the other into two nonempty piles. Solve the standard and miser versions of the game (defined in Problem 6.3).

b) n coins are placed on a line such that they touch each other. A move consists in taking either one coin, or two *adjacent* (touching) coins. Solve the standard and miser versions.

c) The initial position is 1111111011111110111101, where a 1 is a match and 0 an empty space. Players successively remove one match or three adjacent matches. Solve the two versions of the game.

Problem 7

Show that, if a 2×3 matrix has a saddle point, then either one row dominates another, or one column dominates another (or possibly both). Show by a counter-example that this is not true for 3×3 matrices.

Problem 8 *Shapley's criterion*

Consider a game (S, T, u) with finite strategy sets such that for every subsets $S_0 \subset S, T_0 \subset T$ with 2 elements each, the 2×2 game (S_0, T_0, u) has a value. Show that the original game has a value. *Hint: by contradiction. Assume* $\max\min < \min\max$*, and without loss* $\max\min < 0 < \min\max$*. Then find a sub-2x2 matrix of the type* $\begin{bmatrix} + & - \\ - & + \end{bmatrix}$.

7.2 Mixed strategies

Problem 9

In each question you must check that the game in deterministic strategies (given in the matrix form) has no value, then find the value and optimal mixed strategies. Results in section 1.5 will prove useful.

a) $A = \begin{pmatrix} 2 & 3 & 1 & 5 \\ 4 & 1 & 6 & 0 \end{pmatrix}$

b) $A = \begin{pmatrix} 12 & 0 \\ 0 & 12 \\ 10 & 6 \\ 8 & 10 \\ 9 & 7 \end{pmatrix}$

c) $A = \begin{pmatrix} 2 & 0 & 1 & 4 \\ 1 & 2 & 5 & 3 \\ 4 & 1 & 3 & 2 \end{pmatrix}$

d) $A = \begin{pmatrix} 1 & 6 & 0 \\ 2 & 0 & 3 \\ 3 & 2 & 4 \end{pmatrix}$

e) $A = \begin{pmatrix} 0 & 1 & -2 \\ -1 & 0 & 3 \\ 2 & -3 & 0 \end{pmatrix}$

f) $A = \begin{pmatrix} 8 & 4 & 2 \\ 0 & 2 & 4 \\ 0 & 4 & 2 \end{pmatrix}, A = \begin{pmatrix} 5 & 4 & 2 \\ 0 & 2 & 4 \\ 0 & 4 & 2 \end{pmatrix}$

g) $A = \begin{pmatrix} 2 & 4 & 6 & 3 \\ 6 & 2 & 4 & 3 \\ 4 & 6 & 2 & 3 \end{pmatrix}$

Problem 10 *Rock, Paper, Scissors and Well*
Two players choose simultaneously one of 4 pure strategies: Rock, Paper, Scissors and Well. If their choices are identical, no money changes hands. Otherwise the loser pays \$1 to the winner.

The pattern of wins and losses is as follows. The paper is cut by (loses to) the scissors, it wraps (beats) the rock and closes (beats) the well. The scissors break on the rock and fall into the well (lose to both). The rock falls into (loses to) the well . The same choice by both players is a tie (no money changes hand).
a) Solve the game in mixed strategies when the winner gets \$1 from the loser.
b) Solve the game in mixed strategies when losing to the rock or the scissor costs \$2 to the loser, while losing to paper or well only costs \$1.

Problem 11 *Picking an entry*
a) Player 1 chooses either a row or a column of the matrix $\begin{bmatrix} 2 & 1 \\ 4 & 5 \end{bmatrix}$. Player 2 chooses an entry of this matrix. If the entry chosen by 2 is in the row or column chosen by 1, player 1 receives the amount of this entry from player 2. Otherwise no money changes hands. Find the value and optimal strategies.
b) Same strategies but this time if player 2 chooses entry s and this entry is not in the row or column chosen by 1, player 2 gets \$$s$ from player 1; if it is in the row or column chosen by 1, player 1 gets \$$s$ from player 2 as before.

Problem 12 *Guessing a number*
Player 2 chooses one of the three numbers 1,2 or 5. Call s_2 that choice. One of the two numbers not selected by Player 2 is selected at random (equal probability 1/2 for each) and shown to Player 1. Player 1 now guesses Player 2's choice: if his guess is correct, he receives \$$s_2$ form Player 2, otherwise no money changes hand.
Solve this game: value and optimal strategies.
Hint: drawing the full normal form of this game is cumbersome; describe instead the strategy of player 1 by three numbers q_1, q_2, q_5. The number q_1 tells what player 1 does if he is shown number 1: he guesses 2 with probability q_1 and 5 with proba. $1 - q_1$; and so on.

Problem 13.1

Player 1, the catcher, and player 2, the evader, simultaneously and independently pick a node in a given graph. If they choose the same node or two adjacent nodes, player 2 is captured, otherwise he escapes. The payoff is the probability of capture, which Player 1 maximizes, and player 2 minimizes. Solve this game for the following graphs (*hint; use domination arguments*):
a) a line of arbitrary length.
b)

$$\begin{tikzcd}
\circ \arrow[r, leftrightarrow] \arrow[d, leftrightarrow] & \circ \arrow[d, leftrightarrow] & & \\
\circ \arrow[r, leftrightarrow] & \circ \arrow[r, leftrightarrow] & \circ \arrow[r, leftrightarrow] & \circ
\end{tikzcd}$$

c)

$$\begin{tikzcd}
\circ \arrow[r, leftrightarrow] & \circ \arrow[r, leftrightarrow] \arrow[d, leftrightarrow] & \circ \arrow[r, leftrightarrow] \arrow[d, leftrightarrow] & \circ & \\
\circ \arrow[r, leftrightarrow] & \circ \arrow[r, leftrightarrow] & \circ \arrow[r, leftrightarrow] & \circ \arrow[r, leftrightarrow] & \circ
\end{tikzcd}$$

d)

$$\begin{tikzcd}
\circ \arrow[r, leftrightarrow] \arrow[d, leftrightarrow] & \circ \arrow[d, leftrightarrow] & \\
\circ \arrow[r, leftrightarrow] & \circ \arrow[r, leftrightarrow] \arrow[d, leftrightarrow] & \circ \arrow[d, leftrightarrow] \\
& \circ \arrow[r, leftrightarrow] & \circ
\end{tikzcd}$$

Problem 13.2 *Catch me*
a) Player 1 chooses a location x in $[0, 1]$ and player 2 chooses simultaneously a location y. Player 1 is trying to be as far as possible from player 2, and player 2 has the opposite preferences. The payoff (to player 1)is $u(x, y) = (x - y)^2$. Show the game in pure strategies has no value. Find the value and optimal strategies for the game in mixed strategies.
b) Solve the similar game where the "board" is an arbitrary tree (connected graph with no cycles) and the payoff is the square of the distance on the tree.
c) Solve the similar game where the "board" is a circle, using the euclidian distance for defining the payoff..

Problem 14 *Hiding a number*
Fix an increasing sequence of positive numbers $a_1 \le a_2 \le a_3 \le \cdots \le a_p \le \cdots$. Each player chooses an integer, the choices being independent. If they both choose the same number p, player 1 receives \$$p$ from player 2. Otherwise, no money changes hand.
a) Assume first

$$\sum_{p=1}^{\infty} \frac{1}{a_p} < \infty$$

and show that each player has a unique optimal mixed strategy.
b) In the case where

$$\sum_{p=1}^{\infty} \frac{1}{a_p} = \infty$$

show that the value is zero, that every strategy of player 1 is optimal, whereas player 2 has only "ε-optimal" strategies, i.e., strategies guaranteeing a payoff not larger than ε, for arbitrarily small ε.

Problem 15
Asume that both players choose optimal (mixed) strategies $\overline{x}$ and $\overline{y}$ and thus the resulting payoff in the game is v. We know that player 1 would get v if against player 2's choice $\overline{y}$ he would play any pure strategy with positive probability in $\overline{x}$ (i.e. any pure strategy i, such that $\overline{s}_i > 0$), and he would get less then v if he would play any pure strategy i, such that $\overline{x}_i = 0$. Explain why a rational player 1, who assumes that his opponent is also rational, should not choose a pure strategy i such that $\overline{x}_i > 0$ instead of $\overline{x}$.

Problem 16
In a two-person zero-sum game in normal form with a finite number of pure strategies, show that the set of all *mixed* strategies of player 1 which are part of some equilibrium of the game, is a convex subset of the set of player 1's mixed strategies.

Problem 17 *Bluffing game*
At the beginning, players 1 and 2 each put \$1 in the pot. Next, player 1 draws a card from a shuffled deck with equal number of black and red cards in it. Player 1 looks at his card (he does not show it to player 2) and decides whether to raise or fold. If he folds, the card is revealed to player 2, and the pot goes to player 1 if it is red, to player 2 if it is black. If player 1 raises, he must add \$1 to the pot, then player 2 must meet or pass. If she passes the game ends and player 1 takes the pot. If she meets, she puts \$$\alpha$ in the pot. Then the card is revealed and, again, the pot goes to player 1 if it is red, to player 2 if it is black..
Draw the matrix form of this game. Find its value and optimal strategies as a function of the parameter α. Is bluffing part of the equilibrium strategy of player 1?

Problem 17 *A simple poker game*
There are 3 cards, of value Low, Medium and High. Ann is dealt a card face down, with equal probability for each card. After seeing her card, Ann can bid or fold. if she folds she gives \$2 to Bob and the game ends. if she bids, Bill is dealt one of the remaining cards (with equal probability) face down. He looks at his card and can then Fold or See. If he folds he gives \$4 to Ann. If he sees the cards are revealed and the holder of the higher card wins \$10 from the loser.

Solve this game: optimal mixed strategies and value. Do the optimal strategies involve bluffing?

Problem 19 *Another poker game*
There are 3 cards, of value Low, Medium and High. Each player antes \$1 to the pot and Ann is dealt a card face down, with equal probability for each card. After seeing her card, Ann announces "Hi" or "Lo". To go Hi costs her \$2 to the pot, and Lo costs her \$1. Next Bill is dealt one of the remaining cards (with equal probability) face down. he looks at his card and can then Fold or See. If he folds the pot goes to Ann. If he sees he must match Ann's contribution to

the pot; then the pot goes to the holder of the higher card if Ann called Hi, or to the holder of the lower card if she called Lo.
Solve this game: how much would you pay, or want to be paid to play this game as Ann? How would you then play?