Advanced Game Theory

Herve Moulin Baker Hall 263; moulin@rice.edu

Rice University ECO 440 Spring 2009

1 Two person zero sum games

1.1 Introduction: strategic interdependency

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.

1.2 Two-person zero-sum games in strategic form

A two-person zero-sum game in strategic form 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 secure utility level 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. ■

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. ■

1.3 Two-person zero-sum 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 $\sigma : N \rightarrow N \cup \emptyset$ 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 $u : T(N) \rightarrow \mathbf{R}$.

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 $m = \max\limits_{1 \le i \le k} m_i$. 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.*

1.4 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 $X =$*

$\{(s_1, ..., s_n) : \sum_{i=1}^{n} s_i = 1, s_i \geq 0\}$, *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 $s \in X$ and player 2 chooses $y \in Y$, the expected payoff of the game is equal to the matrix product $s^T Ay$:

$$s^T Ay = (s_1, ..., s_n) \begin{pmatrix} a_{11} & ... & a_{1m} \\ ... & ... & ... \\ a_{n1} & ... & a_{nm} \end{pmatrix} \begin{pmatrix} y_1 \\ ... \\ y_m \end{pmatrix} = \sum_{i=1}^{n} \sum_{j=1}^{m} s_i a_{ij} y_j,$$

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

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

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

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

Given the above decomposition of $s^T Ay$, and $v_1(s) = \min_{y \in Y} s^T Ay$, the minimum of $s^T Ay$, will be attained at some pure strategy j (i.e., at some $e^j \in Y$). Indeed, if $s^T Ae^j > v_1(s)$ for all j, then we would have $s^T Ay = \sum y_j \left[s^T Ae^j \right] > v_1(s)$ for all $y \in Y$. Hence, $v_1(s) = \min_j s^T A_{\cdot j}$, and $v_1 = \max_{s \in X} \min_j s^T A_{\cdot j}$. 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 $\overline{s}$ 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 $X = \{(s_1, ..., s_n) : \sum_{i=1}^{n} s_i = 1, s_i \geq 0\}$, while for player 2 it is $Y = \{(y_1, ..., y_m) : \sum_{i=1}^{m} y_j = 1, y_j \geq 0\}$. Let $v_1(s) = \min_{y \in Y} s \cdot Ay$ be the smallest payoff player 1 can get if he chooses to play s. Then $v_1 = \max_{s \in X} v_1(s) = \max_{s \in X} \min_{y \in Y} s \cdot Ay$ is the secure utility level for player 1. Similarly, we define $v_2(y) = \max_{s \in X} s \cdot Ay$, and $v_2 = \min_{y \in Y} v_2(y) = \min_{y \in Y} \max_{s \in X} s \cdot Ay$ 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 $K_v = \{z \in \mathbb{R}^n : s \cdot z \leq v \text{ for all } s \in X\}$.

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

Now, $v_1 = \max_{s \in X} \min_{y \in Y} s \cdot Ay \geq \min_{y \in Y} \overline{s} \cdot Ay \geq \overline{v}$ (since $\overline{s} \cdot z \geq c = \overline{v}$ for all $z \in L$, i.e. for all $z = Ay$, where $y \in Y$). Next, $v_2 = \min_{y \in Y} \max_{s \in X} s \cdot Ay \leq \max_{s \in X} s \cdot A\overline{y} = \max_{s \in X} s \cdot \overline{z} = \max_{i=1, \cdots, n} \overline{z_i} \leq \overline{v}$ (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 $v_1(s)$ is reached at $\overline{s}$, while the minimal value of $v_2(y)$ is reached at $\overline{y}$. Thus, $\overline{s}$ and $\overline{y}$ constructed in the proof are optimal strategies for players 1 and 2 respectively. ■

1.5 Computation of optimal mixed strategies

How can we find the maximin strategy $\overline{s}$, the 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 $\overline{s}, \overline{y}$ be a pair of optimal strategies, and $v = \overline{s}^T A\overline{y}$. Since for all j we have that $\overline{s}^T Ae^j \geq \min_{y\in Y} \overline{s}^T Ay = v_1(\overline{s}) = v_1 = v$, it follows that $v = \overline{s}^T A\overline{y} = \overline{y}_1 \left[\overline{s}^T Ae^1\right] + ... + \overline{y}_m \left[\overline{s}^T Ae^m\right] \geq \overline{y}_1 v + ... + \overline{y}_m v = v\left(\overline{y}_1 + ... + \overline{y}_m\right) = v$, and the equality implies $\overline{s}^T A_{\cdot j} = \overline{s}^T Ae^j = v$ for all j such that $\overline{y}_j \neq 0$. Thus, player 2 receives her minimax value $v_2 = v$ by playing against $\overline{s}$ 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 $\overline{s}$ (i.e. any strategy i, such that $\overline{s}_i \neq 0$). Setting $S^* = \{i|\overline{s}_i > 0\}$ and $T^* = \{j|\overline{y}_j > 0\}$, we see that $\overline{s}, \overline{y}$ solve the following system with unknown s, y

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

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

The difficulty is to find the supports S^*, T^*, because there are 2^{n+m} possible choices, and no systematic way to guess! However we expect the two supports to be of the same size, and in fact for any game there exists an equilibrium (a saddle point in mixed strategies) where both supports have the same cardinality (exercise: prove this claim).

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 $s_i \geq 0$ 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 $s_i \geq 0$ 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 opatimal 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 the choice of the first (dominating) strategy is at least as good as the choice of the second (dominated) strategy, and in some cases it is strictly better. A player can always find an optimal mixed strategy using only undominated strategies.

Proposition 14 *If the row i of a matrix A is strictly dominated, then any optimal strategy $\overline{s}$ of player 1 has $\overline{s}_i = 0$. If the row i of a matrix A is dominated, then player 1 has an optimal strategy $\overline{s}$ such that $\overline{s}_i = 0$. 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. The same is true 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 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.

1.5.1 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 (s_1, s_2) and (y_1, y_2) must have all components positive. Let us repeat the argument above for the 2×2 case. We have $v = s^T A y = a_{11}s_1y_1 + a_{12}s_1y_2 + a_{21}s_2y_1 + a_{22}s_2y_2$, or

$$s_1(a_{11}y_1 + a_{12}y_2) + s_2(a_{21}y_1 + a_{22}y_2) = v.$$

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, $s_1(a_{11}y_1 + a_{12}y_2) + s_2(a_{21}y_1 + a_{22}y_2) \leq s_1v + s_2v = v$. Since $s_1 > 0$ and $s_2 > 0$, 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 $a_{11}s_1 + a_{21}s_2 = v$ and $a_{12}s_1 + a_{22}s_2 = v$. We also know that $s_1 + s_2 = 1$ and $y_1 + y_2 = 1$.

We have a linear system with 6 equations and 5 variables s_1, s_2, y_1, y_2 and v. The minimax theorem guarantees us that this system has a solution with $s_1, s_2, y_1, y_2 \geq 0$. 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 $[a_{11}, a_{22}] \cap [a_{12}, a_{21}] = \varnothing$.

1.5.2 2×n games

By focusing on the player who has two strategies, one computes the value as the solution of a tractable linear program. See the examples in Problem 9.

1.5.3 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 s is an optimal strategy for player 1, then it is also optimal for player 2.

1.6 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 choooses 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 15 *(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:

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

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

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

1.7 Von Neumann's Theorem

It generalizes the minimax theorem. The proof follows from the more general Nash Theorem in Chapter 4.

Theorem 16 *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 $t \in T$*

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

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 $m_i \leq s_i$, 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 $s_i \leq m_i \leq 1$:

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

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

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.

1.8 Problems for two person zero-sum games

1.8.1 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:

2	4	6	3		3	2	2	1
6	2	4	3		2	3	2	1
4	6	2	3		2	2	3	1

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* $\begin{bmatrix} + & - \\ - & + \end{bmatrix}$.

1.8.2 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)

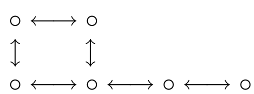

c)

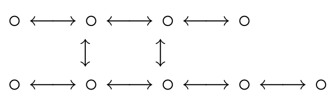

d)

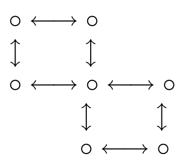

Problem 13.2 *Catch me*

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).
c) Solve the similar game where the "board" is a circle.

Problem 14 *Hiding a number*

Fix an increasing sequence of positive numbers $a_1 \leq a_2 \leq a_3 \leq \cdots \leq a_p \leq \cdots$.
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 18 *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?

2 Nash equilibrium

In a general n-person game in strategic form, interests of the players are neither identical nor completely opposed. As in the previous chapter information about other players' preferences and behavior will influence my behavior. The novelty is that this information may sometime be used *cooperatively,* i.e., to our mutual advantage.

We discuss in this chapter the two most important *scenarios* justifying the Nash equilibrium concept as the consequence of rational behavior by the players:

- the *coordinated scenarios* where players know a lot about each other's strategic opportunities (strategy sets) and payoffs (preferences), and use either deductive reasoning or non binding comunication to coordinate their choices of strategies.
- the *decentralized (competitive) scenarios* where mutual information is minimal, to the extent that a player may not even know how many other players are in the game or what their individual preferences look like.

Decentralized scenarios are realistic in games involving a large number of players, each one with a relatively small influence on the overall outcome, so that the "competitive" assumption that each player ignores the influence of his own moves on other players' strategic choices is plausible. Coordination scenarios are more natural in games with a small number of participants.

This chapter is long on examples and short on abstract proofs. The next chapter is just the opposite.

Definition 17 *A game in strategic form is a list $\mathcal{G} = (N, S_i, u_i, i \in N)$, where N is the set of players, S_i is player i's strategy set and u_i is his payoff, a mapping from $S_N = \prod_{i \in N} S_i$ into $\mathbb{R}$, which player i seeks to maximize.*

An important class of games consists of those where the roles of all players are fully interchangeable.

Definition 18 *A game in strategic form $\mathcal{G} = (N, S_i, u_i, i \in N)$ is symmetrical if $S_i = S_j$ for all i, j, and the mapping $s \to u(s)$ from $S^{|N|}$ into $\mathbb{R}^{|N|}$ is symmetrical.*

In a symmetrical game if two players exchange strategies, their payoffs are exchanged and those of other players remain unaffected.

Definition 19 *A Nash equilibrium of the game $\mathcal{G} = (N, S_i, u_i, i \in N)$ is a profile of strategies $s^* \in S_N$ such that*

$$u_i(s^*) \geq u_i(s_i, s^*_{-i}) \text{ for all } i \text{ and all } s_i \in S_i$$

Note that the above definition uses only the ordinal preferences represented by the utility functions u_i. We use the cardinal representation as payoff (utility) simply for convenience. However when we speak of mixed strategies in the next chapter, the choice of a cardinal utility will matter.

The following inequality provides a useful necessary condition for the existence of at least one Nash equilibrium in a given game $\mathcal{G}$.

Lemma 20 *If s^* is a Nash equilibrium of the game $\mathcal{G} = (N, S_i, u_i, i \in N)$, we have for all i*

$$u_i(s^*) \geq \min_{s_{-i} \in S_{-i}} \max_{s_i \in S_i} u_i(s_i, s_{-i})$$

Example 1 *duopoly a la Hoteling*
The two competitors sell identical goods at fixed prices p_1, p_2 such that $p_1 < p_2$. The consumers are uniformly spread on $[0, 1]$, each with a unit demand. Firms incur no costs. Firms choose independently where to locate a store on the interval $[0, 1]$, then consumers buy from the cheapest store, taking into account a transportation cost of \$$s$ if s is the distance to the store. Assume $p_2 - p_1 = \frac{1}{4}$. Check that

$$\min_{S_2} \max_{S_1} u_1 = p_1; \min_{S_1} \max_{S_2} u_2 = \frac{p_2}{8}$$

where the $\min_{S_2} \max_{S_1} u_1$ obtains from the copycat strategy $s_1 = s_2$ by player 1, and the $\min_{S_1} \max_{S_2} u_2$ is achieved by $s_1 = \frac{1}{2}$, and $s_2 = 0$ or 1. Observe now that the payoff profile $(p_1, \frac{p_2}{8})$ is not feasible, therefore the game has no Nash equilibrium.

2.1 Coordinated scenarios

We now consider games in strategic form involving only a few players who use their knowledge about other players strategic options to form expectations about the choices of these players, which in turn influence their own choices. In the simplest version of this analysis, each player knows the entire strategic form of the game, including strategy sets and individual preferences (payoffs). Yet at the time they make their strategic decision, they act independently of one another, and cannot observe the choice of any other player.

The two main interpretations of the Nash equilibrium are then the *self fulfilling prophecy* and the *self enforcing agreement.*

The former is the meta-argument that if a "Book of Rational Conduct" can be written that gives me a strategic advice for every conceivable game in strategic form, this advice must be to play a Nash equilibrium. This is the "deductive' argument in favor of the Nash concept.

The latter assumes the players engage in "pre-play" communication, and reach a non committal agreement on what to play, followed by a complete break up of communication.

Two conceptual difficulties suggest caution when we apply the Nash equilibrium concept in the coordinated context. First a Nash equilibrium may be *inefficient* (Pareto inferior), as illustrated in the celebrated *Prisoner's Dilemna*:

section 2.1.1. Then communication between the players drives them to move away from the equilibrium, for the benefit of every participant. Second, many games have multiple Nash equilibria, hence a *selection* problem (section 2.1.2). Under either scenario above, it may be unclear how the players will be able to coordinate on one of them.

On the other hand, we can identify large classes of games in which selecting the Nash outcome by deduction (covert communication) is quite convincing, so that our confidence in the predictive power of the concept remains intact. These are the dominance-solvable games in section 2.1.3, and the games with a dominant strategy equilibrium in section 2.1.4.

2.1.1 inefficiency of the Nash equilibrium outcomes

Example 2 *Prisonners Dilemna*

Each player chooses a *selfless* strategy C or a *selfish* strategy D. Choosing C brings a benefit a to every *other* player and a cost of b to me. Playing D brings neither benefit nor cost to anyone. It is a dominant strategy to play D if $b > 0$. If furthermore $b < (n-1)a$, the unique Nash equilibrium is Pareto inferior to the unanimously selfless outcome. This equilibrium is especially credible as each player uses a dominant strategy (see 2.1.4 below).

Example 3 *Pigou traffic example*

There are two roads (country, city) to go from A to B and n commuters want to do just that. The country road entails no congestion: no matter how many users travel on it, each incurs a delay of 1. The city road has linear congestion costs: if x commuters use that road, each of them incurs a delay of $\frac{x}{m}$, where we assume $m \leq n$. A Nash equilibrium is an outcome where m, or $m-1$, agents take the city road, and $n-m$, or $n-m+1$, take the country road, and all get a disutility of 1, or $\frac{m-1}{m}$. However total disutility is minimized by sending only $\frac{m}{2}$ commuters on the city road, for a total delay of $n - \frac{m}{4}$, and a Pareto improvement where $\frac{m}{2}$ city commuters are better off, while the rest are indifferent to the change.

Example 4 *the Braess paradox*

There are two roads to go from A to B, and 6 commuters. The upper road goes through C, the lower road goes through D. The 2 roads only meet at A and B. On each of the four legs, AC, CB, AD, DB, the travel time depends upon the number of users x in the following way:

on AC and $DB : 50 + x$, on CB and $AD : 10x$

Every player must choose a road to travel, and seeks to minimize his travel time. The Nash equilibria of the game are all outcomes with 3 users on each road, and they all give the same disutility 83 to each player. Next we add one more link on the road network, directly between C and D, with travel time $10 + x$. In the new Nash equilibrium outcomes, we have two commuters on each of the paths $ACB, ADB, ADCB$, and their disutility is 92. Thus the new road results in a net increase of the congestion!

2.1.2 the selection problem

When several (perhaps an infinity of) Nash outcomes coexist, and the players' preferences about them do not agree, they will try to force their preferred outcome by means of tactical commitment. Two well known games illustrate the resulting impossibility to predict the outcome of the game.

Example 5 *crossing game (a.k.a. the Battle of the Sexes)*
Each player must stop or go. The payoffs are as follows

stop	$1,1$	$1-\varepsilon,2$
go	$2,1-\varepsilon$	$0,0$
	stop	*go*

Each player would like to commit to go, so as to force the other to stop. A typical way is unilateral communication (schelling): I am going to pass, I cannot hear you anymore. There is a mixed strategy equilibrium as well, but it has its own problems. See Section 3.3.

Example 6 *Nash demand game*
The two players share a dollar by the following procedure: each write the amounts she demands in a sealed envelope. If the two demands sum to no more than \$1, they are honored. Otherwise nobody gets any money. In this game the equal plit outcome stands out because it is fair, and this will suffice in many cases to achieve coordination. However, a player will take advantage of an opportunity to commit to a high demand. More precisely, the pair s_1, s_2 is a Nash equilibrum if and only if $0 \leq s_1, s_2 \leq 1$ and $s_1 + s_2 = 1$, or $s_1 = s_2 = 1$.

Note that Examples 5 and 6 are symmetric games with (many) asymmetric equilibria.

In both above examples and in the next one the key strategic intuition is that the opportunity to commit to a certain strategy by "burning the bridges" allowing us to play anything else, is the winning move provided one convinces the other player that the bridges are indeed gone.

Definition 21 *Given two functions $t \to a(t)$ and $t \to b(t)$, the corresponding game of timing is as follows. Each one of the two players must choose a time to stop the clock between $t = 0$ and $t = 1$. If player i stops the clock first at time t, his payoff is $u_i = a(t)$, that of player j is $u_j = b(t)$. In case of ties, each gets the payoff $\frac{1}{2}(a(t) + b(t))$.*

An example is the noisy duel of chapter 1, where a increases, b decreases, and they intersect at the optimal stopping/shooting time (here *optimality* refers to the saddle point property for this ordinally zero-sum game). Here is another classic example.

Example 7 *War of attrition*
This is a game of timing where both a and b are continuous and decreasing, $a(t) < b(t)$ for all t, and $b(1) < a(0)$. There are two Nash equilibrium outcomes. Setting t^* as the time at which $a(0) = b(t^*)$, one player commits to t^* or more, and the other concedes by stopping the clock immediately (at $t = 0$).

The selection problem can often be alleviated by further arguments of salience, Pareto dominance, or risk dominance. It is easy to agree on an equilibrium more favorable to everyone: the Pareto dominance argument.

Definition 22 *A coordination game is a game $\mathcal{G} = (N, S_i, u_i, i \in N)$ such that all players have the same payoff function: $u_i(s) = u_j(s)$ for all $i \in N, s \in S_N$.*

If, in a coordination game, there is a single outcome maximizing the common payoff, this Nash equilibrium will be selected without explicit comunication. We have no such luck in a coordination game where several outcomes are optimal, as in Schelling's *rendez-vous game*. Two players living in a big city and unable to communicate directly, must meet tomorrow at noon. If they show up at the same *salient* location (e.g., the Eiffel tower in Paris), they both win a prize, otherwise they get nothing. The problem here is that salience may not be a deterministic criteria.

We illustrate finally the risk dominance argument, in an important model where it conflicts with Pareto dominance.

Example 8 *Coordination failure*
This is an example of a public good provision game by voluntary contributions (example 26), where individual contributions enter the common benefit function as perfect complements:

$$u_i(s) = \min_j s_j - C_i(s_i)$$

Examples include the building of dykes or a vaccination program: the safety provided by the dyke is only as good as that of its weakest link. Assume C_i is convex and increasing, with $C_i(0) = 0$ and $C_i'(0) < 1$, so that each player has a stand alone optimal provision level s_i^* maximizing $z - C_i(z)$. Then the Nash equilibria are the outcomes where $s_i = \lambda$ for all i, and $0 \leq \lambda \leq \min_i s_i^*$. They are Pareto ranked: the higher λ, the better for everyone. However the higher λ, the more risky the equilibrium: if other players may make an error and fail to send their contribution, it is prudent not to send anything ($\max_{s_i} \min_{s_{-i}} u_i(s) = 0$ is achieved with $s_i = 0$). Even if the probability of an error is very small, a reinforcement effect will amplify the risk till the point where only the null (prudent) equilibrium is sustainable.

2.1.3 dominance solvable games

Eliminating dominated strategies is the central coordination device performed by independent deductions of players mutually informed about the payoff functions. We repeat a definition already given for two-person zero-sum games (Definition 13).

Definition 23 *In the game $\mathcal{G} = (N, S_i, u_i, i \in N)$, we say that player i's strategy s_i is weakly dominated by his strategy s_i' (or simply dominated) if*

$$\begin{aligned} u_i(s_i, s_{-i}) &\leq u_i(s_i', s_{-i}) \text{ for all } s_{-i} \in S_{-i} \\ u_i(s_i, s_{-i}) &< u_i(s_i', s_{-i}) \text{ for some } s_{-i} \in S_{-i} \end{aligned}$$

We say that strategy s_i is strictly dominated by s'_i if

$$u_i(s_i, s_{-i}) < u_i(s'_i, s_{-i}) \text{ for all } s_{-i} \in S_{-i}$$

Given a subset of strategies $T_i \subset S_i$ we write $\mathcal{WU}_i(T_N)$ (resp. $\mathcal{U}_i(T_N)$) for the set of player i's strategies in the restricted game $(N, T_i, u_i, i \in N)$ that are not dominated (resp. not strictly dominated).

Definition 24 *We say that the game $\mathcal{G}$ is dominance-solvable (resp. strictly dominance-solvable) if the sequence defined inductively by*

$${}^wS_i^0 = S_i; {}^wS_i^{t+1} = \mathcal{WU}_i({}^wS_N^t) \text{ (resp. } S_i^0 = S_i; S_i^{t+1} = \mathcal{U}_i(S_N^t)) \text{ for all } i \text{ and } t = 1, 2, \cdots$$

and called the successive elimination of dominated (resp. strictly dominated) strategies, converges to a single outcome s^*:

$$\cap_{t=1}^{\infty} {}^wS_N^t = \{s^*\} \text{ (resp. } \cap_{t=1}^{\infty} S_N^t = \{s^*\})$$

If the strategy sets are finite, or compact with continuous payoff functions, the set of undominated strategies $\mathcal{U}_i(S_N)$ is non empty and closed, therefore the sequence S_N^t is well defined. On the other hand, the (smaller) set of weakly undominated strategies $\mathcal{WU}_i(S_N)$ is non empty but it may not be closed. Therefore the existence of the sequence ${}^wS_N^t$ is not always guaranteed, except if the strategy sets are finite.

Despite their close similarities, the two types of elimination, of dominated or of strictly dominated strategies, differ in other important ways. The latter never throws away a Nash equilibrium outcome, and so it is not a selection tool, rather a way to identify games with a unique Nash equilibrium. The former, on the other hand, is a genuine selection tool, but one that must be handled with care.

Proposition 25 *For any T the set $\cap_{t=1}^{T} S_N^t$ contains all Nash equilibria of the game. If $\cap_{t=1}^{\infty} S_N^t = \{s^*\}$, then s^* is the single Nash equilibrium outcome of the original game.*

If $\cap_{t=1}^{\infty} {}^wS_N^t = \{s^*\}$, then s^* is a Nash equilibrium of the original game.

The successive elimination of strictly dominated strategies is very robust in the sense that it never loses equilibria , whereas the successive elimination of weakly dominated strategies may lose some, or even all Nash equilibria of the original game (in the latter case, the game reduced to $\cap_{t=1}^{T} {}^wS_N^t$ contains no equilibrium either). Here are two examples

$$\begin{bmatrix} 1,0 & 2,0 & 1,5 \\ 6,2 & 3,7 & 0,5 \\ 3,1 & 2,3 & 4,3 \end{bmatrix}$$

where the elimination of weklay d.s. picks one of the two equilibria, and

$$\begin{bmatrix} 1,3 & 2,0 & 3,1 \\ 0,2 & 2,2 & 0,2 \\ 3,1 & 2,0 & 1,3 \end{bmatrix}$$

where the algorithm throws out the unique Nash equilibrium!

Another difference between the two successive elimination algorithms, based on strict or weak domination, is their robustness with respect to partial elimination. Suppose that at each stage we only drop strictly dominated strategies, i.e., we construct a sequence R_i^t such that $R_i^{t+1} \subseteq \mathcal{U}_i(R_N^t)$ for all i and t. Then it is easy to check that the limit set $\cap_{t=1}^{\infty} R_N^t$ is unaffected, provided we do eliminate some strategies at each round (see Problem 10). On the other hand when we only drop some weakly dominated strategies at each stage, the result of the algorithm may well depend on which ones we drop. Here is an example:

$$\begin{bmatrix} 2,3 & 2,3 \\ 3,2 & 1,2 \\ 1,1 & 0,0 \\ 0,0 & 1,1 \end{bmatrix}$$

Depending on which strategy player 1 eliminates first, we wend up at the $(3,2)$ or the $(2,3)$ equilibrium.

The bottom line is that the successive elimination of strictly dominated strategies can be performed without thinking twice, while we must be cautious in performing the successive elimination of strictly dominated strategies, that can lead to paradoxical examples We use several classic examples to reinforce this point.

Example 9 *Guessing game*
Each one of the n players chooses an integer s_i between 1 and 1000. Compute the average response

$$\overline{s} = \frac{1}{n} \sum_i s_i$$

Each player receives a prize that strictly decreases with the distance of its own strategy s_i to $\frac{2}{3}\overline{s}$

$$u_i(s) = -f(|s_i - \frac{2}{3}\overline{s}|)$$

This game is strictly dominance solvable and

$$\cap_{t=1}^{\infty} S_N^t = \{(1, \cdots, 1)\}$$

Observe that for any $t = 0, 1, \cdots$, if $S_i^t \subseteq \{1, \cdots, p\}$ for some integer p, then $S_i^{t+1} \subseteq \{1, \cdots, \lceil \frac{2}{3}p \rceil\}$. To prove this claim we check that player i's strategy $s_i^* = \lceil \frac{2}{3}p \rceil$ strictly dominates any strategy s_i such that $s_i \geq s_i^* + 1$. Assume player i uses s_i^* and denote by $\widetilde{s}$ the average strategy of players other than i, so that $\overline{s} = \frac{1}{n}s_i^* + \frac{n-1}{n}\widetilde{s}$. Simple computations give

$$\widetilde{s} \leq p \Rightarrow s_i^* \geq \frac{2}{3}\overline{s} \text{ and } s_i^* - \frac{2}{3}\overline{s} < s_i - \frac{2}{3}(\frac{1}{n}s_i + \frac{n-1}{n}\widetilde{s})$$

so s_i^* is strictly closer to $\widetilde{s}$ than s_i. We can now apply the upper bound on S_i^{t+1} repeatedly:

$S_i^1 \subseteq \{1, \cdots, 667\}, S_i^2 \subseteq \{1, \cdots, 445\}, \cdots, S_i^8 \subseteq \{1, \cdots, 40\}, \cdots, S_i^{16} \subseteq \{1, 2\}$. Finally if the game is reduced to the strategies 1 and 2 for everyone, check that strategy 2 is at least $\frac{2}{3}$ away from $\frac{2}{3}\bar{s}$, while strategy 1 is at most $\frac{1}{3}$ away from $\frac{2}{3}\bar{s}$.

The guessing game has been widely tested in the lab, where the participants' limited strategic sophistication lead them to perform only a couple (typically two or three) of rounds of elimination. When playing the guessing game with inexperienced opponents, it is therefore a good idea to choose a number between $(\frac{2}{3})^2 50$ and $(\frac{2}{3})^3 50$.

Example 10 *Cournot duopoly*

Firm i produces s_i units of output, at a unit cost of c_i. The price at which the total supply $s_1 + s_2$ clears is $[A - (s_1 + s_2)]_+$. Hence the profit functions:

$$u_i = [A - (s_1 + s_2)]_+ s_i - c_i s_i \text{ for } i = 1, 2$$

This game is strictly dominance-solvable.

In our next example, weak dominance solvability leads to a mildly paradoxical result.

Example 11 *The chair's paradox*

Three voters choose one of three candidates a, b, c. The rule is plurality with the Chair, player 1, breaking ties. Hence each player i chooses from the set $S_i = \{a, b, c\}$, and the elected candidate for the profile of votes s is

$$s_2 \text{ if } s_2 = s_3; \text{ or } s_1 \text{ if } s_2 \neq s_3$$

Note that the Chair has a dominant strategy (Definition 25 below) to vote for her top choice. The two other players can only eliminate the vote for their bottom candidate as (weakly) dominated.

Assume that the preferences of the voters exhibit the cyclical pattern known as the *Condorcet paradox*, namely

$$u_1(c) < u_1(b) < u_1(a)$$

$$u_2(b) < u_2(a) < u_2(c)$$

$$u_3(a) < u_3(c) < u_3(b)$$

Writing this game in strategic form reveals that after the successive elimination of dominated strategies, the single outcome $s = (a, c, c)$ remains. This is the most plausible Nash equilibrium outcome when players know all preferences. The paradox is that the chair's tie-breaking privilege result in the election of her worst outcome! There are other equilibria; two examples are: everyone votes for a, or everyone for b.

In spite of the shortcomings detailed above, in many important economic games, a couple of rounds of elimination of weakly dominated strategiesf may well be enough to select a unique Nash equilibrium, even though the elimination algorithm is stopped and the initial game is not weakly dominance solvable.

Example 12 *First price auction*
The sealed bid first price auction is strategically equivalent to the Dutch descending auction.An object is auctioned between n bidders who each submit a sealed bid s_i. Bids are in round dollars (so $S_i = \mathbb{N}$). The highest bidder gets the object and pays his bid. In case of a tie, a winner is selected at random with uniform probability among the highest bidders. Assume that the valuations of (willingness to pay for) the object are also integers u_i and that

$$u_1 > u_i \text{ for all } i \geq 2$$

At a Nash equilibrium of this game, the object is awarded to player 1 at a price anywhere between $u_1 - 1$ and u_2, and there is another bid just below player 1's winning bid. However after two rounds of elimination we find a game where the only Nash equilibrium has player 1 paying u_2 for the object while one of the players i, $i \geq 2$, such that $u_i = \max_{j \neq 1} u_j$ bids $u_i - 1$. Thus player 1 exploits his informational advantage to the full.

Example 13 *Steinhaus cake division method*
The referee runs a knife from the left end of the cake to its right end. Each one of the two players can stop the knife at any moment. Whoever stops the knife first gets the left piece,the other player gets the right piece.
If both players have identical preferences over the various pieces of the cake, this is a game of timing structurally equivalent to the noisy duel, and its unique Nash equilibrium is that they both stop the knife at the time t^* when they are indifferent between the two pieces.
When preferences differ, this is a variant of a game of timing. Call t_i^* the time when player i is indifferent between the two pieces, and assume $t_1^* < t_2^*$. The Nash equilibrium outcomes are those where player 1 stops the knife between t_1^* and t_2^* while player 2 is just about to stop it herself: player 1 gets the left piece (worth more than the right piece to him) and player 2 gets the right piece (worth more to her than the left piece). However after two rounds of elimination of weakly dominated strategies, we are left with $S_1^2 = [t_2^* - \varepsilon, 1], S_2^2 = [t_2^*, 1]$ (where our notation is loose; for a precise statement, it is easier to give a discrete version of the model). Although the elimination process stops there, the outcome of the remaining game is not in doubt: $s_1^* = t_2^* - \varepsilon, s_2^* = t_2^*$. Indeed the remaining game is *inessential* (see Problem 29, question a).

2.1.4 dominant strategy equilibrium

One instance where the successive elimination of weakly dominated strategies is convincing is when each player has a dominant strategy. Put differently, the following is a compelling equilibrium selection.

Definition 26 *In the game $\mathcal{G} = (N, S_i, u_i, i \in N)$, we say that player i's strategy s_i^* is dominant if*

$$u_i(s_i^*, s_{-i}) \geq u_i(s_i, s_{-i}) \text{ for all } s_{-i} \in S_{-i}, \text{ all } s_i \in S_i$$

We say that s^* is a dominant strategy equilibrium if for each player i, s_i^* is a dominant strategy.

There is a huge difference in the interpretation of a game where dominance solvability (whether in the strict or weak form) identifies a Nash equilibrium, versus one where a dominant strategy equilibrium exists.

The former requires complete information about mutual preferences and more: I know your preferences, I know that you know that I know your preferences, etc..You know my preferences, I know that you know my preferences, you know that I know ...

In the latter, all a player has to know are the strategy sets of other players; their preferences or their actual strategic choices do not matter at all to pick his dominant strategy. Strategic choices are truly decentralized. Information about other players' payoffs or moves is worthless, as long as our player is unable to influence their choices (no direct communication channel allows to convey a threat of the kind "if you do this I will do that", or this threat is not enforceable).

The Prisoner's Dilemna (Example 1) is the most famous instance of a game with a dominant strategy equilibrium.

Dominant strategy equilibria are rare because the strategic interaction is often more complex. However they are so appealingly simple that when we design a procedure to allocate resources, elect one of the candidates to a job, or divide costs, we would like the corresponding strategic game to have a dominant strategy equilibrium as often as possible. In this way we are better able to predict the behavior of our participants. The two most important examples of such *strategy-proof* allocation mechanisms follow. In both cases the game has a (weakly but not strictly) dominant strategy equilibrium for all preference profiles, and the corresponding outcome is efficient (Pareto optimal).

Example 14 *Vickrey's second price auction*
An object is auctioned between n bidders who each submit a sealed bid s_i. Bids are in round dollars (so $S_i = \mathbb{N}$). The highest bidder gets the object and pays *the second highest bid*. In case of a tie, a winner is selected at random with uniform probability among the highest bidders (and pays the highest bid). If player i's valuation of the object is u_i, it is a dominant strategy to bid "sincerely", i.e., $s_i^* = u_i$. The corresponding outcome is the same as in the Nash equilibrium of the first price auction that we selected by dominance-solvability in example 12. But to justify that outcome we needed to assume complete information, in particular the highest valuation player must know precisely the second highest valuation. By contrast in the Vickrey auction, each player knows what bid to slip in the envelope, whether or not she has any information about other players' valuations, or even their number.

Note that in the second price auction game, there is a distressing variety of other Nash equilibrium outcomes. In particular any player, even the one with the lowest valuation of all, receives the object in some equilibrium. It is easy to check that for any player i and for any price p, $0 < p < a_i$ there is a Nash equilibrium where player i gets the object and pays p.

Example 15 *voting under single-peaked preferences*

The n players vote to choose an outcome x in $[0,1]$. Preferences of player i over the outcomes are single-peaked with the peak at v_i: they are strictly increasing on $[0, v_i]$ and strictly decreasing on $[v_i, 1]$. Assume for simplicity n is odd. Each player submits a ballot $s_i \in [0,1]$, and the *median* outcome among $s_1, \cdots, s_n$ is elected: this is the number $x = s_{i^*}$ such that more than half of the ballots are no less than x, and more than half of the ballots are no more than x.

It is a dominant strategy to bid "sincerely", i.e., $s_i^* = v_i$. Again, any outcome x in $[0,1]$ results from a Nash equilibrium, so the latter concept has no predictive power at all in this game.

2.2 Decentralized behavior and dynamic stability

In this section we interpret a Nash equilibrium as the resting point of a dynamical system. The players behave in a simple myopic fashion, and learn about the game by exploring their strategic options over time. Their behavior is compatible with total ignorance about the existence and characteristics of other players, and what their behavior could be.

Think of Adam Smith's *invisible hand* paradigm: the price signal I receive from the market looks to me as an exogenous parameter on which my own behavior has no effect. I do not know how many other participants are involved in the market, and what they could be doing. I simply react to the price by maximizing my utility, without making assumptions about its origin.

The analog of the *competitive behavior* in the context of strategic games is the *best reply behavior*. Take the profile of strategies s_{-i} chosen by other players as an exogeneous parameter, then pick a strategy s_i maximizing your own utility u_i, under the assumption that this choice will not affect the parameter s_{-i}.

The deep insight of the invisible hand paradigm is that decentralized price taking behavior will result in an efficient allocation of resources (a Pareto efficient outcome of the economy). This holds true under some specific microeconomic assumptions in the Arrow-Debreu model, and consists of two statements. First the invisible hand behavior will converge to a competitive equilibrium; second, this equilibrium is efficient. (The second statement is much more robust than the first).

In the much more general strategic game model, the limit points of the best reply behavior are the Nash equilibrium outcomes. Both statements, the best reply behavior converges, the limit point is an efficient outcome, are problematic. The examples below show that the best reply behavior may not converge at all If it converges, the limit Nash equilibrium outcome may well be inefficient (as we saw in section 2.1). Decentralized behavior may diverge, or it may converge toward a socially suboptimal outcome.

2.2.1 Stable and unstable equilibria

Definition 27 *Given the game in strategic form $\mathcal{G} = (N, S_i, u_i, i \in N)$, the best-reply correspondence of player i is the (possibly multivalued) mapping br_i*

from $S_{-i} = \prod_{j \in N \setminus \{i\}} S_j$ *into* S_i *defined as follows*

$$s_i \in br_i(s_{-i}) \Leftrightarrow u_i(s_i, s_{-i}) \geq u_i(s_i', s_{-i}) \text{ for all } s_i' \in S_i$$

Definition 28 *We say that the sequence* $s^t \in S_N, t = 0, 1, 2, \cdots,$ *is a* ***best reply dynamics*** *if for all* $t \geq 1$ *and all* i, *we have*

$$s_i^t \in \{s_i^{t-1}\} \cup br_i(s_{-i}^{t-1}) \text{ for all } t \geq 1$$

$$\text{and } s_i^t \in br_i(s_{-i}^{t-1}) \text{ for infinitely many values of } t$$

We say that s^t *is a* ***sequential best reply dynamics*** *if in addition at each step at most one player is changing her strategy.*

The best reply dynamics is very general, in that it does not require the successive adjustments of the players to be synchronized. If all players use a best reply at all times, we speak of *myopic adjustment;* if our players take turn to adjust, we speak of *sequential adjustment*. For instance with two players the latter dynamics is:

$$\text{if } t \text{ is even: } s_1^t \in br_i(s_2^{t-1}), s_2^t = s_2^{t-1}$$

$$\text{if } t \text{ is odd: } s_2^t \in br_i(s_1^{t-1}), s_1^t = s_1^{t-1}$$

But the definition allows much more complicated dynamics, where the timing of best reply adjustments varies accross players. An important requirement is that at any date t, every player will be using his best reply adjustment some time in the future.

The first observation is an elementary result.

Proposition 29 *Assume the strategy sets* S_i *of each player are compact and the payoff functions* u_i *are continuous (this is true in particular if the sets* S_i *are finite). If the best reply dynamics* $(s^t)_{t \in \mathbb{N}}$ *converges to* $s^* \in S_N$, *then* s^* *is a Nash equilibrium.*

Proof. Pick any $\varepsilon > 0$. As u_i is uniformly continuous on S_N, there exists T such that

$$\text{for all } i, j \in N \text{ and } t \geq T\text{: } |u_i(s_j^t, s_{-j}) - u_i(s_j^*, s_{-j})| \leq \frac{\varepsilon}{n} \text{ for all } s_{-j} \in S_{-j}$$

Fix an agent i. By definition of the b.r. dynamics, there is a date $t \geq T$ such that $s_i^{t+1} \in br_i(s_{-i}^t)$. This implies for any $s_i \in S_i$

$$u_i(s^*) + \varepsilon \geq u_i(s_i^{t+1}, s_{-i}^t) \geq u_i(s_i, s_{-i}^t) \geq u_i(s_i, s_{-i}^*) - \frac{n-1}{n}\varepsilon$$

where the left and right inequality follow by repeated application of uniform continuity. Letting ε go to zero ends the proof. ■

Observe that a limit point s^* of the best reply dynamics $(s^t)_{t \in \mathbb{N}}$ is typically not a Nash equilibrium! The second game in Example 16 below is a case in point; see also Example 19.

Definition 30 *We call a Nash equilibrium s* ***strongly stable*** *if any best reply dynamics (starting form any initial profile of strategies in S_N) converges to s. Such an equilibrium must be the unique equilibrium.*

We call a Nash equilibrium ***sequentially stable*** *if any sequential best reply dynamics (starting form any initial profile of strategies in S_N) converges to it. Such an equilibrium must be the unique equilibrium.*

We give a series of examples illustrating these definitions.

Example 16: *Two-person zero sum games*
Here a Nash equilibrium is precisely a saddle point. In the following game, a saddle point exists and is strongly stable

$$\begin{bmatrix} 4 & 3 & 5 \\ 5 & 2 & 0 \\ 2 & 1 & 6 \end{bmatrix}$$

Check that 3 is the value of the game. To check stability check that from the entry with payoff 1, any b.r. dynamics converges to the saddle point; then the same is true from the entry with payoff 6; then also from the entry with payoff 0, and so on.

In the next game, a saddle point exists but is not even sequentially stable:

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

Starting from (top,left), say, we cycle on the four coprners of the matrix, each one of them a limit point of the sequential b.r. dynamics, but we never reach the saddle point (middle,middle).

Stability in finite a (not necessarily zero-sum) two person game (S_1, S_2, u_1, u_2) is easy to analyze. Define $f = br_2 \circ br_1$ the composition of the two best reply correspondences. A fixed point of f is $s_2 \in S_2$ such that $s_2 \in f(s_2)$, and a cycle of length T is a sequence of *distinct* elements $s_2^t, t = 1, \cdots, T$ such that $s_2^{t+1} \in f(s_2^t)$ for all $t = 1, \cdots, T-1$, and $s_2^1 \in f(s_2^T)$.

Proposition 31 *The Nash equilibrium s^* of the finite game (S_1, S_2, u_1, u_2) is strongly stable if and only if it is sequentially stable. This happens if and only if f has a unique fixed point and no cycle of length 2 or more.*

Proof. If the game is sequentially stable, a sequence s_2^t with an arbitrary starting point s_2 and such that $s_2^{t+1} \in f(s_2^t)$, converges to the same limit s_2^*, and the corresponding sequence s_1^t also has a unique limit s_1^*. Thus (s_1^*, s_2^*) is the unique Nash equilibrium outcome and s_2^* the unique fixed point of f. To check strong stability, consider any best reply dynamics s^t. At some $t \geq 1$, $s_1^{t+1} \in br_1(s_2^t)$, so s^t reaches the set $br_1(S_2) \times S_2$, and never leaves it thereafter. at some $t' > t$, $s_2^{t'+1} \in br_2(s_1^{t'})$, so the sequence s^t reaches the set $br_1(S_2) \times f(S_2)$, never to leave it. Repeating the argument, we see that the sequence reaches $br_1 \circ f(S_2) \times f(S_2)$, then $br_1 \circ f(S_2) \times f^2(S_2)$, and so on, which ensures its convergence to (s_1^*, s_2^*).

The easy proof of the second statement is omitted. ■

Example 17 *price cycles in the Cournot oligopoly*
The demand function and its inverse are

$$D(p) = (a - bp)_+ \Leftrightarrow D^{-1}(q) = \frac{(a-q)_+}{b}$$

Firm i incurs the cost $C_i(q_i) = \frac{q_i^2}{2c_i}$ therefore its competitive supply given the price p is $O_i(p) = c_i p$, and total supply is $O(p) = (\sum_N c_i)p$. Assume there are many agents, each one small w.r.t. the total market size (i.e., each c_i is small w.r.t. $\sum_N c_j$), so that the competitive price-taking behavior is a good approximation of the best reply behavior. Strategies here are the quantities q_i produced by the firms, and utilities are

$$u_i(q) = D^{-1}(\sum_N q_j)q_i - C_i(q_i)$$

The equilibrium is unique, at the intersection of the O and D curves. If $\frac{b}{c} > 1$ it is strongly stable; if $\frac{b}{c} < 1$ it is sequentially but not strongly stable.

Example 18: *Schelling's model of binary choices*
Each player has a binary choice, $S_i = \{0, 1\}$, and the game is symmetrical, therefore it is represented by two functions $a(.), b(.)$ as follows

$$\begin{aligned} u_i(s) &= a(\frac{1}{n}\sum_N s_i) \text{ if } s_i = 1 \\ &= b(\frac{1}{n}\sum_N s_i) \text{ if } s_i = 0 \end{aligned}$$

Assuming a large number of agents, we can draw a, b as continuous functions and check that the Nash equilibrium outcomes are at the intersections of the 2 graphs, at $s = (0, \cdots, 0)$ if $a(0) \leq b(0)$, and at $s = (1, \cdots, 1)$ if $a(1) \geq b(1)$.

Whether a cust b from above or below makes a big difference in the stability of the corresponding equilibrium outcome.

Example 18a: vaccination Strategy 1 is to take the vaccine, strategy 0 to avoid it. Both a and b are strictly increasing: the risk of catching the disease diminishes as more people around us vaccinate. If $\frac{1}{n}\sum_N s_i$ is very small, $a > b$, as the risk of catching the disease is much larger than the risk of complications from the vaccine; this inequality is reversed when $\frac{1}{n}\sum_N s_i$ is close to 1. So the intersection of the two curves is the sequentially stable, but not strongly stable, equilibrium outcome[1].

[1] Note that this is a statement *in utilities*, as the Nash equilibrium property only determines the number of players using each strategy, but not their identity. Yet all equilibrium outcomes yield the same utility profile, which allows us to state those stability properties.

Example 18b: traffic Each player chooses to use the bus ($s_i = 1$) or his own car ($s_i = 0$); for a given congestion level $\frac{1}{n}\sum_N s_i$, traffic is equally slow in either vehicle, but more comfortable in the car, so $a(t) < b(t)$ for all t; however a and b both increase in t, as more people riding the bus decreases congestion. If $a(1) > b(0)$ the equilibrium in dominant strategies $s_i = 0$ for all i is Pareto inferior.

Example 18c Now a and b intersect only once, and a cuts b from below. Now we have three equilibrium outcomes, at $t = 0, 1$ and at the intersection of a and b. The latter is unstable, and the former two are stable in a "local" sense[2].

2.2.2 strictly dominance-solvable games

We saw in section 2.1.3 that in such games the Nash equilibrium exists and is unique. We can say more.

Proposition 32 *If the game $\mathcal{G} = (N, S_i, u_i, i \in N)$ is strictly dominance solvable, its unique Nash equilibrium $\cap_{t=1}^{\infty} S_N^t = \{s^*\}$ is globally stable.*

The unique equilibrium obtains both as the result of the (timeless) deductive process of successive elimination of strategies by fully informed players, and also as the limit of any best reply dynamics by players with very limited knowledge of their environment who naively "best reply" to the observed behavior of the other players (unaware of those players' preferences).

See section 2.1.3 for examples.

2.2.3 potential games

Potential games generalize the pure coordination games (Definition 21) where all players have the same payoff functions. As shown in the following example, strong stability is problematic in a coordination game but sequential stability is not.

Example 19 *a simple coordination game*
The game is symmetrical and the common strategy space is $S_i = [0, 1]$; the payoffs are identical for all n players

$$u_i(s) = g(\sum_{i=1}^{n} s_i)$$

where g is a continuous function on $[0, n]$.
Suppose first that g has a unique maximum z^* and no other local maxima (g is single-peaked). All s such that $\sum_{i=1}^{n} s_i = z^*$ are Nash equilibria, therefore none is strongly stable. The single exceptions are $z^* = 0$ or 1, because then the

[2] We measure the deviation from an equilibrium by the number of agents who are not playing the equilibrium strategy. We say that a a Nash equilibrium s^* is *locally stable in population* if for any number $\lambda, 0 < \lambda < 1$, there exists $\mu, 0 < \mu < 1$, such that if a fraction not larger than μ of the agents change strategies, any sequential b.r. dynamics converges to an equlibrium where at most λ of the players have changed from the original equilibrium.

Nash equilibrium is unique, and strongly stable. For z^* such that $0 < z^* < 1$, the *game* is sequentially stable *in utilities* (Example 18), because along any sequential best reply dynamics, the common utility increases and converges to $g(z^*)$. However, even in the restricted sense of convergence in utilities, the game is not strongly stable, because, for instance, simultaneous best reply sequences cycle around z^* without reaching it.

Definition 33 *A game in strategic form $\mathcal{G} = (N, S_i, u_i, i \in N)$ is a potential game if there exists a real valued function P defined on S_N such that for all i and $s_{-i} \in S_{-i}$ we have*

$$u_i(s_i, s_{-i}) - u_i(s'_i, s_{-i}) = P(s_i, s_{-i}) - P(s'_i, s_{-i}) \text{ for all } s_i, s'_i \in S_i$$

or equivalently there exists P and for all i a real valued function h_i defined on $S_{N \diagdown \{i\}}$ such that

$$u_i(s) = P(s) + h_i(s_{-i}) \text{ for all } s \in S_N$$

The original game $\mathcal{G} = (N, S_i, u_i, i \in N)$, and the game $\mathcal{P} = (N, S_i, P, i \in N)$ with the same strategy sets as $\mathcal{G}$ and identical payoffs P for all players, have the same best reply correspondences therefore the same Nash equilibria. Call s^* a *coordinate-wise maximum* of P if for all i, $s_i \to P(s_i, s^*_{-i})$ reaches its maximum at s^*_i. Clearly s is a Nash equilibrium (of $\mathcal{G}$ and $\mathcal{P}$) if and only if it is a coordinate-wise maximum of P.

If P reaches its global maximum on S_N at s, this outcome is a Nash equilibrium of $\mathcal{P}$ and therefore of $\mathcal{G}$. Thus potential games with continuous payoff functions and compact strategy sets always have at least a Nash equilibrium. Moreover, these equilibria have appealing stability properties.

Proposition 34 *Let $\mathcal{G} = (N, S_i, u_i, i \in N)$ be a potential game where the sets S_i are compact and the payoff functions u_i are continuous. If the best reply function of every player is single valued and continuous, and the Nash equilibrium is unique, the game $\mathcal{G}$ is sequentially stable.*

Proof. (sketch) For any sequential b.r. dynamics s^t if $s^t \neq s^{t+1}$, we have $P(s^t) < P(s^{t+1})$.If the sequence s^t has more than one limit point, one constructas a cycle of the sequential best reply dynamics, which contradicts the fact that P strictly increases along such dynamics. Thus s^t converges, and by continuity of P it must be a coordinate-wise maximum of P, namely a Nash equilibrium. ■

Example 20 *public good provision by voluntary contributions*
Each player i contributes an amount of input s_i toward the production of a public good, at a cost $C_i(s_i)$. The resulting level of public good is $B(\sum_i s_i) = B(s_N)$. Hence the payoff functions

$$u_i = B(s_N) - C_i(s_i) \text{ for } i = 1, \cdots, n$$

The potential function is

$$P(s) = B(s_N) - \sum_i C_i(s_i)$$

therefore existence of a Nash equilibrium is guaranteed if B, C_i are continuous and the potential is bounded over $\mathbb{R}_+^N$.

Remark The public good provision model is a simple and compelling argument in favor of centralized control of the production of pure public goods. To see that in equilibrium the level of production is grossly inefficient, assume for simplicity identical cost functions $C_i(s_i) = \frac{1}{2}s_i^2$ and $B(z) = z$. The unique Nash equilibrium is $s_i^* = 1$ for all i, yielding total utility

$$\sum_i u_i(s^*) = nB(s_N^*) - \sum_i C_i(s_i^*) = n^2 - \frac{n}{2}$$

whereas the outcome maximizing total utility is $\widetilde{s_i} = n$, bringing $\sum_i u_i(\widetilde{s}) = \frac{n^3}{2}$, so each individual equilibrium utility is less than $\frac{2}{n}$ of its "utilitarian" level.

The much more general version of Example 20 where the common benefit is an arbitrary function $B(s) = B(s_1, \cdots, s_n)$, remains a potential game for $P = B - \sum_i C_i$, therefore existence of a Nash equilibrium is still guaranteed. See Example 8 and Problem 7 for two alternative choices of B, respectively $B(s) = \min s_i$ and $B(s) = \max s_i$.

Example 21 *congestion games*

These games generalize both Pigou's model (Example 3) and Schelling's model (Example 18). Each player i chooses from the same strategy set and her payoff only depends upon the number of other players making the same choice. Examples include choosing a travel path between a source and a sink when delay is the only consideration, choosing a club for the evening if crowding is the only criteria, and so on.

$S_i = S$ for all i; $u_i(s) = f_{s_i}(n_{s_i}(s))$ where $n_x(s) = |\{j \in N | s_j = x\}|$ and f_x is arbitrary. If f is decreasing, we have a negative congestion externality, as in traffic examples. If f is increasing we have the opposite effect where we want more players to choose the same strategy as our own, as in the club example. In the latter we can also think of f as single-peaked (some crowding is good, up to a point).

Here the potential function is

$$P(s) = \sum_{x \in S} \sum_{m=1}^{n_x(s)} f_x(m)$$

See Problems 16 to 19 for illustrations and variants.

2.3 problems on chapter 2

Problem 1

In Schelling's model (example 18) find the Nash equilibrium outcomes and analyze their stability in the following cases:

a) $a(t) = 8t(1-t); b(t) = t$

b) $a(t) = 8t(1-t); b(t) = 1-t$

c) $a(t) = 8t(1-t); b(t) = \frac{1}{2}$

Problem 2 *Games of timing (Definition 20)*

a) We have two players, a and b both increase, are continuous, and a intersects b from below only once. Perform the successive elimination of (weakly and strictly) dominated strategies, and find all Nash equilibria. Can they be Pareto improved?

b) We extend the war of attrition (example 7) to n players. If player i stops the clock first at time t, his payoff is $u_i = a(t)$, that of all other players is $u_j = b(t)$. Both a and b are continuous and decreasing, $a(t) < b(t)$ for all t, and $b(1) < a(0)$. Answer the same questions as in a).

c) We have n players as in question b), but this time a increases, b decreases, and they intersect.

Problem 3 *Example 13 continued*

The interval $[0, 1]$ is a nonhomogeneous cake to be divided between two players. The utility of player 1 for a share $A \subset [0, 1]$ is $v_1(A) = \int_A \left(\frac{3}{2} - x\right) dx$. The utility of player 2 for a share $B \subset [0, 1]$ is $v_2(B) = \int_B \left(\frac{1}{2} + x\right) dx$. When time runs from $t = 0$ to $t = 1$, a knife is moved at the speed 1 from $x = 0$ to $x = 1$. Each player can stop it at any time. If the knife is stopped at time t by player i, this player gets the share $[0, t]$, while the other player gets the share $[t, 1]$.

Analyze the game as in Example 13. What strategic advice would you give to each player? Distinguish the two cases where this player knows his opponent's utility and that where she does not.

Problem 4

One hundred people live in the village, of whom 51 support the conservative candidate and 49 support the liberal candidate. A villager gets utility +9 if her candidate wins, -11 if her candidate looses, and 0 if they are tied. In addition, she gets a disutility of -1 for actually voting, but no disutility for staying home (so if her canddate wins and she voted, net utility is 10, etc..).

a) Why it is not Nash equilibrium for everybody to vote?

b) Why it is not Nash equilibrium for nobody to vote?

c) Find a Nash equilibrium where all conservatives use the same strategy, and all liberals use the same strategy.

d) What can you say about other possible Nash equilibria of this game?

Problem 5 *third price auction*

We have n bidders, $n \geq 3$, and bidder i's valuation of the object is u_i. Bids are independent and simultaneous. The object is awarded to the highest bidder at the third highest price. Ties are resolved just like in the Vickrey auction, with the winner still paying the third highest price. We assume for simplicity that the profile of valuations is such that $u_1 > u_2 > u_3 \geq u_i$ for all $i \geq 4$.

a) Find all Nash equilibria.

b) Find all dominated strategies of all players and all Nash equilibria in undominated strategies.
c) Is the game dominance-solvable?

Problem 6 *tragedy of the commons*

A pasture produces 100 units of grass, and a cow transforms x units of grass into x units of meat (worth \$$x$), where $0 \leq x \leq 10$, i.e., a cow eats at most 10 units of grass. It cost \$2 to bring a cow to and from the pasture (the profit from a cow that stays at home is \$2). Economic efficiency requires to bring exactly 10 cows to the pasture, for a total profit of \$80. A single farmer owning many cows would do just that.

Our n farmers, each with a large herd of cows, can send any number of cows to the commons. If farmer i sends s_i cows, s_N cows will share the pasture and each will eat $\min\{\frac{100}{s_N}, 10\}$ units of grass.

a) Write the payoff functions and show that in any Nash equilibrium the total number s_N of cows on the commons is bounded as follows

$$50\frac{n-1}{n} - 1 \leq s_N \leq 50\frac{n-1}{n} + 1$$

b) Deduce that the commons will be overgrazed by at least 150% and at most 400%, depending on n, and that almost the entire surplus will be dissipated in equilibrium. (*Hint: start by assuming that each farmer sends at most one cow*).

Problem 7 *a public good provision game.*
The common benefit function is $B(s) = \max_j s_j$: a single contributor is enough. Examples include R&D, ballroom dancing (who will be the first to dance) and dragon slaying (a lone knight must kill the dragon). Costs are quadratic, so the payoff functions are

$$u_i(s) = \max_j s_j - \frac{1}{2\lambda_i} s_i^2$$

where λ_i is a positive parameter differentiating individual costs.
a) Show that in any Nash equilibrium, only one agent contributes.
b) Show that there are p such equilibria, where p is the number of players i such that

$$\lambda_i \geq \frac{1}{2} \max_j \lambda_j$$

c) Compute strictly dominated strategies for each player. For what profiles (λ_i) is our game (strictly) dominance-solvable?

Problem 8 *the lobbyist game*
The two lobbyists choose an 'effort' level $s_i, i = 1, 2$, measured in money (the amount of bribes distributed) and the indivisible prize worth \$$a$ is awarded randomly to one of them with probabilities proportional to their respective efforts (if the prize is divisible, no lottery is necessary). Hence the payoff functions

$$u_i(s) = a\frac{s_i}{s_1 + s_2} - s_i \text{ if } s_1 + s_2 > 0; u_i(0, 0) = 0$$

a) Compute the best reply functions and show there is a unique Nash equilibrium.
b) Perform the successive elimination of strictly dominated strategies, and check the game is not dominance-solvable. However, if we eliminate an arbitrarily small interval $[0, \varepsilon]$ from the strategy sets, the reduced game is dominance solvable.
c) Show that the Nash equilibrium (of the full game) is strongly stable.

Problem 9

Two players share a *well* producing x liters of water at a cost $C(x) = \frac{1}{2}x^2$. Player i requests x_i liters of water, and the cost $C(x_1 + x_2)$ of pumping the total demand of water is divided in proportion to individual demands: player i pays $x_i \frac{C(x_1+x_2)}{x_1+x_2}$.

Player i's utility for x_i liters of water at cost c_i is

$$v_i(x_i) = 84 \log(1 + x_i) - c_i$$

a) Write the normal form of the game where each player chooses independently how much water to request.

b) Compute the quantities $\min_{x_j} \max_{x_i} u_i(x_i, x_j)$ and find the Nash equilibrium outcome. Show it is unique.

c) Is the Nash equilibrium outcome Pareto optimal? If not, compute the outcome maximizing total utility and compute the welfare loss at the equilibrium outcome.

d) Perform the successive elimination of strictly dominated strategies and comment on the result.

Problem 10

a) Prove Proposition 24.
b) Prove the statement discussed two paragraphs later. Given a game $(N, S_i, u_i, i \in N)$ with finite strategy sets, we write $S_N^\infty = \cap_{t=1}^\infty S_N^t$ for the result of the successive elimination of strictly dominated strategies. Consider any finite decreasing sequence $R_N^t \subseteq S_N$ such that $R_N^0 = S_N$ and $R_i^{t+1} \subseteq \mathcal{U}_i(R_N^t)$ for all i and t. Then show that $(R_N^T)^\infty = S_N^\infty$.
c) Prove Proposition 31.

Problem 11

There are 10 locations with values $0 < a_1 < a_2 < ... < a$. Player i $(i = 1.2)$ has $n_i < 10$ soldiers and must allocate them among the locations (no more then one soldier per location). The payoff at location p is a_p to the player whose soldier is unchallenged, and $-a_p$ to his opponent; if they both have a soldier at location p, or no one does, the payoff is 0. The total payoff of the game is the sum of all locational payoffs.
Show that this game has a unique equilibrium in dominant strategies. What if some a_p are equal?

Problem 12 *price competition*

The two firms have constant marginal cost $c_i, i = 1, 2$ and no fixed cost. They sell two substitutable commodities and compete by choosing a price $s_i, i = 1, 2$.

The resulting demands for the 2 goods are

$$D_i(s) = (\frac{s_j}{s_i})^{\alpha_i}$$

where $\alpha_i > 0$. Show that there is an equilibrium in dominant strategies and discuss its stability.

Problem 13 *examples of best reply dynamics*
a) We have a symmetric two player game with $S_i = [0, 1]$ and the common best reply function

$$br(s) = \min\{s + \frac{1}{2}, 2 - 2s\}$$

Show that we have three Nash equilibria, all of them locally unstable, even for the sequential dynamics.
b) We have three players, $S_i = \mathbb{R}$ for all i, and the payoffs

$$u_1(s) = -s_1^2 + 2s_1s_2 - s_2^2$$

$$u_2(s) = -9s_2^2 + 6s_2s_3 - s_3^2$$

$$u_3(s) = -16s_1^2 - 9s_2^2 - s_3^2 + 24s_1s_2 - 6s_2s_3 + 8s_1s_3$$

Show there is a unique Nash equilibrium and compute it. Show the sequential best reply dynamics where players repeatedly take turns in the order $1, 2, 3$ does not converge to the equilibrium, whereas the dynamics where they repeatedly take turns in the order $2, 1, 3$ does converge from any initial point. What about the myopic adjustment where each player uses his best reply at each turn?

Problem 14 *stability analysis in two symmetric games*
a) This symmetrical n-person game has the strategy set $S_i = [0, +\infty[$ for all i and the payoff function

$$u_1(s) = s_2s_3 \cdots s_n(s_1e^{-(s_1+s_2+\cdots+s_n)} - 1)$$

(other payoffs deduced by the symmetry of the game).
Find all dominated strategies if any, and all Nash equilibria (symmetric or not) in pure strategies. Is this a potential game? Discuss the stability of the best reply and sequential best reply dynamics in this game.
b) Answer the same questions as in a) for the following symmetric game with the same strategy sets:

$$u_1(s) = s_2s_3 \cdots s_n(2e^{-(s_1+s_2+\cdots+s_n)} + s_1)$$

Problem 15
Consider the following N players game. The set of pure strategies for each player is $C_i = \{1, ..., N\}$, thus the game consists in each player announcing (simultaneously and independently) an integer between 1 and N. To each pair of players i, j corresponds a number $v_{ij}(= v_{ji})$, interpreted as the utility both players could derive from being together (note that v_{ij} can be negative). Players

are together if and only if they announce the same number. Thus, the payoff to each player i is the sum of v_{ij} over all players j who announced the same number as i. Prove that this game is a potential game and find all Nash equilibria.

Problem 16 *Congestion (variant of Example 3)*

We have n agents who travel from A to B at the same time. Agent i can use a private road at cost $c_i = i$ (that does not depend upon other agents' actions), or use the public road. If k agents travel on that road, they each pay a congestion cost k.

a) Describe the Nash equilibrium outcome (or outcomes) of this game.

b) Is this (are these) equilibrium outcome (s) Pareto optimal? Does it maximize total surplus? If not, compute the fraction of the efficient surplus wasted in equilibrium.

c) Show this is a potential game. Discuss the stability of the equilibrium outcome(s). Is the game dominance solvable (strictly or weakly)?

d) Now the private costs c_i are arbitrary numbers s.t. $1 \leq c_i \leq n$. Answer questions a) and c) above.

Problem 17 *Cost sharing*

We have n agents labeled $1, 2, \cdots, n$, who want to send a signal from A to B. Agent i can send her message via a private carrier at cost $c_i = \frac{1}{i}$ (independently of other agents' choices), or use the public link. If k agents use the public link, they each pay $\frac{1}{k}$.

a) Show that there is one Pareto inferior Nash equilibrium outcome and one Pareto optimal one. Show that the game is a potential game. Discuss the stability of these equilibrium outcomes.

c) Variant: the public link costs $\frac{1+\varepsilon}{k}$ to each user, where ε is a small positive number. Show that the game is now strictly dominance solvable. Compute the inefficiency loss, i.e., the ratio of the total cost in equilibrium to the efficient (minimal) cost of sending all messages.

e) Now the private costs c_i are arbitrary numbers s.t. $0 \leq c_i \leq 1$. Find the Nash equilibrium or equilibria, discuss their efficiency and whether the game is a potential game, or is dominance solvable.

Problem 18 *more congestion games*

We generalize the congestion games of Example 21. Now each player chooses among *subsets* of a fixed finite set S, so that $s_i \subset 2^S$. The same congestion function $f_x(m)$ applies to each element x in S. The payoff to player i is

$$u_i(s) = \sum_{x \in s_i} f_x(n_x(s)) \text{ where } n_x(s) = |\{j \in N | x \in s_j\}|$$

Interpretation: each commuter chooses a different route (origin and destination) on a common road network represented by a non oriented graph. Her own delay is the sum of the delays on all edges of the network.

Show that this game is still a potential game.

Problem 19 *A different congestion game*

There are m men and n women who must choose independently which one of two discos to visit. Let n_a, n_b be the number of women choosing to visit respectively

disco A and disco B, and define similarly m_a, m_b. Each player only cares about the number of visitors of the opposite gender at the disco he or she visits.
a) Assume first the following payoff functions:

$$u_i = n_x \text{ if } i \text{ is a man choosing disco } X; v_j = m_x \text{ if } j \text{ is a woman choosing disco } X$$

Discuss the Nash equilibria of the game and their stability (strong and weak). It will help to show first that this game is a potential game.
b) Now the strategies of the $m + n$ players are the same but the payoffs are:

$$u_i = n_x \text{ if } i \text{ is a man choosing disco } X; v_j = -m_x \text{ if } j \text{ is a woman choosing disco } X$$

In other words men want to be in the disco with more women, while women seek the disco with fewer men (remember this is a theoretical example).
Discuss the Nash equilibria of the game and their stability (strong and weak). Show that this game is **not** a potential game.

Problem 20 *ordinal potential games*
Let σ be the sign function $\sigma(0) = 0, \sigma(z) = 1$ if $z > 0, = -1$ if $z < 0$. Call a game $\mathcal{G} = (N, S_i, u_i, i \in N)$ an *ordinal potential game* if there exists a real valued function P defined on S_N such that for all i and $s_{-i} \in S_{-i}$ we have

$$\sigma\{u_i(s_i, s_{-i}) - u_i(s'_i, s_{-i})\} = \sigma\{P(s_i, s_{-i}) - P(s'_i, s_{-i})\} \text{ for all } s_i, s'_i \in S_i$$

a)Show that the following Cournot oligopoly game is an ordinal potential game. Firm i chooses a quantity s_i, and D^{-1} is the inverse demand function. Costs are linear and identical:

$$u_i(s) = s_i D^{-1}(s_N) - cs_i \text{ for all } i \text{ and all } s$$

b) Show that Proposition 33 still holds for ordinal potential games.

Problem 21 *Cournot duopoly with increasing or U-shaped returns*
In all three questions the duopolists have identical cost functions C.
a) The inverse demand is $D^{-1}(q) = (150 - q)_+$ and the cost is

$$C(q) = 120q - \frac{2}{3}q^2 \text{ for } q \leq 90; = 5,400 \text{ for } q \geq 90$$

Show that we have three equilibria, two of them strongly locally stable.
b) The inverse demand is $D^{-1}(q) = (130 - q)_+$ and the cost is

$$C(q) = \min\{50q, 30q + 600\}$$

Compute the equilibrium outcomes and discuss their stability.
c)The inverse demand is $D^{-1}(q) = (150 - q)_+$ and the cost is

$$C(q) = 2,025 \text{ for } q > 0; = 0 \text{ for } q = 0$$

Show that we have three equilibria and discuss their stability.

Problem 22 *Cournot oligopoly with linear demand and costs*

The inverse demand for total quantity q is

$$D^{-1}(q) = \overline{p}(1 - \frac{q}{\overline{q}})_+$$

where $\overline{p}$ is the largest feasible price and $\overline{q}$ the supply at which the price falls to zero. Each firm i has constant marginal cost c_i and no fixed cost.
a) If all marginal costs c_i are identical, show there is a unique Nash equilibrium, where all n firms are active if $\overline{p} > c$, and all are inactive otherwise.
b) If the marginal costs c_i are arbitrary and $c_1 \leq c_2 \leq \cdots \leq c_n$, let m be zero if $\overline{p} \leq c_1$ and otherwise be the largest integer such that

$$c_i < \frac{1}{m+1}(\overline{p} + \sum_1^i c_k)$$

Show that in a Nash equilibrium outcome, exactly m firms are active and they are the lowest cost firms.

Problem 23 *Hoteling competition in location*
The consumers are uniformly spread on $[0,1]$, and each wants to buy one unit. Each firm charges the fixed exogenous price p and chooses its location s_i in the interval. Production is costless. Once locations are fixed, each consumer shops in the nearest store. The tie-breaking rule: the demand is split equally between all stores choosing the same location
a) Show that with two competing stores, the unique Nash equilibrium is that both locate in the center. Show the game is not dominance-solvable. However, it is dominance solvable if each firm must locate in one of the $n+1$ points $0, \frac{1}{n}, \frac{2}{n}, \cdots, 1$.
b) Show that with three competing stores, the game has no Nash equilibrium.
c) Show that with four competing stores, the game has a Nash equilibrium. Is it unique?
d) What is the situation with five stores?

Problem 24 *Hoteling competition in location: probabilistic choice*
a) Two stores choose a location on the interval $[0,100]$. Customers are uniformly distributed on this interval, with at most a unit demand, and will shop from the nearest store if at all. If the distance between a customer and the store is t, he will buy with probability $p(t) = \frac{2}{\sqrt{t+4}}$. Thus if a store is located at 0 and is the closest store to all customers in the interval $[0,x]$, it will get from these customers the revenue

$$r(x) = \int_0^x p(t)dt = 4\sqrt{x+4} - 8$$

Stores maximize their revenues. Analyze the competition between the two stores and compute their equilibrium locations. Compare them to the collusive outcome, namely the choice of locations maximizing the total revenue of the two stores.

b) Generalize the model of question a). Now $p(t)$ is unspecified and so is its primitive $r(t)$. We assume that p is continuous, strictly positive, and strictly decreasing from $p(0) = 1$.
Under what condition on p do both stores locate at the midpoint in the Nash equilibrium of the game?
Show that if in equilibrium the stores choose different locations, they will never locate on $[0, 25]$ or $[75, 100]$.

Problem 25 *Hoteling competition in prices: two firms*
The 1000 consumers are uniformly spread on $[0, 3]$ and each wants to buy one unit and has a very large reservation price. The two firms produce costlessly and set arbitrary prices s_i. Once these prices are set consumers shop from the cheapest firm, taking into account the unit transportation cost t. A consumer at distance d_i from firm i buys

from firm 1 if $s_1 + td_1 < s_2 + td_2$, from firm 2 if $s_1 + td_1 > s_2 + td_2$

(the tie-breaking rule does not matter)
a) If the firms are located at 0 and 3, show that there is a unique Nash equilibrium pair of prices. Is it strongly/sequentially stable?
b) If the firms are located at 1 and 2, show that there is no Nash equilibrium (*hint: check first that a pair of two different prices can't be an equilibrium*).

Problem 26 *Hoteling competition in prices: three firms*
The consumers are uniformly spread over the interval $[0, 3]$ and each wants to buy one unit of the identical good produced by the three firms. The firms are located respectively at $0, 1$ and 3 and they produce costlessly. The transportation cost is 1 per unit. As usual consumers shop at the firm where the sum of the price and the transportation cost is smallest.
a) Write the strategic form of the game where the three firms choose the prices s_1, s_2, s_3 respectively.
b) Show that the game has a unique Nash equilibrium and compute it.
c) Is the equilibrium computed in b) strongly/sequentially stable?.

Problem 27 *price war*
Two duopolists (a la Bertrand) have zero marginal cost and capacity c. The demand d is inelastic, with reservation price $\overline{p}$. Assume $c < d < 2c$. We also fix a small positive constant ε ($\varepsilon < \frac{\overline{p}}{10}$).
The game is defined as follows. Each firm chooses a price $s_i, i = 1, 2$ such that $0 \leq s_i \leq \overline{p}$. If $s_i \leq s_j - \varepsilon$, firm i sells its full capacity at price s_i and firm j sells $d - c$ at price s_j. If $|s_i - s_j| < \varepsilon$ the firms split the demand in half and sell at their own price (thus ε can be interpreted as a transportation cost between the two firms). To sum up

$$\begin{aligned} u_1(s) &= cs_1 \text{ if } s_1 \leq s_2 - \varepsilon \\ &= (d - c)s_1 \text{ if } s_1 \geq s_2 + \varepsilon \\ &= \frac{d}{2}s_1 \text{ if } s_2 - \varepsilon < s_1 < s_2 + \varepsilon \end{aligned}$$

with a symmetric expression for firm 2.
Set $p^* = \frac{d-c}{c}\overline{p}$ and check that the best reply correspondence of firm 1 is

$$\begin{aligned} br_1(s_2) &= \overline{p} \text{ if } s_2 < p^* + \varepsilon \\ &= \{\overline{p}, p^*\} \text{ if } s_2 = p^* + \varepsilon \\ &= s_2 - \varepsilon \text{ if } s_2 > p^* + \varepsilon \end{aligned}$$

Show that the game has no Nash equilibrium, and that the sequential best reply dynamics describes a cyclical price war.

Problem 28 *Bertrand duopoly*
The firms sell the same commodities and have the same cost function $C(q)$, that is continuous and increasing. They compete by setting prices $s_i, i = 1, 2$. The demand function D is continuous and decreasing. The low price firm captures the entire demand; if the 2 prices are equal, the demand is equally split between the 2 firms. Hence the profit function for firm 1

$$u_1(s) = s_1 D(s_1) - C(D(s_1)) \text{ if } s_1 < s_2; = 0 \text{ if } s_1 > s_2$$

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

and the symmetrical formula for firm 2.
a) Show that if s^* is a Nash equilibrium, then $s_1^* = s_2^* = p$ and

$$AC(\frac{q}{2}) \leq p \leq 2AC(q) - AC(\frac{q}{2})$$

where $q = D(p)$ and $AC(q) = \frac{C(q)}{q}$ is the average cost function.
b) Assume increasing returns to scale, namely AC is (strictly) decreasing. Show there is no Nash equilibrium $s^* = (p, p)$ where the corresponding production q is positive. Find conditions on D and AC such that there is an equilibrium with $q = 0$.
c) In this and the next question assume decreasing returns to scale, i.e., AC is (strictly) increasing. Show that if $s^* = (p, p)$ is a Nash equilibrium, then $p_- \leq p \leq p_+$ where p_- and p_+ are solutions of

$$p_- = AC(\frac{D(p_-)}{2}) \text{ and } p_+ = 2AC(D(p_+)) - AC(\frac{D(p_+)}{2})$$

Check that the firms have zero profit at (p_-, p_-) but make a positive profit at (p_+, p_+) if $p_- < p_+$. *Hint: draw on the same figure the graphs of* $D^{-1}(q), AC(\frac{q}{2})$ *and* $2AC(q) - AC(\frac{q}{2})$.
d) To prove that the pair (p_+, p_+) found in question c) really is an equilibrium we must check that the revenue function $R(p) = pD(p) - C(D(p))$ is non decreasing on $[0, p_+]$. In particular p_+ should not be larger than the monopoly price.

Assume $C(q) = q^2$, $D(p) = (\alpha - \beta p)_+$ and compute the set of Nash equilibrium outcomes, discussing according to the parameters α, β.

Problem 29

In the game $\mathcal{G} = (N, S_i, u_i, i \in N)$ we write

$$\alpha_i = \max_{s_i} \min_{s_{-i}} u_i(s_i, s_{-i}); \beta_i = \min_{s_{-i}} \max_{s_i} u_i(s_i, s_{-i})$$

and assume the existence for each player of a prudent strategy $\overline{s}_i$, namely $\alpha_i = \min_{s_{-i}} u_i(\overline{s}_i, s_{-i})$.
a) Assume $\alpha = (\alpha_i)_{i \in N}$ is a Pareto optimal utility profile: there exists $\widetilde{s} \in S_N$ such that

$$\alpha = u(\widetilde{s}) \text{ and for all } s \in S_N : \{u(s) \geq u(\widetilde{s})\} \Rightarrow u(s) = u(\widetilde{s})$$

Show that $\alpha = \beta$ and that any profile of prudent strategies is a Nash equilibrium. Then we speak of an *inessential* game.
b) Assume that the strategy sets S_i are all finite, and $\beta = (\beta_i)_{i \in N}$ is a Pareto optimal utility profile. Show that if each function u_i is one-to-one on S_N then the outcome $\widetilde{s}$ such that $\beta = u(\widetilde{s})$ is a Nash equilibrium. Give an example of a game with finite strategy sets (where payoffs are not one-to-one) such that β is Pareto optimal and yet the game has no Nash equilibrium.