Advanced Game Theory

Herve Moulin Baker Hall 263; moulin@rice.edu

Rice University ECO 440 Spring 2008

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".

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.■

Examples:

- *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.

- 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 <u>optimal</u>—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.*

Examples:

- *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 $n \neq m$.

- *Chess* and Zermelo's theorem: the game of Chess has three payoffs, $+1, -1, 0$. Although we do which one, one of these 3 numbers is the value of the game, i.e., either Win can guarantee a win, or Black can, or both can secure a draw.

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

1) a tree, with a particular node taken as the origin;
2) for each non-terminal node, a specification of who has the move;
3) for each terminal node, a payoff attached to it.

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

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

For each non-terminal node n, $n \in N_i$ $(i = 1, 2)$ means that player i has the move at this node. A move consists of picking a successor to this node.

The game starts at the origin n_0 of the tree and continues until some terminal node n_t is reached. Then the payoff $u(n_t)$ attached to this node is realized (i.e., player 1 gains $u(n_t)$ and player 2 looses $u(n_t)$).

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

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

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

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

Proof:

The proof is by induction in the length l of the tree Γ. For $l = 1$ the theorem holds trivially, since it is a one-person one-move game (say, player 1 is to choose a move at n_0, and any of his moves leads to a terminal node). Thus, a prudent strategy for the player 1 is a move which gives him the highest payoff, and this payoff is the value of the game.

Assume now that the theorem holds for all games of length at most $l - 1$, and consider a game G of length l. Without loss of generality, $n_0 \in N_1$, i.e., player 1 has a move at the origin.

Let $\{n_1, ..., n_k\}$ be the set of successors of the origin n_0. Each subtree Γ_i, with the origin n_i, is of length $l-1$ at most. Hence, by the induction hypothesis, any subgame G_i associated with a Γ_i has a value, say, m_i. We claim that the value of the original game G is $m = \max\limits_{1 \leq i \leq k} m_i$.

Indeed, by moving first to n_i and then playing optimally at G_i, player 1 can guarantee himself at least m_i. Thus, player 1 can guarantee that he will gain at least m in our game G. But, by playing optimally in each game G_i, player 2 can guarantee herself the loss of not more than m_i. Hence, player 2 can guarantee that she will lose at most m in our game G. Thus max-min and min-max payoffs coincide and m is the value of the game G.■

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

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

1.4 Mixed strategies

Motivating examples:

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.

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.

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

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 $Y = \{(y_1, ..., y_m) : \sum_{i=1}^m y_j = 1, y_j \geq 0\}$.

Claim 11 *When player 1 chooses $s \in X$ and player 2 chooses $y \in Y$, the expected payoff of the game is equal to $s^T Ay$.*

Proof: $s^T Ay = (s_1, ..., s_n) \begin{pmatrix} a_{11} & ... & a_{1m} \\ ... & ... & ... \\ a_{n1} & ... & a_{nm} \end{pmatrix} \begin{pmatrix} y_1 \\ ... \\ y_m \end{pmatrix} = \sum_{i=1}^n \sum_{j=1}^m s_i a_{ij} y_j,$
and each element of this double sum is $s_i a_{ij} y_j = a_{ij} s_i y_j =$
$a_{ij} \times$Pro[1 chooses i]$\times$Pro[2 chooses j] $= a_{ij} \times$Pro[1 chooses i and 2 chooses 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 $v_1(s) = \min\limits_{y \in Y} s^T Ay$ be the minimum payoff player 1 can get if he chooses to play s. Then $v_1 = \max\limits_{s \in X} v_1(s) = \max\limits_{s \in X} \min\limits_{y \in Y} s^T Ay$ is the secure utility level for player 1.

Similarly, we define $v_2(y) = \max\limits_{s \in X} s^T Ay$, and $v_2 = \min\limits_{y \in Y} v_2(y) = \min\limits_{y \in Y} \max\limits_{s \in X} s^T Ay$, the secure utility level for player 2.

Claim 12 *The number $s^T Ay$ can be viewed as 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).*

Proof:

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

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 $A_{\cdot j} = Ae^j$.

Now, $s^T A e^j$ is the expected payoff to player 1, when he uses (mixed) strategy s and player 2 uses (pure) strategy e^j. Hence, $s^T A y = \sum_{j=1}^{m} y_j \left[s^T A e^j\right]$ is a weighted average of player 1's payoffs against pure strategies of player 2 (when player 1 uses strategy s). In this weighted sum, weights y_j are equal to the probabilities that player 2 would choose these pure strategies e^j.

Given this claim, $v_1(s) = \min_{y \in Y} s^T A y$, the minimum of $s^T A y$, will be attained at some pure strategy j (i.e., at some $e^j \in Y$). Indeed, if $s^T A e^j > v_1(s)$ for all j, then we would have $s^T A y = \sum y_j \left[s^T A e^j\right] > v_1(s)$ 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 13 *(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 $K_v = \{z \in \mathbb{R}^n : s \cdot z \leq v$ 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, $v_1 = \max\limits_{s \in X} \min\limits_{y \in Y} s \cdot Ay \geq \min\limits_{y \in Y} \overline{s} \cdot Ay \geq \overline{v}$ (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, $v_2 = \min\limits_{y \in Y} \max\limits_{s \in X} s \cdot Ay \leq \max\limits_{s \in X} s \cdot A\overline{y} = \max\limits_{s \in X} s \cdot \overline{z} = \max\limits_{i=1,\cdots,n} \overline{z_i} \leq \overline{v}$ (since $\overline{z} \in K_{\overline{v}}$).

We obtain that $v_2 \leq \overline{v} \leq v_1$. Together with the fact that $v_1 \leq v_2$, it gives us $v_2 = \overline{v} = v_1$, the desired statement.

Note also, that the maximal value of $v_1(s)$ 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 strategies

How can we find the maximin (mixed)strategy $\overline{s}$, the minimax (mixed) 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 $\overline{s}^T A e^j \geq \min\limits_{y \in Y} \overline{s}^T A y = v_1(\overline{s}) = v_1 = v$, it follows that $v = \overline{s}^T A \overline{y} = \overline{y}_1 \left[\overline{s}^T A e^1\right] + ... + \overline{y}_m \left[\overline{s}^T A e^m\right] \geq \overline{y}_1 v + ... + \overline{y}_m v = v\left(\overline{y}_1 + ... + \overline{y}_m\right) = v$, and the equality implies $\overline{s}^T A_{\cdot j} = \overline{s}^T A e^j = v$ for all j such that $\overline{y}_j \neq 0$. Thus, player 2 receives her minimax value $v_2 = v$ by playing against $\overline{s}$ 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 can often simplify the task of computing the supports of optimal mixed strategies by successively eliminating *dominated* rows and columns.

Definition 14 *We say that i-th row of a matrix A dominates its k-th row, if $a_{ij} \geq a_{kj}$ for all j and $a_{ij} > a_{kj}$ for at least one j. Similarly, we say that j-th column of a matrix A dominates its l-th column, if $a_{ij} \geq a_{il}$ for all i and $a_{ij} > a_{il}$ for at least one 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 15 *If rows $i_1, ..., i_k$ of a matrix A are dominated, then player 1 has an optimal strategy $\overline{s}$ such that $\overline{s}_{i_1} = ... = \overline{s}_{i_k} = 0$; moreover, any optimal strategy for the game obtained by removing dominated rows from A will also be an optimal strategy for the original game. The same is true for dominated columns and player 2.*

Given this, we can proceed as follows. 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 $A = \begin{pmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{pmatrix}$ and this game does not have saddle point. 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, it can be seen that $a_{11}s_1 + a_{21}s_2 = v$ 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 thus have the linear system with 6 equations and 5 variables s_1, s_2, y_1, y_2 and v. 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

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

1.5.2 2×n games

By fiocusing on the player who has two strategies, one computes the value as the solution of a tractable liner program. See the examples in class.

1.5.3 Symmetric games

The game with matrix A is symmetric if $A = -A^T$ (Exercise:check this). Like in a general 2 person zero-sum game, the value of a symmetric game is zero. 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: *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, $\max_{x \in X} \min_{y \in Y} u(x, y) = -1 < +1 = \min_{y \in Y} \max_{x \in X} u(x, y)$.

An important result, known as Glicksberg Theorem, states that it will 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.

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: *Mixed strategies in the silent gunfight*
In the silent gunfight (Problem 5; see also the noisy version 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. We *guess* 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 *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:

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

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

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. It 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 $t \in T$*

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

Example: *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:

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

Find all saddle points.

Problem 3. *Gale's roulette*

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

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

Find the value and optimal strategies of this game

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

Problem 4 *Land division game.*

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

Problem 5 *Silent gunfight*
Now the duellists cannot hear when the other player shoots. Payoffs are computed in the same way. If v is the value of the *noisy* gunfight, show that in the silent version, the values $\overline{m} = \min\max$ and $\underline{m} = \max\min$ are such that $\underline{m} < v < \overline{m}$.

Problem 6.1
Two players move in turn and the one who cannot move loses. Find the winner (1-st or 2-nd player) and the winning strategy.

In questions a) and b), both players move the same piece.
a) A castle stays on the square a1 of the 8×8 chess board. A move consists in moving the castle according to the chess rules, but only in the directions up or to the right.
b) The same game, but with a knight instead of a castle.

In questions c) and d), a move consists of adding a new piece on the board.
c) A move consists in placing a castle on the 8 by 8 chess board in such a way, that it does not threatens any of the castles already present.
d) The same game, but bishops are to be placed instead of castles.

Problem 6.2
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.
a) The initial position is 11111110111110111101, where a 1 is a match and 0 an empty space. Players successively remove one match or three adjacent matches. Who wins if the player removing the last match loses ? if the player removing the last match wins?
b) 20 coins are placed on the table in a chain (such that they touch each other), so that they form either a straight line, or a circle. A move consists in taking either one or two adjacent (touching) coins. Two versions: whoever removes the last coin wins, or that person loses.
c) The game starts with two piles, of respectively 20 and 21 coins. A move consists in taking one pile away and dividing the other into two nonempty piles. Two versions: the position 1,1 is losing; or it is winning. Generalization: now the two piles are of sizes n and m.
d) From a pile of n coins, the players take turns to remove *one* or *four* coins. Solve the same two versions as above.
h) From two piles of sizes n and m, the players take turns removing either *one* or *two* coins from a pile of their choice. Solve the usual two versions.

Problem 6.3
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:

- $n = 2, m = 4$

- $n = 3, m = 3$
- $n = 4, m = 4$
- n and m even
- n even, m odd

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.

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 *Stone, Paper, Scissors and Well*

The paper is cut by (loses to) the scissors, it wraps (beats) the stone and closes (beats) the well. The scissors break on the stone and fall into the well (loses to both). The stone falls into the well.
Solve the game where both chose one of these 4 strategies simultaneously and the loser pays \$1 to the winner. Discuss pure and mixed strategies.

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 *Catch me*
Player 1 chooses a location x in $[0, 1]$ and player 2 chooses simultaneously a location y. Player1 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.

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 *decentralized 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;
- 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.

Decentralized scenarios are well suited to games involving a large number of players, each one with a relatively small influence on the overall outcome (competitive context). Coordination scenarios are more natural in games with a small number of participants.

This chapter is long on examples and short on abstract proofs (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. 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 *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 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 not only the best reply behavior may not converge at all, or if it converges, the limit equilibrium outcome may well be inefficient (Pareto inferior). Decentralized behavior may diverge, or it may converge toward a socially suboptimal outcome.

Definition 21 *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 22 *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, also called an improvement path, 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 23 *Assume the strategy sets S_i of each player are compact and the payoff functions u_i are continuous. 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!

Note that the topological assumptions in the Proposition hold true if the strategy sets are finite.

Definition 24 *We call a Nash equilibrium s* ***strongly globally 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 ***strongly locally stable*** *if for any neighborhood $\mathcal{N}$ of s in S_N there is a sub-neighborhood $\mathcal{M}$ of s such that any best reply dynamics starting in $\mathcal{M}$ stays in $\mathcal{N}$.*

We call a Nash equilibrium ***weakly globally 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 call a Nash equilibrium ***weakly locally stable*** *if for any neighborhood $\mathcal{N}$ of s in S_N there is a sub-neighborhood $\mathcal{M}$ of s such that any sequential best reply dynamics starting in $\mathcal{M}$ stays in $\mathcal{N}$.*

Note that if strategy sets are finite, the concept of local stability (in both versions) has no bite (every equilibrium is strongly locally stable).

2.1.1 stable and unstable equilibria

We give a series of examples illustrating these definitions. The actual analysis of each game is done in class.

Example 1: *two-person zero sum games*
Here a Nash equilibrium is precisely a saddle point. In the following game, a saddle point exists and is globally 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 weakly stable:

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

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 25 *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 weakly stable, and this happens if and only if f has a unique fixed point and no cycle of length 2 or more.*

Proof: Assume both br_1 and br_2 are single-valued. If s^* is weakly stable, any best reply dynamics starting from $s \in \{br_1(S_2) \times S_2\} \cup \{S_1 \times br_2(S_2)\}$ must converge to s^*; clearly any best reply dynamics reaches this set in one step, therefore s^* is globally stable. Details of the proof are the subject of problem 2.1.

Example 2 *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 globally stable; if $\frac{b}{c} < 1$ it is not strongly stable yet weakly globally stable.

Example 3: *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

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

Several possible interpretations. Vaccination: strategy 1 is to take the vaccine, strategy 0 to avoid it. 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. Traffic game: 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.

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 $a(1) \geq b(1)$. If the equilibrium is unique it is weakly (globally) stable but not strongly stable.

The strategy sets are finite so local stability has no bite. However when the number of agents is large we can measure the deviation from an equilibrium by the number of agents who are not playing the equilibrium strategy. This leads to the concept of *local stability in population.* For a Nash equilibrium s^* this requires that for any parameter $\lambda, 0 < \lambda < 1$ 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 fo the players have changed from the original equilibrium.

In example 3, an equlibrium is locally stable in population if a cuts b from above, and unstable if a cuts b from below.

2.1.2 potential games

We introduce in the next three subsections three classes of games where some form of stability is guaranteed. Potential games generalize the pure coordination games where all players have the same payoff functions. As shown in example 2 above, in such games strong stability is problematic but weak stability is not.

Definition 26 *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 equilibiurm (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, the best reply dynamics has very appealing stability properties.

Proposition 27 *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.*

i) If there is a unique Nash equilibrium, it is weakly globally stable.

ii) If a Nash equilibrium of $\mathcal{G}$ is a local maximum, locally unique, of the potential P, this equilibrium is weakly locally stable.

The simplest examples of a potential game are the coordination games, where the potential is the common utility. Our next example shows why, even in this simple case, strong stability is out of reach.

Example 4 *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 globally stable, even weakly. The single exceptions are $z^* = 0$ or 1, because then the Nash equilibrium is unique, and globally stable. For an arbitrary z^*, we can still say that the *game* is weakly globally stable in *utilities*, because along any 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 globally stable, because myopic best reply sequences cycle around z^* without reaching it.

Example 5 *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 $\mathbb{R}_+^N$.

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 the game 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 15 and problem 7 for two alternative choices of B.

Example 6 *congestion games*

These games generalize both examples 3 and 4. 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. It is easy to think of examples where f is single-peaked, as when we choose a restaurant.

Here the potential function is

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

Example 7: *the Braess paradox*

There are two roads to go from A to B, and 6 commuters want to do just that. 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 m in the following way:

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

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. We now add one more link on the road network, directly between C and D, with travel time $10 + m$. 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!

In view of example 6, these two games are potential games. One checks that the potential function has a unique coordinate-wise maximum, so that this unique Nash equilibrium is weakly globally stable.

2.1.3 strictly dominance-solvable games

Definition 28 *In the game in strategic form $\mathcal{G} = (N, S_i, u_i, i \in N)$, we say that player i's strategy s_i is strictly dominated by his strategy 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 $\mathcal{U}_i(T_N)$ for the set of player i's strategies in the restricted game $\mathcal{G}(T_N) = (N, T_i, u_i, i \in N)$ that are not strictly dominated.

Definition 29 *We say that the game $\mathcal{G}$ is strictly dominance-solvable if the sequence defined inductively by*

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

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

$$\cap_{t=1}^{\infty} S_N^t = \{s^*\}$$

Proposition 30 *Under this assumption, s^* is the single Nash equilibrium outcome of the game, and it is strongly globally stable.*

Even if the successive elimination of strictly dominated strategies does not converge to a singleton, it does simplify the search for Nash equilibria, because of the following important fact, of which the easy proof is omittted:

the set $\cap_{t=1}^{\infty} S_N^t$ contains all Nash equilibria of the game
and all limit points of the best reply dynamics

The computation of the limit set $\cap_{t=1}^{\infty} S_N^t$ is often simplified by the fact that if we perform successive rounds of partial elimination of strictly dominated strategies, for instance by eliminating only the strictly dominated strategies of a single player, we reach eventually the same limit set $\cap_{t=1}^{\infty} S_N^t$ (provided the initial strategy sets are finite). We discuss this property in section 2.2.2, where it is contrasted with the lack of robustness of the successive elimination of weakly dominated strategies.

Example 8 *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 $\frac{2}{3}\overline{s}$, while strategy 1 is at most $\frac{1}{3}$ away from $\frac{2}{3}\overline{s}$.

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 9 *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.

2.1.4 games with increasing best reply

A class of games closely related to dominance-solvable games consist of those where the best reply functions (or correspondences) are non decreasing. By way of illustration consider a symmetric game where $S_i = [0, 1]$ and the (symmetric) best reply function $s \to br(s, \cdots, s)$ is non decreasing. This function must cross the diagonal; it provides a simple illustration of the first statement in the next result.

Proposition 31 *Let the strategy sets S_i be either finite, or real intervals $[a_i, b_i]$, and the utility functions u_i be continuous. Assume the best reply functions in the game $\mathcal{G} = (N, S_i, u_i, i \in N)$ are single valued and non decreasing*

$$s_{-i} \leq s'_{-i} \Rightarrow br_i(s_{-i}) \leq br_i(s'_{-i}) \text{ for all } i \text{ and } s_{-i} \in S_{-i}$$

Then the game has a smallest Nash equilibrium outcome s_- and s_+ a largest one s_+. Any best reply dynamics starting from a converges to s_-; any best reply dynamics starting form b converges to s_+.

Proposition 32 *Say that the payoff functions u_i satisfy the single crossing property if for all i and all $s, s' \in S_N$ such that $s \leq s'$ we have*

$$u_i(s'_i, s_{-i}) > u_i(s_i, s_{-i}) \Rightarrow u_i(s'_i, s'_{-i}) > u_i(s_i, s'_{-i})$$

$$u_i(s'_i, s_{-i}) \geq u_i(s_i, s_{-i}) \Rightarrow u_i(s'_i, s'_{-i}) \geq u_i(s_i, s'_{-i})$$

Under the SC property, define br_i^- *and* br_i^+ *to be respectively the smallest and largest element of the best reply correspondence. They are both non-decreasing. The sequences* s_-^t *and* s_+^t *defined as*

$$s_-^0 = a; s_-^{t+1} = br_i^-(s_-^t); s_+^0 = b; s_+^{t+1} = br_i^+(s_+^t)$$

are respectively non decreasing and non increasing, and they converge respectively to the smallest Nash equilibrium s_- *and to the largest one* s_+*. Finally the successive elimination of strictly dominated strategies converges to* $[s_-, s^+]$

$$\cap_{t=1}^{\infty} S_N^t = [s_-, s_+]$$

In particular if the game has a unique equilibrium outcome, it is strictly dominance-solvable.

Note that if u_i is twice differentiable the SCP holds if and only if

$$\frac{\partial^2 u_i}{\partial s_i \partial s_j} \geq 0 \text{ on } [a, b].$$

Example 5 *Voluntary contribution to a public good (continued)*
In Example 5 we assume that B is convex. Then the game has the SC prtoperty, therefore all the properties spelled above apply. As the game is also a potential game, we conclude that it is strictly dominance solvable if the potential function $P(s) = B(s_N) - \sum_i C_i(s_i)$ has a unique coordinate-wise maximum.

Example 10 *A search game*
Each player exerts effort searching for new partners. The probability that player i finds any other player is $s_i, 0 \leq s_i \leq 1$, and when i and j meet, they derive the benefits α_i and α_j respectively. The cost of the effort is $C_i(s_i)$. Hence the payoff functions

$$u_i(s) = \alpha_i s_i s_{N \diagdown \{i\}} - C_i(s_i) \text{ for all } i$$

Assuming only that C_i is increasing, we find that the game satisfies the single crossing property. The strategy profile $s_- = 0$ is always an equilibrium, and the largest equilibrium s_+ is Pareto superior.

Example 11 *price competition*
Each firm has a linear cost production (set to zero without loss of generality) and chooses a non negative price p_i. The resulting demand and net payoff for firm i are

$$D_i(p) = (A_i - \frac{\alpha_i}{3} p_i^2 + \sum_{j \neq i} \beta_j p_j)_+ \text{ and } u_i(p) = p_i D_i(p)$$

The game has increasing best reply functions. In the symmetric case its equilibrium is unique hence the game is dominance-solvable.

2.2 coordination and Nash equilibrium

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.

Schelling's rendez-vous game illustrates both interpretations.

If a game has multiple Nash equilibria we have a selection problem: under either scenario above, it is often unclear how the players will be able to coordinate on one of them. Then even if a Nash equilibrium is unique, it may be challenged by other strategic choices that are safer or appear so.

On the other hand in dominance-solvable games, selecting the Nash outcome by deduction (covert communication) is quite convincing, and our confidence in the predictive power of the concept remains intact.

2.2.1 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. This fundamental difficulty is illustrated by the two following celebrated games.

Example 12 *crossing game (a.k.a. the Battle of the Sexes)*
Each player must stop or go. The payoffs are as follows

$$\begin{array}{ccc} stop & 1,1 & 1-\varepsilon,2 \\ go & 2,1-\varepsilon & 0,0 \\ & stop & go \end{array}$$

Each player would like to commit to go, so as to force the other to stop. There is a mixed strategy equilibrium as well, but it has its own problems. See Section 3.3.

Example 13 *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.

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 does it first and other players are sure to notice.

A *game of timing* takes the following form. 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.

Example 14 *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 be solved by further arguments of Pareto dominance, or risk dominance. Sometimes the selection problem is facilitated because the players agree on the most favorable equilibrium: the Pareto dominance argument. A simple example is any *coordination game*: if a single outcome maximizes the common payoff it will be selected without explicit comunication. When several outcomes are optimal, we may hope that one of them is more salient, as in Schelling's rendez-vous game.

Finally prudence may point to some particular equilibrium outcome. But this criterion may conflict with Pareto dominance as in Kalai's hat story, and in the following important game.

Example 15 *coordination failure*

This is an example of a public good provision game by voluntary contributions (example 6), 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.

2.2.2 dominance solvable games

Eliminating dominated strategies is the central coordination device performed by independent deductions of completely informed agents.

Definition 33 *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*

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

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

Definition 34 *We say that the game $\mathcal{G}$ is dominance-solvable if the sequence defined inductively by*

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

and called the successive elimination of dominated strategies, converges to a single outcome s^*:

$$\cap_{t=1}^{\infty} {}^wS_N^t = \{s^*\}$$

If the strategy sets are finite, or compact with continuous payoff functions, the set of weakly undominated strategies is non empty.

Notice an important difference between the elimination of strictly versus weakly dominated strategies. As noted in section 2.1.3, the elimination of strictly dominated strategies never "loses" a Nash equilibrium:

$$\{s \text{ is a Nash equilibrium of } \mathcal{G}\} \Rightarrow s \in \cap_{t=1}^{\infty} S_N^t$$

By contrast the elimination of weakly dominated strategies may loose some, or even all, Nash equilibria along the way. Compare the two person game

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

where the algorithm picks the unique equilibrium, to the following example

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

where the algorithm may throw out the baby with the water!

On the other hand if the game reduced to the strategy space $\cap_{t=1}^{\infty}{}^{w}S_N^t$ has a Nash equilibrium, this outcome is an equilibrium of the original game as well (exercise: prove this fact).

Another difference between the two successive elimination algorithms, based on strict or weak domination, is their robustness with respect to partial elimination. Suppose, in the case where we only drop strictly dominated strategies, that at each stage we choose S_i^{t+1} as a subset of $\mathcal{U}_i(S_N^t)$: then it is easy to check that the limit set $\cap_{t=1}^{\infty}S_N^t$ is unaffected (provided we eventually take all elimination opportunities)(exercise: prove this claim). On the other hand when we only drop some weakly dominated strategies at each stage, the result of the algorithm may well depend on the choice of subsets ${}^{w}S_i^{t+1}$ in $\mathcal{WU}_i({}^{w}S_N^t)$. 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. Despite the difficulty above, in many instances the elimination algorithm in Definition 33 leads to a convincing equilibrium selection.

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

Example 16 *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 to vote for her top choice. The two other players can only eliminate the vote for their bottom candidate.

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 a Nash equilibrium outcome. The paradox is that the chair's tie-breaking privilege result in the election of her worst outcome!

Often a couple of rounds of elimination are enough to select a unique Nash equilibrium, even though the elimination algorithm is stopped and the initial game is not (weakly) dominance solvable.

Example 17 *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. 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 18 *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, 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 dominated strategies, we are left with $S_1^2 = [t_2^* - \varepsilon, 1], S_2^2 = [t_2^*, 1]$. Although the elimination process stops there, the outcome of the remaining game[1] is not in doubt: $s_1^* = t_2^* - \varepsilon, s_2^* = t_2^*$.

2.2.3 dominant strategy equilibrium

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

Definition 35 *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.

[1] This game is an inessential game, as discussed in question a) of problem 18.

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. 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. Information about other players' payoffs or moves is worthless, as long as our player is unable to influence their choices (for instance a threat of the kind "if you do this I will do that" is not enforceable).

A game with an equilibrium in dominant strategies is weakly, but not necessarily strictly, dominance-solvable.

A very famous game illustrates the fact that an equilibrium in dominant strategies may not be Pareto optimal:

Example 19 *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 dominant strategy equilibrium is Pareto inferior to the unanimously selfless outcome.

Dominant strategy equilibria do not happen in very many games 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 celebrated examples of such *strategy-proof* allocation mechanisms follow. In both cases the game has a dominant strategy equilibrium in all cases, and the corresponding outcome is efficient (Pareto optimal).

Example 20 *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 that we selected by dominance-solvability in the first price auction (example 17). 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.

It is interesting to note that in the second price auction game, there is a distressing variety of Nash equilibrium outcomes and in particular any player, even the one with the lowest valuation of all, receives the object in some equilibrium. More precisely, 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 21 *voting under single-peaked preferences*
The n players vote to choose an outcome x in $[0,1]$. 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. 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 $[v_i, 1]$.

Here again, 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.

2.3 problems on chapter 2

Problem 1
In Schelling's model (example 3) find examples of the functions a and b such that the equilibrium is unique and strongly globally stable; such that it is unique and weakly but not strongly globally stable.

Problem 2.1 *Example 2 continued*
Assume the function g has a local maximum at z^*, to be precise z^* is the unique maximum of g over the interval $[\frac{n-1}{n}z^*, \frac{n-1}{n}z^* + 1]$. Show that z^* is a Nash equilibrium outcome, and that each corresponding Nash equilibrium strategy profile s^* is weakly locally stable but not strongly locally stable.

Problem 2.2
Complete the proof of Proposition 25, in particular by dealing with the case where the best reply correspondence can be multivalued.

Problem 3 *games of timing*
a) We have two players, a and b both increase, and a intersects b from below. Perform the successive elimination of dominated strategies, and find all Nash equilibria. Can they be Pareto improved?
b) We extend the war of attrition (example 14) 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 4 *Example 18 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 $[t,1]$.

What strategic advice would you give to each player? Discuss the cases where a player knows his opponent's utility and that where she does not.

Problem 5

One hundred people live in the village, of whom 51 support conservative candidate and 49 support liberal candidate. A villager gets payoff +10 if her candidate wins and -10 if her candidate looses. But since voting process is time-consuming, a villager gets -1 simply from the fact that she votes.
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 6 *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 7 *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 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 8

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. Find the potential and Nash equilibria of this game.

Problem 9 *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 search game (example 10) and the symmetric case of the price competition (example 11) are ordinal potential games.
b)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:

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

c) Show that Proposition 27 still holds for ordinal potential games.

Problem 10 *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 11 *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 12 *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

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

Show that each equilibrium is weakly locally stable.
c) Compute strictly dominated strategies for each player. For what profiles (λ_i) is our game (strictly) dominance-solvable?

Problem 13 *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 globally stable.

Problem 14 *more congestion games*
We generalize the congestion games of example 7. 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 15 *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 16
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 17 *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 18 *Cournot duopoly with increasing or U-shaped returns*
In all 3 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 (local) 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 (local) stability.

Problem 19 *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 20 *Hoteling competition in location*
The consumers are uniformly spread on $[0, 1]$, and each wants to buy one unit. Each firm charges the fixed 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 does not matter).
a) Show that with two competing stores, the unique Nash equilibrium is that both locate in the center. Is the game dominance-solvable?
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 21 *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 22 *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

$$\text{from firm 1 if } s_1 + td_1 < s_2 + td_2, \text{ from firm 2 if } s_1 + td_1 > s_2 + td_2$$

(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. Analyze its stability properties.
b) If the firms are located at 1 and 2, show that there is no Nash equilibrium (*hint: check first that a pair of 2 different prices can't be an equilibrium*).

Problem 23 *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 3 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) Discuss the stability of the equilibrium computed in b).

Problem 24 *price war*
Two duopolists (a la Bertrand) have zero marginal cost and capacity c. The demand d is inelsatic, 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 captures a cyclical price war.

Problem 25 *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 26
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.
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.

Problem 27
In the Cournot model of example 3 where we do not assume a large number of agents and do not identify best reply behavior and competitive behavior, show the Nash equilibrium is unique and describe it.

3 Existence results and mixed strategies

3.1 Nash's theorem

Nash's theorem generalizes Von Neumann's theorem to n-person games.

Theorem 36 *(Nash) If in the game $\mathcal{G} = (N, S_i, u_i, i \in N)$ the sets S_i are convex and compact, and the functions u_i are continuous over X and quasi-concave in s_i, then the game has at least one Nash equilibrium.*

For the proof we use the following **mathematical preliminaries.**
1) Upper hemi-continuity of correspondences
A correspondence $f : A \rightarrow\rightarrow \mathbb{R}^m$ is called *upper hemicontinuous* at $x \in A$ if for any open set U such that $f(x) \subset U \subset A$ there exists an open set V such that $x \in V \subset A$ and that for any $y \in V$ we have $f(y) \subset U$. A correspondence

$f : A \to\to \mathbb{R}^m$ is called *upper hemicontinuous* if it is upper hemicontinuous at all $x \in A$.

Note that for a single-valued function f, this definition is just the continuity of f.

Proposition

A correspondence $f : A \to\to \mathbb{R}^m$ is upper hemicontinuous if and only if it has a closed graph and the images of the compact sets are bounded (i.e. for any compact $B \subset A$ the set $f(B) = \{y \in \mathbb{R}^m : y \in f(x) \text{ for some } x \in B\}$ is bounded).

Note that if $f(A)$ is bounded (compact), then the upper hemicontinuity is equivalent to the closed graph condition. Thus to check that $f : A \to\to A$ from the premises of Kakutani's fixed point theorem is upper hemicontinuous it is enough to check that it has closed graph. I.e., one needs to check that for any $x^k \in A$, $x^k \to x \in A$, and for any $y^k \to y$ such that $y^k \in f(x^k)$, we have $y \in f(x)$.

2) Two fixed point theorems

Theorem *(Brouwer's fixed point theorem)*

Let $A \subset \mathbb{R}^n$ be a nonempty convex compact, and $f : A \to A$ be single-valued and continuous. Then f has a fixed point : there exists $x \in A$ such that $x = f(x)$.

Extension to correspondences:

Theorem *(Kakutani's fixed point theorem)*

Let $A \subset \mathbb{R}^n$ be a nonempty convex compact and $f : A \to\to A$ be an upper hemicontinuous convex-valued correspondence such that $f(x) \neq \varnothing$ for any $x \in A$. Then f has a fixed point: there exists $x \in A$ such that $x \in f(x)$.

Proof of Nash Theorem.

For each player $i \in N$ define a best reply correspondence $R_i : S_{-i} \to\to S_i$ in the following way: $R_i(s_{-i}) = \arg\max\limits_{\sigma \in S_i} u_i(\sigma, s_{-i})$. Consider next the best reply correspondence $R : S \to\to S$, where $R(s) = R_1(s_{-1}) \times ... \times R_N(s_{-N})$. We will check that R satisfies the premises of the Kakutani's fixed point theorem.

First $S = S_1 \times ... \times S_N$ is a nonempty convex compact as a Cartesian product of finite number of nonempty convex compact subsets of $\mathbb{R}^p$.

Second since u_i are continuous and S_i are compact there always exist $\max\limits_{\sigma \in S_i} u_i(\sigma, s_{-i})$. Thus $R_i(s_{-i})$ is nonempty for any $s_{-i} \in S_{-i}$ and so $R(s)$ is nonempty for any $s \in S$.

Third $R(s) = R_1(s_{-1}) \times ... \times R_N(s_{-N})$ is convex since $R_i(s_{-i})$ are convex. The last statement follows from the (quasi-) concavity of $u_i(\cdot, s_{-i})$. Indeed if $s_i, t_i \in R_i(s_{-i}) = \arg\max\limits_{\sigma \in S_i} u_i(\sigma, s_{-i})$ then $u_i(\lambda s_i + (1-\lambda)t_i, s_{-i}) \geq \lambda u_i(s_i, s_{-i}) + (1-\lambda)u_i(t_i, s_{-i}) = \max\limits_{\sigma \in S_i} u_i(\sigma, s_{-i})$, and hence $\lambda s_i + (1-\lambda)t_i \in R_i(s_{-i})$.

Finally given that S is compact to guarantee upper hemicontinuity of R we only need to check that it has closed graph. Let $s^k \in S$, $s^k \to s \in S$, and $t^k \to t$ be such that $t^k \in R(s^k)$. Hence for any k and for any $i = 1, ..., N$ we have that $u_i(t^k, s^k_{-i}) \geq u_i(\sigma, s^k_{-i})$ for all $\sigma \in S_i$. Given that $(t^k, s^k_{-i}) \to (t, s_{-i})$ continuity of u_i implies that $u_i(t, s_{-i}) \geq u_i(\sigma, s_{-i})$ for all $\sigma \in S_i$. Thus $t \in \arg\max\limits_{\sigma \in S_i} u_i(\sigma, s_{-i}) = R(s)$ and so R has closed graph.

Now, Kakutani's fixed point theorem tells us that there exists $s \in S = S_1 \times ... \times S_N$ such that $s = (s_1, ..., s_N) \in R(s) = R_1(s_{-1}) \times ... \times R_N(s_{-N})$. I.e. $s_i \in R(s_{-i})$ for all players i. Hence, each strategy in s is a best reply to the vector of strategies of other players and thus s is a Nash equilibrium of our game.■

A useful variant of the theorem is for symmetrical games.

Theorem 37 *If in addition to the above assumptions, the game is symmetrical, then there exists a symmetrical Nash equilibrium $s_i = s_j$ for all i, j.*

Proof: The game is (N, S_0, u) with S_0 the common strategy set, and $u : S_0 \times S_0^{N \setminus \{1\}} \to \mathbb{R}$ its common payoff function. Check that we can apply Kakutani's theorem to the mapping R_0 from S_0 into itself:

$$R_0(s_0) == \arg\max_{\sigma \in S_0} u_i(\sigma; s_0, s_0, \cdots, s_0)$$

A fixed point of R_0 is a symmetric Nash equilibrium.

The main application of Nash's theorem is to finite games in strategic form where the players use mixed strategies.

Consider a normal form game $\Gamma_f = (N, (C_i)_{i \in N}, (u_i)_{i \in N})$, where N is a (finite) set of players, C_i is the (nonempty) **finite** set of **pure** strategies available to the player i, and $u_i : C = C_1 \times ... \times C_N \to \mathbb{R}$ is the payoff function for player i. Let $S_i = \Delta(C_i)$ be the set of all probability distributions on C_i (i.e., the set of all **mixed** strategies of player i). We extend the payoff functions u_i from C to $S = S_1 \times ... \times S_N$ by expected utility.

In the resulting game S_i will be convex compact subsets of some finite-dimensional vector space. Extended payoff functions $u_i : S \to \mathbb{R}$ will be continuous on S, and $u_i(\cdot, s_{-i})$ will be be concave (actually, linear) on S_i. Thus we can apply the theorem above to show that

Theorem 38 Γ_f *always has a Nash equilibrium in mixed strategies.*

Note that a Nash equilibrium of the initial game remains an equilibrium in its extension to mixed strategies.

3.2 Von Neumann Morgenstern utility

We axiomatize expected utility over random outcomes.
Notation:

C is the finite set of outcomes (consequences), $C = \{c_1, \cdots, c_m\}$

Δ is the set of lotteries on C with generic element $L = (p_1, \cdots, p_m), p_j \geq 0$ for all j and $\sum_1^m p_j = 1$

Definition: compound lottery

Given K (simple) lotteries $L_k \in \Delta, k = 1, \cdots, K$, and a probability distribution $\pi = (\pi_1, \cdots, \pi_K)$, the compound lottery $(L_k, k = 1, \cdots, K; \pi)$ is the random choice of an outcome in C where we pick first a lottery L_k according to π, then an outcome in C according to L_k.

The simple lottery $L = \sum_1^K \pi_k L_k$ give the same ultimate probability distribution over outcomes as the compound lottery $(L_k, k = 1, \cdots, K; \pi)$, yet it is not unreasonable to distinguish these two objects from a decision-theoretic viewpoint.

Consequentialist axiom: *the preferences of our decision maker over a compound lottery do not distinguish it from the associated simple lottery.*

In view of this axiom, the preferences of our agent over the random outcomes in C, obtained via compuond lotteries of arbitrary order, are represented by a rational preference (complete, transitive) $\preceq$ over Δ.

Continuity axiom: *upper and lower contour sets of $\preceq$ are closed in Δ.*

By the classic Debreu theorem, the continuity axiom implies that these preferences can be represented by a continuous utility function.

Independence axiom: *for all $L, L', L'' \in \Delta$,* for all $\alpha \in [0,1]$

$$L \succeq L' \Leftrightarrow \alpha L + (1-\alpha)L'' \succeq \alpha L' + (1-\alpha)L''$$

The independence axiom is very intuitive given consequentialism, and yet extremely powerful. It is the mathematical engine driving th VNM theorem.

Definition: the utility function $U : \Delta \to \mathbb{R}$ has the Von Neumann Morgenstern expected utility form if there exists real numbers $u_1, \cdots, u_m$ such that

$$U(L) = \sum_{j=1}^{m} u_j p_j \text{ for all } L = (p_1, \cdots, p_m) \in \Delta$$

An equivalent definition is that the function U is *affine* on Δ, namely

$$U(\alpha L + (1-\alpha)L') = \alpha U(L) + (1-\alpha)U(L') \text{ for all } L, L' \in \Delta, \text{ and all } \alpha \in [0,1]$$

An important invariance property of the VNM representation of a preference relation on Δ: if U has the VNM form and represents $\preceq$, so does $\beta U + \gamma$ for any numbers $\beta > 0$ and $\gamma \in \mathbb{R}$. Conversely, such utility functions are the only alternative VNM representations of $\preceq$.

A consequence of this invariance is that differences in cardinal utilities have meaning:

$$u_1 - u_2 > u_3 - u_4 \Leftrightarrow \frac{1}{2}u_1 + \frac{1}{2}u_4 > \frac{1}{2}u_2 + \frac{1}{2}u_3$$

Theorem: *(Von Neumann and Morgenstern)*
The preferences $\preceq$ over Δ meet the Continuity and Independence axioms if and only if they are representable in the expected utility form.

Critique of the independence axiom: *the Allais paradox*
Consider three outcomes

- c_1: win a prize of 800K
- c_2: win a prize of 500K
- c_3: no prize.

Now consider the two choices between two pairs of lotteries

$$L_1 = (0, 1, 0) \text{ versus } L_1' = (0.1, 0.89, 0.01)$$

$$L_2 = (0, 0.11, 0.89) \text{ versus } L_2' = (0.1, 0, 0.9)$$

A common observation is the following preferences:

$$L_1 \succ L_1', \; L_2' \succ L_2$$

but these preferences are not compatible with VNM expected utility!

3.3 mixed strategy equilibrium

Here we discuss a number of examples to illustrate both the interpretation and computation of mixed strategy equilibrium in n-person games. We start with *two-by-two games,*namely where two players have two strategies each.

Example 1 *crossing games*
We revisit the example 12 from chapter 2

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

and compute the (unique) mixed strategy equilibrium

$$s_1^* = s_2^* = \frac{1-\varepsilon}{2-\varepsilon}\text{stop} + \frac{1}{2-\varepsilon}\text{go}$$

with corresponding utility $\frac{2-2\varepsilon}{2-\varepsilon}$ for each player. So an accident (both player go) occur with probability slightly above $\frac{1}{4}$. Both players enjoy an expected utility only slightly above their secure (guaranteed) payoff of $1-\varepsilon$. Under s_1^*, on the other hand, player 1 gets utility close to $\frac{1}{2}$ about half the time: for a tiny increase in the expected payoff, our player incur a large risk.

The point is stronger in the following variant of the crossing game

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

where the (unique) mixed strategy equilibrium is

$$s_1^* = s_2^* = \frac{1+\varepsilon}{2+\varepsilon}\text{stop} + \frac{1}{2+\varepsilon}\text{go}$$

and gives to each player exactly her guaranteed utility level in the mixed game. Indeed a (mixed) prudent strategy of player 1 is

$$\widetilde{s}_1 = \frac{2}{2+\varepsilon}\text{stop} + \frac{\varepsilon}{2+\varepsilon}\text{go}$$

and it guarantees the expected utility $\frac{2+2\varepsilon}{2+\varepsilon}$, which is also the mixed equilibrium payoff. Now the case for playing the equilibrium strategy in lieu of the prudent one is even weaker.

In general finite games computing the mixed equilibrium or equilibria follows the same general approach as for two-person zero-sum games. The difficulty is to identify the support of the equilibrium strategies, typically of equal sizes[2]. Once this is done we need to solve a linear system and check a few inequalities.

Unlike in two-person zero-sum games, we may have several mixed equilibria with very different payoffs. A general theorem shows that for "most games", the number of mixed or pure equilibria is *odd*.

Example 2 *public good provision (Bliss and Nalebuff)*
Each one of the n players can provide the public good (hosting a party, slaying the dragon, or any other example where only one player can do the job) at a cost $c > 0$. The benefit is b to every agent if the good is provided. We assume $c < nb$: the social benefit justifies providing the good. The palyers can divide the burden of providing the good by the following use of lotteries. Each player chooses to step forward (volunteer) or not. If nobody volunteers, the good is not provided; if some players volunteer, we choose one of them with uniform probability to provide the good.

If $b < c$, the game in pure strategies is a classic Prisoner's Dilemna (section 2.2.3) . If $b > c$, it resembles the war of attrition (sectione 2.2.1) in that we have n pure strategy equilibria where one player provides the good and the other free ride.

If there is a symmetrical equilibrium in mixed strategies in which every player steps forward with probability p^*, then p^* solves

$$\frac{nb}{c}p = \frac{1-(1-p)^n}{(1-p)^{n-1}} = f(p)$$

Notice that f is convex, increasing, from $f(0) = 0$ to $f(1) = \infty$, and $f'(0) = n$. Therefore if $b < c$, the only solution of the equation above is $p = 0$ and we are back to the Prisoner's Dilemna. But if $b > c$, there is a unique equilibrium in mixed strategies. For instance if $n = 2$, we get

$$p_2^* = \frac{2(b-c)}{2b-c} \text{ and } u_i(p^*) = \frac{2b(b-c)}{2b-c}$$

One checks that as n grows, p_n^* goes to zero as $\frac{K}{n}$ where K is the solution of

$$\frac{c}{b} = \frac{Ke^K}{1-e^{-K}}$$

therefore the probability that the good be provided goes to $1 - e^{-K}$, but the probability of volunteering of each player goes to zero.

[2] If a finite game has a mixed equilibrium with supports of different sizes, then an arbitrarily small change in the payoffs will eliminate such "abnormal" equilibria.

Example 3 *war of attrition (a.k.a. all-pay second price auction)*
The n players compete for a prize worth \$$p$ by "hanging on" longer than everyone else. Hanging on costs \$1 per unit of time. Once a player is left alone, he wins the prize without spending any more effort. The game in pure strategies is a game of timing as in Example 14 chapter 2:

$$u_i(s) = p - \max_{j \neq i} s_j \text{ if } s_i > \max_{j \neq i} s_j; = -s_i \text{ if } s_i < \max_{j \neq i} s_j; = \frac{p}{K} - s_i \text{ if } s_i = \max_{j \neq i} s_j$$

where K is the number of largest bids.

In addition to the pure equilibria described in the previous chapter, we have one symmetrical equilibrium in completely mixed strategies where each player independently chooses s_i in $[0, \infty[$ according to the cumulative distribution function

$$F(x) = (1 - e^{-\frac{x}{p}})^{\frac{1}{n-1}}$$

In particular the support of this distribution is $[0, \infty[$ and for any $B > 0$ there is a positive probability that a player bids above B. The payoff to each player is zero so the mixed strategy is not better than the prudent one (zero bid) payoffwise. It is also more risky.

Example 4 *lobbying game (a.k.a. all-pay first price auction)*
The n players compete for a prize of \$$p$ and can spend \$$s_i$ on lobbying (bribing) the relevant jury members. The largest bribe wins the prize; all the money spent on bribes is lost to the players. Hence the payoff functions

$$u_i(s) = p - s_i \text{ if } s_i > \max_{j \neq i} s_j; = -s_i \text{ if } s_i < \max_{j \neq i} s_j; = \frac{p}{K} - s_i \text{ if } s_i = \max_{j \neq i} s_j$$

The game has no equilibrium in pure strategies. In the symmetrical mixed Nash equilibrium each player independently chooses a bid in $[0, p]$ according to the cumulative distribution function

$$F(x) = (\frac{x}{p})^{\frac{1}{n-1}}$$

As in the above example the equilibrium payoff is zero, just like the guaranteed payoff from a null bid.

3.4 correlated equilibrium

Given a finite n-players game in strategic form $\Gamma = (N, (C_i)_{i \in N}, (u_i)_{i \in N})$, a correlation device is a *lottery* L over the set $C = C_1 \times ... \times C_n$ of strategy profiles. The interpretation is that the lottery itself is a *non binding* agreement to play according to its outcome. Thus the lottery is built jointly by the players (much like we say that the players jointly reach an agreement to play a certain Nash equilibrium), and once it draws an outcome $x \in C$, the players are *supposed to* play accordingly, namely player i chooses x_i in C_i.

If the outcome of the lottery is publicly known, the agreement will be self enforcing if and only if the support of the lottery consists of Nash equilibrium outcomes (in pure strategies). Then the lottery is a simple coordination device over a set of equilibria in pure strategies. This is a useful coordination device, for instance to achieve a fair compromise between asymetric equil;ibria in a symmetric game. In the crossing game of example 1, tossing a fair coin between the two equilibria yields a payoff of $1.5 \pm \varepsilon$, much better than the payoff of the only symmetric equilibrium, in mixed strategies. We can interpret a red light as achieving precisely this kind of coordination when two lines of traffic cross.

More interesting is the scenario where the distribution L is known to everyone, but the outcome of the lottery is only *partially revealed* to each player. Specifically player i learns the i-th coordinate of the outcome x, but no more: then she evaluates the random strategies chosen by other players according to the conditional probability of L given x_i. If other players are indeed following the recommendation of the correlation device, this evaluation is correct. Now the equilibrium (self-enforcing) property of the lottery L states that player i's best reply to any recommendation x_i is to comply.

Given a lottery $L \in \Delta(C)$ we write its support $[L] \subset C$ and the projection of the support on C_i as $proj_i\{[L]\}$. This set contains the strategies of player i that the device recommends to play with positive probability. For any i and $x_i \in C_i$, we denote by $L(x_i)$ the corresponding conditional probability of L on $C_{N\diagdown\{i\}}$. Thus if L_x denotes the probability that L selects outcome x, we have

$$L(x_i)_{x_{-i}} = \frac{L_{(x_i,x_{-i})}}{\sum_{y_{-i}\in C_{N\diagdown\{i\}}} L_{(x_i,y_{-i})}} \text{ for all } x_i \in proj_i\{[L]\} \text{ all } x_{-i} \in C_{N\diagdown\{i\}}$$

Definition *A lottery $L \in \Delta(C)$ is a correlated equilibrium of the game $(N, (C_i)_{i\in N}, (u_i)_{i\in N})$ if for all $i \in N$ we have*

$$u_i(x_i, L(x_i)) \geq u_i(y_i, L(x_i)) \text{ for all } y_i \in C_i \text{ and all } x_i \in proj_i\{[L]\}$$

$$\Leftrightarrow \sum_{y_{-i}\in C_{N\diagdown\{i\}}} u_i(x_i, y_{-i}) L_{(x_i,y_{-i})} \geq \sum_{y_{-i}\in C_{N\diagdown\{i\}}} u_i(y_i, y_{-i}) L_{(x_i,y_{-i})} \text{ for all } y_i, x_i \in C_i$$

If $s \in \Delta(C_1) \times ... \times \Delta(C_n)$ is an equilibrium in mixed strategies, then the lottery $L = s_1 \oplus s_2 \oplus \cdots \oplus s_n$ is a correlated equilibrium. This remark establishes that a correlated equilibrium always exists in a finite game.

The most important feature of the set $\mathcal{C}$ of correlated equilibria is that it is a convex, compact subset of $\Delta(C)$. Indeed $\mathcal{C}$ is defined by a finite set of linear inequalities in $\Delta(C)$. Thus it contains all convex combinations of Nash equilibria, pure and mixed.

In some games, that is all. For instance suppose each player has a strictly dominant strategy: then the unique Nash equilibrium is also the unique correlated equilibrium. Furthermore the support of any correlated equilibrium must resist the successive elimination of strictly dominated strategies, and there is always one correlated equilibrium of which the support resists the successive elimination of weakly dominated strategies.

On the other hand in some cases correlation allows a considerable improvement upon the Nash equilibrium outcomes.

Example 5 *another Battle of the Sexes*

home	$10,10$	$5,13$
theater	$13,5$	$0,0$
	home	theater

One of the spouses must stay home, lest they are both very unhappy to call for a baby sitter. Both would prefer to go to the theater if the other stays home. Each must commit to one of the two strategies *before* returning home, and without the possibility to communicate with each other.
There are two equilibria in pure strategies, and a mixed equilibrium where each player goes out with probability $\frac{3}{8}$. The expected payoff of the latter is 8.1 for each. Tossing a fair coin *before leaving to work* between the two equilibria yields the payoff 9 for each spouse.
There is a better correlated equilibrium, choosing (theater, home) and (home, theater) each with probability $\frac{3}{11}$, and (home,home) with probability $\frac{5}{11}$. The expected payoff is now 9.45 for each.

Example 6 *musical chairs*
We have n players and 2 "chairs" (locations), with $n \geq 5$. The game is symetrical. Each player chooses a chair. His payoff is $+4$ if he is alone to make this choice, 1 if one other player (exactly) makes the same choice, and 0 otherwise (i.e., if his choice is shared by at least 2 other players).

In a pure strategy equlibria of the game, each chair is filled by two or more players and all such outcomes are equilibria. The total payoff is 2 or 0. In the symmetric mixed equilibrium each player chooses a chair with probability 0.5, and the resulting expected payoff is

$$4\frac{1}{2^n}\binom{n}{1} + 2\frac{1}{2^n}\binom{n}{2} = \frac{n(n+3)}{2^n} \ll 2$$

(there are no other mixed equilibria)

The best symmetric correlated equilibrium (i.e., the one giving the highest total payoff) selects with probability $\pi = \frac{2}{n-3}$ a distribution where one player sits alone (and chooses with uniform probability among all such distributions), and with probability $1-\pi = \frac{n-5}{n-3}$ it picks a distribution where two players share one chair (and chooses with uniform probability among all such distributions). The total payoff is $2 + \frac{12}{n-3}$.

3.5 games of incomplete information

A game in Bayesian form(or Bayesian game) specifies

- the set N of players
- the set of pure strategies X_i for each player i
- the set of *types* T_i of each player i
- the set of *beliefs* of each player i, represented by a probability distribution $\pi_i(\cdot|t_i)$ over $T_{N \diagdown \{i\}}$: one distribution for each possible type of player i

- the payoff function $u_i(x, t)$ for each player i, where $x \in X_N$ and $t \in T_N$.

A Bayesian equilibrium is decribed by a mixed strategy for each player, conditional on his type: $s_i(t_i) \in \Delta(X_i)$. The equilibrium property is

$$\forall i, t_i \in T_i, \forall s'_i \in \Delta(X_i):$$

$$\sum_{t_{-i} \in T_{N \setminus \{i\}}} \pi_i(t_{-i}|t_i) u_i(s(t), t) \geq \sum_{t_{-i} \in T_{N \setminus \{i\}}} \pi_i(t_{-i}|t_i) u_i(s'_i; s_{-i}(t_{-i}), t)$$

where we use the notation

$$s(t) \in \Pi_{i \in N} \Delta(X_i), s_{-i}(t_{-i}) \in \Pi_{j \in N \setminus \{i\}} \Delta(X_j) : s_j(t) = s_j(t_j)$$

It is enough in the equilibrium property to consider deviations to pure strategies $x_i \in X_i$. Therefore the number of inequalities characterizing the equilibrium is $\sum_i |T_i||X_i|$.

Theorem: *If the sets X_i and T_i are finite, the game possesses at least one Bayesian equilibrium.*

This is a direct consequence of Nash's theorem, after observing that a Bayesian equilibrium is a Nash equilibrium (in pure strategies) of the game with $\mathcal{N} = \oplus_i T_i$, strategy set $\Delta(X_i)$ for each player $(i, t_i) \in \mathcal{N}$ and payoffs

$$\widetilde{u}_{(i,t_i)}(s) = \sum_{t_{-i} \in T_{N \setminus \{i\}}} \pi_i(t_{-i}|t_i) u_i(s_{(i,t_i)}; s_{(j,t_j)} \ j \in N \setminus \{i\})$$

This game meets all the assumptions of Nash's Theorem (in particular utility is linear in own strategy).

The common prior, common knowledge assumption

In most examples , the individual beliefs are consistent, they are derived from a common prior, namely a probability distribution π over T_N, such that $\pi_i(\cdot|t_i) = \pi(\cdot|t_i)$ is simply the conditional probability induced by π once a player learns his type. This distribution π is *common knowledge*, which means that player i knows it, i knows that player j knows it, j knows that player i knows that player j knows it, and so on. More generally, for any sequence $i, j, k, \cdots, l$ of players (possibly with repetition): i knows that j knows that k knows that $\cdots$ that l knows it.

The classic story of the 40 villagers illustrates the subtle role of the common knowledge assumption.

In a Bayesian game where the beliefs are not consistent, the interpretation of the equilibrium notion is more difficult.

Example 7:

Two players, player 1's type is known, that of player 2 is t_1 with probability 0.6, t_2 with probability 0.4:

T	$1,2$	$0,1$
B	$0,4$	$1,3$
t_1	L	R

T	$1,3$	$0,4$
B	$0,1$	$1,2$
t_2	L	R

Note that player 2 has a dominant strategy, hence the unique equilibrium is

$$x_1 = T; x_2 = L \text{ if } t_1, \ = R \text{ if } t_2$$

Note that this is not the same as playing the unique B equlibrium in each matrix separately.

Another example with the same information structure:

T	$0,2$	$2,0$		T	$1,1$	$5,0$
B	$2,0$	$0,2$		B	$0,5$	$3,3$
t_1	L	R		t_2	L	R

Here the game under t_1 is essentially matching pennies, and under t_2 player 2 has a dominant strategy to play L. There is no pure strategy equilibrium, as the sequences of best replies are: $LL \to B \to RL \to T \to LL$, and $RR \to T$, $LR \to B$. In the unique Bayesian equilibrium player 1 mixed strategy is the optimal play for matching pennies, because under t_2 player 2 plays L for sure:

$$s_1 = \frac{1}{2}T + \frac{1}{2}B; s_2 = \frac{2}{3}L + \frac{1}{3}R \text{ if } t_1, \ = L \text{ if } t_2$$

Another example with the same information structure:

T	$0,2$	$2,0$		T	$2,0$	$1,2$
B	$2,0$	$0,2$		B	$0,3$	$2,0$
t_1	L	R		t_2	L	R

Here again we have no pure strategy equilibrium, as the best reply sequence is $T \to LR \to B \to RL \to T$. In the unique B equilibrium, player 1's mixed strategy neutralizes player 2 in one but not both of the two 2x2 matrix games. One computes:

$$s_1 = \frac{1}{2}T + \frac{1}{2}B; s_2 = \frac{5}{6}L + \frac{1}{6}R \text{ if } t_1, \ = L \text{ if } t_2$$

Example 8 *a two-person zero sum betting game*
Bob (column player) draws a card High or Low with equal probability $\frac{1}{2}$. Ann (row player) has a Medium card (a fact known to Bob). Bob can raise (R) or stay put (P). After seeing Bob's mover, Ann can see (S) or fold (F). Payoffs are as follows

S	$-10,10$	$-4,4$		S	$10,-10$	$4,-4$
F	$-1,1$	$1,-1$		F	$-1,1$	$1,-1$
$High$	R	P		Low	R	P

Here Ann has 4 pure strategies denoted XY for do X if Bob raises, do Y if he stays; Bob's strategy depends on his type, and is written similarly XY for do X if High, do Y if Low.

Check first there is no pure strategy equilibrium, as the sequence of best replies is

$$RR \to SS \to RP \to FS \to RR; PR \to SF \to RP \to \cdots; PP \to FF \to RP \to \cdots$$

Bob has a dominant strategy to raise if his card is high; thus his P strategy reveals to Ann that he is Low, in which case she wants to see. Therefore the

Bayesian equilibrium takes the form

Ann:	$p\delta_S + p'\delta_F$ if Bob raises; S if Bob stays put
Bob:	R if High; $q\delta_R + q'\delta_P$ if Low

The equilibrium conditions are

$$\text{for Ann:} \quad \frac{1}{1+q}(-10) + \frac{q}{1+q}(10) = -1 \Rightarrow q = \frac{9}{11}$$

$$\text{for Bob:} \quad p(-10) + p'(1) = -4 \Rightarrow p = \frac{5}{11}$$

In equilibrium Ann expects to pay $\$\frac{6}{11}$ to Bob: private information is more valuable than second move.

Example 9: *first price auction (Vickrey)*
Each player draws a valuation in the $[0, 100]$ interval. The draws are IID with cumulative distribution function F. We asume that F is continuous: the underlying distribution has no atoms.

The *symmetrical* equilibrium has player i bid $x(t_i)$ where t_i is his (privately known) valuation.The expected payoff to player i from bidding y, given that other players use the equilibrium strategy $x(\cdot)$ is

$$u_i(y|t_i) = (t_i - y)\pi\{x(t_j) < y \text{ for all } j \neq i\}$$

Assuming $x(\cdot)$ is increasing and continuous (this can be shown to hold true if $x(\cdot)$ is an equilibrium), player i must choose his bid $y = x(t)$ so as to maximize $(t_i - x(t))F^{n-1}(t)$. The equilibrium property is that $t = t_i$ is such a maximizer. Differentiating:

$$x'(t)F^{n-1}(t) - (t - x(t))\{F^{n-1}(t)\}' = 0$$

The boundary condition is $x(0) = 0$. A zero valuation player does not want to bid any positive amount. The differential equation writes

$$\{x(t)F^{n-1}(t)\}' = t\{F^{n-1}(t)\}'; \quad x(0) = 0$$

Therefore

$$x(t) = \frac{\int_0^t z dF^{n-1}(z)}{F^{n-1}(t)} = E[t_{(2)}|t_{(1)} = t]$$

where $t_{(k)}$ is the k-th order statistics of the n variables t_i. To check the second equality, observe that for all $a, t, a < t$

$$\pi\{t_{(2)} = a|t_{(1)} = t\} = \pi\{t_{-1} \leq a|t_{-1} \leq t; t_1 = t\} = \pi\{t_{-1} \leq a|t_{-1} \leq t\} = \frac{F^{n-1}(a)}{F^{n-1}(t)}$$

The equilibrium bid is the expected value of the second highest bid, conditional on your own bid winning the object.

For instance if F is the uniform distribution on $[0,100]$, $x(t) = \frac{n-1}{n}t$ and the expected highest bid (revenue of the seller) is $E[x(t_{(1)})] = \frac{n-1}{n}E[t_{(1)}] = \frac{n-1}{n+1}100$ while the expected joint surplus to the seller and bidders is $E[t_{(1)}] = \frac{n}{n+1}100$, because the efficient buyer (the one with the highest valuation) gets the object. This leaves only an expected gain of $\frac{1}{n(n+1)}100$ per bidder!

Interestingly this sharing of the surplus between buyers and the seller is the same as in Vickrey's second price auction, because there the revenue of the seller is

$$E[t_{(2)}] = \int_0^{100} E[t_{(2)}|t_{(1)} = t]dF^n(t) = \int_0^{100} x(t)dF^n(t) = E[x(t_{(1)})]$$

Example 10 *bilateral trade (Myerson and Satterthwaite)*
The object is worth a to the seller, b to the buyer. Both a and b are IID on $[0,300]$ with uniform distribution. They play the *sealed bid double auction* game: they independently and simulatneously send an ask price x (seller) and an offer price y (buyer). If $x > y$, no trade takes place; if $x \leq y$, trade takes place at price $p = \frac{x+y}{2}$.

One checks first that $x(a) = a, y(b) = b$ is not an equilibrium. Then we compute the *linear equilibrium,* i.e., each player uses a bid function that is linear in own valuation

$$x(a) = \alpha a + \beta; y(b) = \gamma b + \delta$$

There is a unique linear equilibrium

$$x(a) = \frac{2}{3}a + 75; y(b) = \frac{2}{3}b + 25$$

Actually there are many other non linear equilibria. See problem 15 for an example. Computing them all is an open problem.

3.6 Problems for Chapter 3

Problem 1
a) In the two-by-two game

T	$5,5$	$4,10$
B	$10,4$	$0,0$
	L	R

Compute all Nash equilibria. Show that a slight *increase* in the (B,L) payoff to the row player results in a *decrease* of his mixed equilibrium payoff.
b) Consider the crossing game of example 1

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

and its variant where strategy "go" is more costly by the amount $\alpha, \alpha > 0$, to the row player:

$$\begin{array}{ccc} stop & 1,1 & 1-\varepsilon, 2 \\ go & 2-\alpha, 1-\varepsilon & -\alpha, 0 \\ & stop & go \end{array}$$

Show that for α and ε small enough, row's mixed equilibrium payoff is *higher* if the go strategy is more costly.

Problem 2

Three plants dispose of their water in the lake. Each plant can send clean water ($s_i = 1$) or polluted water ($s_i = 0$). The cost of sending clean water is c. If only one firm pollutes the lake, there is no damage to anyone; if two or three firms pollute, the damage is a to everyone, $a > c$.

Compute all Nash equilibria in pure and mixed strategies.

Problem 3

Give an example of a two-by-two game where no player has two equivalent pure strategies, and the set of Nash equilibria is infinite.

Problem 4

A two person game with finite strategy sets $S_1 = S_2 = \{1, \cdots, p\}$ is represented by two $p \times p$ payoff matrices U_1 and U_2, where the row player is labeled 1 and the column player is 2. The entry $U_i(j,k)$ is player i's payoff when row chooses j and column chooses k. Assume that both matrices are invertible and denote by $|A|$ the determinant of the matrix A. Then write $\widetilde{U}_i(j,k) = (-1)^{j+k}|U_i(j,k)|$ the (j,k) cofactor of the matrix U_i, where $U_i(j,k)$ is the $(p-1) \times (p-1)$ matrix obtained from U_i by deleting the j row and the k column.

Show that if the game has a completely mixed Nash equilibrium, it gives to player i the payoff

$$\frac{|U_i|}{\sum_{1 \le j,k \le p} \widetilde{U}_i(j,k)}$$

Problem 5

In this symmetric two-by-two-by-two (three-person) game, the mixed strategy of player i takes the form $(p_i, 1-p_i)$ over the two pure strategies. The resulting payoff to player 1 is

$$u_1(p_1,p_2,p_3) = p_1p_2p_3 - 3p_1(p_2+p_3) + p_2p_3 - p_1 - 2(p_2+p_3)$$

Find the symmetric mixed equilibrium of the game. Are there any non symmetric equilibria (in pure or mixed strategies)?

Problem 6

Let$(\{1,2\}, C_1, C_2, u_1, u_2)$ be a finite two person game and $\mathcal{G} = (\{1,2\}, S_1, S_2, u_1, u_2)$ be its mixed extension. Say that the set $\mathcal{NE}(\mathcal{G})$ of mixed Nash equilibrium outcomes of $\mathcal{G}$ has the *rectangularity property* if we have for all $s, s' \in S_1 \times S_2$

$$s, s' \in \mathcal{NE}(\mathcal{G}) \Rightarrow (s'_1, s_2), (s_1, s'_2) \in \mathcal{NE}(\mathcal{G})$$

a) Prove that $\mathcal{NE}(\mathcal{G})$ has the rectangularity property if and only if it is a convex subset of $S_1 \times S_2$.
b) In this case, prove there exists a Pareto dominant mixed Nash equilibrium s^*:

$$\text{for all } s \in \mathcal{NE}(\mathcal{G}) \Rightarrow u(s) \leq u(s^*)$$

Problem 7 *all-pay second price auction*
This is a variant of example 3 with only two players who value the prize respectively at a_1 and a_2. The payoff are

$$u_i(s_1, s_2) = a_i - s_j \text{ if } s_j < s_i; = -s_i \text{ if } s_i < s_j; = \frac{1}{2}a_i - s_i \text{ if } s_j = s_i;$$

For any two numbers b_1, b_2 in $[0, 1]$ such that $\max\{b_1, b_2\} = 1$, consider the mixed strategy of player i with cumulative distribution function

$$F_i(x) = 1 - b_i e^{-\frac{x}{a_j}}, \text{ for } x \geq 0$$

Show that the corresponding pair of mixed strategies (s_1, s_2) is an equilibrium of the game.

Riley shows that these are the only mixed equilibria of the game.

Problem 8 *all-pay first price auction*
This is a variant of example 4 with only two players who value the prize respectively at a_1 and a_2. The payoffs are

$$u_i(s_1, s_2) = a_i - s_i \text{ if } s_j < s_i; = -s_i \text{ if } s_i < s_j; = \frac{1}{2}a_i - s_i \text{ if } s_j = s_i$$

Assume $a_1 \geq a_2$. Show that the following is an equilibrium:

player 1 chooses in $[0, a_2]$ with uniform probability;

player 2 bids zero with probability $1 - \frac{a_2}{a_1}$, and with probability $\frac{a_2}{a_1}$ he chooses in $[0, a_2]$ with uniform probability.
Riley shows this is the unique equilibrium if $a_1 > a_2$.

Problem 9 *first price auction*
We have two players who value the prize respectively at a_1 and a_2. The payoffs are

$$u_i(s_1, s_2) = a_i - s_i \text{ if } s_j < s_i; = 0 \text{ if } s_i < s_j; = \frac{1}{2}(a_i - s_i) \text{ if } s_j = s_i$$

a) Assume $a_1 = a_2$. Show that the only Nash equilibrium of the game in mixed strategies is $s_1 = s_2 = a_i$.
b) Assume $a_1 > a_2$. Show there is no equilibrium in pure strategies. Show that in any equilibrium in mixed strategies

player 1 bids a_2

player 2 chooses in $[0, a_2]$ according to some probability distribution π such that for any interval $[a_2 - \varepsilon, a_2]$ we have $\pi([a_2 - \varepsilon, a_2]) \geq \frac{\varepsilon}{a_2 - a_1}$.
Give an example of such an equilibrium.

Problem 10 *a location game*
Two shopowners choose the location of their shop in $[0,1]$. The demand is inelastic; player 1 captures the whole demand if he locates where player 2 is, and player 2's share increases linearly up to a cap of $\frac{2}{3}$ when he moves away from player 1. The sets of pure strategies are $C_i = [0,1]$ and the payoff functions are:

$$\begin{aligned} u_1(x_1, x_2) &= 1 - |x_1 - x_2| \\ u_2(x_1, x_2) &= \min\{|x_1 - x_2|, \frac{2}{3}\} \end{aligned}$$

a) Show that there is no Nash equilibrium in pure strategies.
b) Show that the following pair of mixed strategies is an equilibrium of the mixed game:

$$\begin{aligned} s_1 &= \frac{1}{3}\delta_0 + \frac{1}{6}\delta_{\frac{1}{3}} + \frac{1}{6}\delta_{\frac{2}{3}} + \frac{1}{3}\delta_1 \\ s_2 &= \frac{1}{2}\delta_0 + \frac{1}{2}\delta_1 \end{aligned}$$

and check that by using such a strategy, a player makes the other one indifferent between all his possible moves.

Problem 11 *Correlated equilibrium*
In the crossing game of example 1, compute all correlated eqilibria. Show that the best symmetric one is a simple "red light".

Problem 12 *more musical chairs*
Consider three variants of example 6 where

- there are two chairs and 3 players
- there are two chairs and 4 players
- there are three chairs and n players, $n \geq 7$

In each case discuss the equilibria in pure strategies, in mixed strategies, and the best symmetric correlated equilibrium.

Problem 13 *Correlated equilibrium*
We have three players named $1, 2, 3$, each with two strategies labeled A, B. The game is symmetrical, and the payoffs are as follows:

$$\begin{aligned} (B, B, A) &\rightarrow (2, 2, 0) \\ (A, A, A) \text{ or } (B, B, B) &\rightarrow (1, 1, 1) \\ (B, A, A) &\rightarrow (0, 0, 0) \end{aligned}$$

a) Find all equilibria in pure strategies, and all equilibria in mixed strategies.
b) Find the symmetrical correlated equilibrium with the largest common payoff.

Problem 14 *a coordination game*

There are q locations equally distributed on the oriented unit circle, $q \geq 3$, and each of the two players chooses one location. The payoff to both players is 1 if they choose the same location, 0 if they choose two different locations that are not adjacent. If the two choices are adjacent, the player who precedes the other (given the orientation of the circle) gets a payoff of 3, the other one gets a payoff of 2.
Show that the game has no pure strategy equilibrium; compute its symmetric equilibrium in mixed strategies and the corresponding payoffs.
Show there is no other equilibrium in mixed strategies.
Construct a correlated equilibrium where total payoff is maximal, anmely 2.5 for each player.

Problem 15

Find all equilibria in pure and mixed strategies of the following three person game. Each player has two pure strategies, $C_i = \{x_i, y_i\}$ for all $i = 1, 2, 3$. The payoff is zero to everybody, unless exactly one player i chooses y_i, in which case this player i gets 5, the player before i in the $1 \to 2 \to 3 \to 1$ cycle gets 6, and the player after i in this cycle gets 4. Note that the game is not symmetric in the sense of Definition 21 (Chapter 2), yet it is *cyclically* symmetric, i.e., with respect to the cycle $1 \to 2 \to 3 \to 1$.
Compute the (fully) *symmetric* correlated equilibria of the game and compare their payoffs to those of the pure and mixed equilibria.

Problem 16 *Bayesian equilibrium*

a) The strategy sets and information structure is as in Example 7, and the payoffs are

T	$1,2$	$0,0$		T	$0,0$	$3,1$
B	$0,0$	$2,1$		B	$1,3$	$0,0$
t_1	L	R		t_2	L	R

Check that we have two pure strategy equilibria. How many Bayesian equilibria involving mixed strategies?
b) The payoffs are now

T	$1,2$	$0,0$		T	$4,1$	$0,0$
B	$0,0$	$2,1$		B	$0,0$	$2,3$
t_1	L	R		t_2	L	R

Find all Bayesian equilibria.
c) Player 1 chooses a row and his type is known, player 2 chooses a column and his type is t_1 with probability $\frac{2}{3}$, t_2 with probability $\frac{1}{3}$. Payoffs are:

T	$2,0$	$0,2$		T	$0,0$	$2,2$
B	$0,2$	$2,0$		B	$3,3$	$0,0$
t_1	L	R		t_2	L	R

Find all equilibria in pure strategies and all Bayesian equilibria.

Problem 17

Two opposed armies are poised to seize an island. Each army's general chooses (simultaneously and independently) either to attack or not to attack. In addition, every army is either strong or weak, with equal probability, and the army's type is known to its general (but not to the general of the opposed army). An

army captures the island if either it attacks it while its opponent does not attack, or if it attacks while strong, whereas its rival is weak. If two armies of equal strength both attack, neither captures the island.
Payoffs are zero initially; the island is worth 8 if captured; an army incurs a cost of fighting, which is 3 if it is strong and 6 if it is weak. There is no cost of attacking if the rival does not attack, and no cost to not attacking.
Give the normal form of the game, eliminate dominated strategies if any, and compute all Bayesian equilibria.

Problem 18 *all-pay first price auction*
The game is identical to that in example 4, except for the fact that the valuation t_i of the object to player i is known only to this agent. Other agents know that t_i is drawn from the uniform probability distribution over $[0, 100]$, and that all draws are stochastically independent.
a) Show that if bidder i observes his type t_i, contemplates the bid y and knows that other bidders all use the same bidding function $x(t)$, bidder i's expected pay-off is

$$t_i \pi\{x(t_j) < y \text{ for all } j \neq i\} - y$$

b) Deduce the unique symmetrical equilibrium bidding function $x(\cdot)$. Compare it to the symmetrical equilibrium of the first price auction.
c) Show that the expected revenue to the seller is the same as in the first price auction (example 8) and in the second price auction. Compare the expected profit of a bidder in these three auctions.

Problem 19 *sealed bid double auction*
In the game of example 10, consider the folowing pair of strategies, where α is a number in [0,300]:

$$\text{seller } x(a) = \alpha \text{ if } 0 \leq a \leq \alpha; = 300 \text{ if } \alpha < a \leq 300$$

$$\text{buyer } y(b) = 0 \text{ if } 0 \leq b < \alpha; = \alpha \text{ if } \alpha \leq b \leq 300$$

Show that this pair is a Bayesian equilibrium.
Compute its welfare loss and choose α so that it is minimal. Then compare it to the welfare loss of the linear equilibrium found in example 10.

Problem 20 *the lemon problem*
The seller's reservation price t is drawn in $[0, 100]$ with uniform probability. The buyer does not see t. Her reservation price for the object is $\frac{3}{2}x$.
a) Suppose the buyer makes a "take it or leave it" offer which the seller can only accept or reject. Show that the only Bayesian equilibrium of this game has the buyer offering a price of zero, which the seller always refuses.
b) What is the Bayesian equilibrium of the game where the seller makes a "take it or leave it" offer which the buyer can only accept or reject?