Lecture notes on Game Theory: Chapters 3,4
Econ 440

Hervé Moulin

Spring 2009

1 Chapter 3: mixed strategies, correlated and Bayesian equilibrium

1.1 Nash's theorem

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

Theorem 1 *(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 $f : A \rightarrow\rightarrow \mathbb{R}^m$ 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 2 *A correspondence $f : A \rightarrow\rightarrow \mathbb{R}^m$ 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 $f(B) = \{y \in \mathbb{R}^m : y \in f(x)$ 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 $f : A \rightarrow\rightarrow A$ 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 3 *(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 4 *(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: $R_i(s_{-i}) = \arg\max_{\sigma \in S_i} u_i(\sigma, s_{-i})$. 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 $\max_{\sigma \in S_i} u_i(\sigma, s_{-i})$. 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 $s_i, t_i \in R_i(s_{-i}) = \arg\max_{\sigma \in S_i} u_i(\sigma, s_{-i})$ then $u_i(\lambda s_i + (1-\lambda)t_i, s_{-i}) \geq \lambda u_i(s_i, s_{-i}) + (1-\lambda)u_i(t_i, s_{-i}) = \max_{\sigma \in S_i} u_i(\sigma, s_{-i})$, 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 $t \in \arg\max_{\sigma \in S_i} u_i(\sigma, s_{-i}) = R(s)$ 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 5 *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 $u_i : C = C_1 \times ... \times C_N \rightarrow \mathbb{R}$ 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. The normative assumptions justifying this type of preferences over uncertain outcomes are the subject of the next section.

In the resulting game S_i will be convex compact subsets of some finite-dimensional vector space. Extended payoff functions $u_i : S \rightarrow \mathbb{R}$ 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 6 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.

The Problems offer several applications of Nash's theorem, in particular problems ?/

1.2 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. In those games existence of a Nash equilibrium is guaranteed by the general fixed point theorem of Tarski, stating that an increasing function in a lattice must have at least a fixed point. A simple instance of this result is that any non decreasing function f from $[0,1]^n$ into itself (i.e., $x \leq x' \Rightarrow f(x) \leq f(x')$) has a fixed point. We also know that it has a smallest fixed point, and a largest fixed point.

By way of illustration of Tarski's theorem, consider a symmetric game where $S_i = [0,1]$ and the (symmetric) best reply function $s \rightarrow br(s, \cdots, s)$ is non decreasing. This function must cross the diagonal, which shows that a symmetric Nash equilibrium exists.

Proposition 7 *Let the strategy sets S_i be either finite, or real intervals $[a_i, b_i]$. 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 8 *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 $[s_-, s^+]$

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

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

Note that if u_i is twice differentiable the SC property holds if and only if

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

Example 1 *Voluntary contribution to a public good (continued)*
Consider Example 20 of chapter 2 where $z \to B(z)$ is convex over $\mathbb{R}_+$. Then the game has the SC property, 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. An example is $B(x) = \frac{1}{2}x^2, C_i(x) = \frac{1}{3}\frac{x^3}{a_i^2}$.

Example 2 *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 to s_-.
The game is a potential game as well, provided we rescale the utility functions as

$$v_i(s) = \frac{1}{\alpha_i} u_i(s) = s_i s_{N \diagdown \{i\}} - \frac{1}{\alpha_i} C_i(s_i)$$

so the potential is

$$P(s) = \sum_{i \neq j} s_i s_j - \sum_i \frac{1}{\alpha_i} C_i(s_i)$$

Example 3 *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

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

Check that for any p_{-i}, the best reply of player i is

$$br_i(s_{-i}) = \frac{1}{\sqrt{\alpha_i}} \sqrt{A_i + \sum_{j \neq i} \beta_j p_j}$$

so that the game has increasing best reply functions. On the other hand it does not have the single crossing property.
In the symmetric case ($A_i = A, \alpha_i = \alpha, \beta_i = \beta$), one checks that its equilibrium is unique and is strongly stable.

1.3 Von Neumann Morgenstern utility

We axiomatize preferences over random outcomes represented by an expected utility function.
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 9 *(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 compound 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 10 *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 11 *(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.*

A consequence of the Independence axiom is the property that indifference contours of these preferences are straight lines; this is the key argument in the proof of the Theorem.

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 commonly observed set of preferences are:

$$L_1 \succ L'_1,\ L'_2 \succ L_2$$

but these preferences are not compatible with VNM expected utility!

1.4 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 (* two players have two strategies each).

Example 4 *crossing games*
We revisit the example 12 from chapter 2

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

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

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

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, unless we maintain a strict interpretation of the VNM preferences.

Computing the mixed equilibrium or equilibria of a finite n-person game follows the same general approach as for two-person zero-sum games. Here too the difficulty is to identify the support of the equilibrium strategies. In a two-person games, we can always find at least one equilibrium with two supports of equal sizes, but this is not true any more with three or more players. Once this is done we need to solve a system of linear equalities and inequalities.

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

Example 5 *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 players 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.

We look for a symmetrical equilibrium in mixed strategies in which every player steps forward with probability $p^*, 0 < p^* < 1$. Then each player is indifferent between stepping forward or not. The latter gives the expected utility $b(1-(1-p)^{n-1})$, the former gives[1]

$$b - c(\sum_{k=0}^{n-1} \frac{\binom{n-1}{k}}{k+1} p^k (1-p)^{n-1-k}) = b - c\frac{1-(1-p)^n}{np}$$

Therefore 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.

[1]Setting $f(p) = p \cdot (\sum_{k=0}^{n-1} \frac{\binom{n-1}{k}}{k+1} p^k (1-p)^{n-1-k})$, one checks $f'(p) = (1-p)^{n-1}$ so that $f(p) = \frac{1-(1-p)^n}{n}$.

Note that the game has many other equilibria, where only a subset of k players step forward with the corresponding probability $p^*(k)$.

Infinite sets of pure strategies

Existence of a Nash equilibrium in mixed strategies holds under the same assumptions as Glicksberg theorem for two person zero-sum games, namely strategy sets are convex and compact, and utility functions are continuous. Here is an example.

Example 6 *war of attrition (a.k.a. all-pay second price auction)*
We revisit the game of timing in Example 7 Chapter 2, specifying VNM utilities. 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.

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

where K is the number of largest bids.

In addition to the pure equilibria described in Example 7, Chapter 2, we have one symmetrical equilibrium in completely mixed strategies where each player independently chooses s_i in $[0, \infty[$ according to a cumulative distribution function F. To compute F we assume that all players $2, \cdots, n$ choose s_i according to F and consider the expected payoff of player 1 using the pure strategy s_1:

$$\int_0^{s_1} (p-t)\dot{G}(t)dt - s_1(1 - G(s_1)), \text{ where } G(t) = F^{n-1}(t)$$

Then we write that all pure strategies s_1 give the same payoff to player 1, i.e.e, the above expression is constant in s_1. This gives $p\dot{G}(t) + G(t) = 1$ for all t, hence

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

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 7 *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. As in the previous example we compute the expected payoff to player 1 using his pure strategy s_1 against the mixed strategy of everyone else: $(p-s_1)F^{n-1}(s_1)-s_1(1-F^{n-1}(s_1))$. That this payoff is independent of $s_1 \in [0,p]$ gives

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

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

1.5 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 $1.5 \pm \frac{\varepsilon}{2}$, 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 12 *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*

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

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

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. Indeed the support of any correlated equilibrium must resist the successive elimination of strictly dominated strategies. Furthermore, there is always one correlated equilibrium of which the support resists the successive elimination of weakly dominated strategies.

But as soon as we have several Nash equilibria (pure or mixed) not in a rectangular position, there are more correlated equilibria. In some games this only helps to average between pure equilibria, as in Example 4 above. In other games, correlation allows a considerable improvement upon the Nash equilibrium outcomes.

Example 8 *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 9 *musical chairs*

We have n players and 2 "chairs" (locations), with $n > 5$. 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

$$4\frac{1}{2^n}+1\frac{n-1}{2^n}=\frac{2n-1}{2^n}\ll 2$$

(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 $2+\frac{4}{n-3}$.

1.6 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

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

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

where we use the notation

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

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

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

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, and each player i learns her own type t_i. Thus player i's beliefs are described by the conditional probability $\pi_i(\cdot|t_i) = \pi(\cdot|t_i)$ of π upon learning one's 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[2].

Example 10:
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	$1,2$	$0,1$		T	$1,3$	$0,4$
B	$0,4$	$1,3$		B	$0,1$	$1,2$
t_1	L	R		t_2	L	R

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 Bayesian equlibrium in each matrix separately, which makes no sense given player 1's information.

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's 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

[2]Consider a 2×2 two-person zero-sum game where if $t_1 = t_2$ the game has a value of $+1$, whereas if $t_1 \neq t_2$ the value is -1. If player 1 (resp. player 2) believes $t_1 = t_2$ (resp. $t_1 \neq t_2$) for sure, both players, ex ante, "win".

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 Bayesian 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 11 *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 move, 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 does not; 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

$$RR \to SS \text{ (or } SF) \to RP \text{ (revealing)} \to FS \to RR; \ PR \to SF \to RP \to \cdots; \ PP \to FF \to RP \to \cdots$$

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

$$\begin{array}{ll} \text{Ann:} & p\delta_S + p'\delta_F \text{ if Bob raises; } S \text{ if Bob stays put} \\ \text{Bob:} & R \text{ if High; } q\delta_R + q'\delta_P \text{ if Low} \end{array}$$

The equilibrium conditions are

$$\begin{array}{ll} \text{for Ann:} & \frac{1}{1+q}(-10) + \frac{q}{1+q}(10) = -1 \Rightarrow q = \frac{9}{11} \\ \text{for Bob:} & p(-10) + p'(1) = -4 \Rightarrow p = \frac{5}{11} \end{array}$$

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

Example 12: *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

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

Player i chooses 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.

Check first that $x(\cdot)$ must be increasing. Fix $t, t', t < t'$, and set $p = \pi\{x(t_j) < x(t) \text{ for all } j \neq i\}$, $p' = \pi\{x(t_j) < x(t') \text{ for all } j \neq i\}$. The equilibrium conditions at t and t' give respectively

$$\{(t - x(t))p > (t - x(t')p', \text{ and } (t' - x(t'))p' > (t' - x(t)p\} \Rightarrow (t' - t)(p' - p) \geq 0$$

and the desired conclusion. Similar arguments show that $x(\cdot)$ must be continuous and differentiable.

Now we write that $z \to (t - x(z))F^{n-1}(z)$ reaches its maximum at t, for all t. 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 $a, t, a < t$

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

(where the first inequality follows from the fact that types are identically distributed, and the second from the fact they are stochastically independent). This says that the equilibrium bid is the expected value of the second highest bid, conditional on your own bid winning the object.

For instance assume the uniform distribution on $[0, 100]$, so that $F(t) = t$, then $x(t) = \frac{n-1}{n}t$ and the expected highest bid (revenue of the seller) is

$$E[x(t_{(1)})] = \frac{n-1}{n}E[t_{(1)}] = \frac{n-1}{n+1}100$$

Moreover the efficient buyer (the one with the highest valuation) gets the object, therefore the expected joint surplus to the seller and bidders is $E[t_{(1)}] = \frac{n}{n+1}100$. 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 13 *sealed bids double auction (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. Suppose the buyer plays $y(b) = b$, and the seller is of type a; his profit $\int_0^x (a - \frac{x+b}{2})db$ is maximized at $x = \frac{2}{3}a$.

We compute the *linear equilibrium*, where each player uses a bid function that is linear in own valuation

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

If the buyer uses $y(\cdot)$ above, trade will occur if the seller's offer x is such that $x \leq y(b) \Leftrightarrow b \geq \frac{x-\delta}{\gamma}$. The expected profit of a type a seller offering x is

$$\int_{\frac{x-\delta}{\gamma}}^{300} (\frac{x + \gamma b + \delta}{2} - a)db = \frac{1}{2\gamma}\{-\frac{3}{2}x^2 + (300\gamma + 2a - \delta)x + \text{ constant}\}$$

It is maximized at $x = \frac{1}{3}(2a + 300\gamma + \delta)$. Similarly if the seller uses $x(\cdot)$ above, the expected profit of a type b buyer offering y is

$$\int_0^{\frac{y-\beta}{\alpha}} (b - \frac{y + \alpha a + \beta}{2})da = \frac{1}{2\alpha}\{-\frac{3}{2}y^2 + (2b + \beta)y + \text{ constant}\}$$

maximized at $y(b) = \frac{2b+\beta}{3}$. Thus the unique candidate linear equilibrium is

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

It remains to check that participation is voluntary, i.e., no one would prefer to abstain from bidding. A buyer of type $b < 75$ bids *above* his own valuation, $y(b) > b$, but as the seller's offer is never below \$75, such an offer is never accepted. Similarly a seller of type $a > 225$ bids $x(a) < a$, but again, this offer is irrelevant as $y(b) \leq 225$ for all b.

Finally we compute the welfare loss at this equilibrium. Trade occurs only if $x(a) \leq y(b) \Leftrightarrow b \geq a + 75$. Therefore the loss is

$$\frac{1}{300^2} \int\int_{a \leq b \leq a+75} (b - a)dadb = \frac{125}{16} \simeq 7.8$$

so about 16% of the efficient expected surplus

$$\frac{1}{300^2} \int\int_{a \leq b} (b - a)dadb = 50$$

It is important to keep in mind that the liner equilibrium is but one equilibrtium among many others, non linear equilibria. Computing all equilibria of the double auction game is an open problem. See problem 20 for an example and Problem 21 for alternative trade mechanisms in the same context.

1.7 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 4

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:

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

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

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

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^*:

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

Problem 7 *all-pay second price auction*

This is a variant of example 6 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

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

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 7 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 $a_1 \ge a_2$. 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*
This is a variant of Example 12 Chapter 2 where the two players value the prize respectively at a_1 and a_2. Each player bids \$$s_i$, where $s_i \in \mathbb{R}_+$ (instead of integers in Example 12, Chapter 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 shop owners 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 4, 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 9 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:

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

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 $1 \rightarrow 2 \rightarrow 3 \rightarrow 1$ 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 $1 \rightarrow 2 \rightarrow 3 \rightarrow 1$.
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 10, and the payoffs are

$$\begin{array}{ccc} T & 1,2 & 0,0 \\ B & 0,0 & 2,1 \\ t_1 & L & R \end{array} \qquad\qquad \begin{array}{ccc} T & 0,0 & 3,1 \\ B & 1,3 & 0,0 \\ t_2 & L & R \end{array}$$

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

Mob becomes very strong in fighting on the day he uses drugs, otherwise he is weak. No matter, whether he used drugs or not, Mob is often involved in conflicts of the type described below.

Bob has just insulted Mob in the bar, and Mob must decide whether to fight Bob immediately, or to leave and try to beat Bob after Bob leaves the bar in a couple of hours. If Mob leaves and tries to catch up with Bob later outside, then, if Mob is strong today, he beats Bob and gets utility 10. However, if Mob is weak, Bob beats him and Mob gets -10.

If Mob decides to fight immediately then it is Bob's choice whether to fight or to leave. If Bob leaves, Mob gets utility 5 from humiliating Bob. If they fight in the bar, then on the day Mob is strong he would beat Bob publicly and get utility 20. However, on the day Mob is weak he would loose to Bob publicly and get utility -30.

Mob knows whether he took drugs this day. Bob does not know it, but he was told by the bar owner that Mob uses drugs on average one day out of three.

If Bob is challenged, he gets -10 from leaving, -15 if he fights and looses and 5 if he fights and wins. If the fight is postponed Bob gets -6 from loosing it and 3 from winning.

a) Describe the set of pure strategies for each player, and write the game matrix. Eliminate dominated strategies.

b) Find all Nash equilibria of this game.

Problem 19 *all-pay first price auction*
The game is identical to that in Example 7, 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 11) and in the second price auction. Compare the expected profit of a bidder in these three auctions.

Problem 20 *sealed bid double auction*
In the game of Example 13, consider the following 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 it 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 13.

Problem 21 *alternative trade mechanisms*
As in Example 13, we have a buyer and a seller with IID valuations in $[0, 300]$.
a) Consider the following take it or leave it mechanism: the seller chooses a price $x \in [0, 300]$, which the buyer accepts or not. Compute its unique Bayesian equilibrium, and compare its welfare loss to that found in Example 13, and in Problem 19. also compare the division of the surplus between the two players.
b) Consider the following mechanism. After the seller and buyer independently bid respectively x and y

trade occurs at price $\frac{y}{2}$ if $y \geq 3x$ and $x + y \leq 300$
trade occurs at price $\frac{x}{2} + 150$ if $y \geq \frac{x}{3} + 200$ and $x + y > 300$
no trade occurs, and no money changes hands, in every other case

Show that sincere report of one's valuation ($x(a) = a$ and $y(b) = b$ for all a, b) is a Bayesian equilibrium. Compare the welfare loss of this mechanisms to those found in Example 13 and in Problem 19.

Problem 22 *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?

2 Chapter 4: extensive form games

The general model of n-person games in extensive form is a straightforward extension of the model in sectiion 1.3 for two-person zero sumgames.

2.1 Definition

An n-person game in extensive form Γ^e is given by:

1) a set of players $N = \{1, ..., n\}$;

2) a tree (a connected graph without cycles), with a particular node taken as the root;

3) for each non-terminal node, a specification of who has the move (one of real players or "chance");

4) a partition of all nodes, corresponding to each particular player, into information states, which specify what players know about their location on the tree;

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

Formally, a rooted tree is a pair (M, σ) where M is the finite set of nodes, and $\sigma : M \to M \cup \emptyset$ associates to each node its predecessor. A (unique) node m_0 with no predecessors (i.e., $\sigma(m_0) = \emptyset$) is the root of the tree. Terminal nodes are those which are not predecessors of any node. Denote by $T(M)$ the set of terminal nodes. For any non-terminal node r, the set $\{m \in M : \sigma(m) = r\}$ is the set of successors of r. We call the edges, which connect m with its successors, "alternatives" at m. The maximal possible number of edges in a path from the root to some terminal node is called the length of the tree.

Given a rooted tree (M, σ), the game in extensive form is specified once we label all the nodes and edges according to the following rules.

(a) Each non-terminal node (including the root) is labeled by number from $\{0, 1, ..., n\}$, where $i \in \{1, ..., n\} = N$ represents a real player in the game, and 0 represents a "nature" or "chance" player. We denote by M_i the set of nodes labeled by the player i. The interpretation is that when the game is played, we start at the root and then for each node $m \in M_i$ the player i is choosing which edge to follow from this node.

(b) The alternatives at a node labeled by the chance player 0 are labeled by numbers from $[0, 1]$, so that those numbers over all the alternatives sum to 1. They represent probabilities that chance would choose those alternatives.

(c) The alternatives at a node $m \in M_i,\ i \in \{1, ..., n\}$ are labeled by "move labels". Different alternatives at the same node are labeled with different labels.

(d) Each $M_i,\ i \neq 0$, is partitioned into information sets $P^i_1, ..., P^i_{k_i}$, $M_i = \bigcup\limits_j P^i_j$, $P^i_{j_1} \cap P^i_{j_2} = 0$, with the following condition: any two nodes x, y from the same information set must have the same number of successors, and the set of move labels on the alternatives at x should coincide with the set of move labels on the alternatives at y. The interpretation is that when a player i has to choose an alternative at the node $m \in M_i$, he knows in what information set he is, but he does not know at what exact node from this information set he is making his choice.

(e) Each terminal node m is labeled by a vector $u(m) = (u_1, ..., u_n)$ which specifies the payoffs for players $1, ..., n$, if the game ends at this node. This defines the payoff function $u : T(M) \to \mathbb{R}$.

The game starts at the root m_0 of the tree. For each non-terminal node m, $m \in M_i$ means that player i has the move at this node. A move for the chance player consists in choosing the successor of m randomly according to the probability distribution on the alternatives at m. A move for a real player $i \in N$ consists in picking a move label for the successor of this node. Note that when making the move, a player does not know where exactly he stands. He only knows the information set he is at, and hence the set of the move labels. Once a move label is picked, the game moves to the successor of the node m which is connected to m by the alternative with the chosen move label. The game continues until some terminal node m_t is reached. Then a payoff $u(m_t)$, attached to this node, is realized.

An important special case: When each information set of each player consists of a single node, we say that this game has "**perfect information**". This term refers to the fact that, when a player has to move, he possesses perfect information about where exactly in the tree he is.

Normal form games as extensive form games: any normal form game can be represented in extensive form, by ordering the players arbitrarily say $1, 2, \cdots, n$, have player 1 move first, after which the information set of other players "hides" 1's move, then player 2 moves, after which the information set of the remaining players hides the first two moves, etc.. In this fashion we can also represent multi stage games where at some nodes, several players move simultaneously.

Conversely there is a canonical normal form representation Γ of any extetensive form game Γ^e. A strategy for a player i is a complete specification of what move to choose at each and every information set from $P = \{P_1^i, ..., P_{k_i}^i\}$. The set of all such possible specifications is the strategy set C_i for player i in Γ. The payoff $u_i(c_1, ..., c_n)$ is the payoff to player i at the terminal node which is reached after all players have chosen all their moves according to the strategies $c_1, ..., c_n$. It is important to note that, since there are chance players in the extensive form game who make their choice at random, the game could have an uncertain outcome even when all real players use pure strategies. In this case the game could end in different terminal nodes, but we can calculate the probability of our game to end in each terminal node (given choice of strategies $c_1, ..., c_n$). Then, the payoff payoff $u_i(c_1, ..., c_n)$ will be the expected payoff according to those probabilities.

As usual, we assume that players evaluate different outcome on the basis of a VNM (expected) utility function.

As for normal form games we define the mixed strategy $s_i \in \Delta(C_i)$ for player i as a probability distribution on his set of pure strategies C_i. The best response correspondence is defined by $br_i(s_{-i})$ to be the set of strategies for player i that give him the best (expected) payoff against the vector s_{-i} of strategies of other players. A Nash equilibrium of the extensive form game Γ^e is the vector $s = (s_1, ..., s_n)$ of strategies, where each one is a best response to the others.

2.2 Subgame perfection

In a game in extensive form, the set of the Nash equilibria is often very big and some of those equilibria make little sense.

Consider for instance the extensive form variant of the Nash demand game (example 6 in Chapter 2) with perfect information. Demands are in cents (they divide $1), player 1 chooses his demand x, which is revealed to player 2, who can only accept or reject it. For any integer $x, 0 \leq x \leq 100$, the pair of strategies where player 1 demands x, and player 2 rejects if $s_1 > x$, and accepts if $s_1 \leq x$, is a Nash equilibrium. But for $x \leq 50$, this equilibrium involves the unrealistic refusal of a fair share of the pie.

The key concept of subgame perfection is an important *refinement* that will eliminate many such unrealistic outcomes. We define it first, before illustrating its predictive power and its limits in a handful of examples.

We assume that our game has *perfect recall*. Thus, in the course of the game each player remembers his past moves. In particular, it implies some restrictions on the information sets. Two nodes x, y cannot belong to the same information set of the player i, if the choices in the game he made before reaching x or y allow him to distinguish between the two. For example, no game path (a path from the root to a terminal node) could contain several nodes from the same informational set.

A proper subgame of an extensive form game Γ^e is a subtree starting from some non-terminal node, with all the labels, such that any information set which intersects with the set of nodes in this subtree, is fully contained in that set of nodes. Thus, the fact that a player knows that a subgame is being played does not give him any additional information to refine his information structure.

Definition 13 *A subgame perfect equilibrium for the extensive form game Γ^e is a Nash equilibrium whose restriction to any subgame is also a Nash equilibrium of this subgame.*

in the variant of the Nas demand game just discussed, there are exactly two subgame perfect equilibria: player 1 demands 100, and player 2 accepts any demand; player 1 demands 99 and player 2 accepts any demand of 99 or less, but rejects the demand 100. Note that an extensive experimental testing of this game reveals that such a strategy typically fails, because the utility of player 2 depends on more than the amount of money he takes home.

Example 1 Consider the following extensive form game with perfect information. Player 1 decides whether to go left or right. Knowing his choice, player 2decides whether to go up or down. The payoffs are $u(left, up) = (3, 1)$, $u(left, down) = (0, 0)$, $u(right, up) = (0, 0)$, $u(right, down) = (1, 3)$.

In this game player 1 has two strategies (left and right), while player 2 has four strategies, since one has to specify for her what to do if player 1 chooses left as well as what to do if he chooses right. Thus, her strategy set is $\{(up_l, up_r), (up_l, down_r), (down_l, up_r), (down_l, down_r)\}$, where subindex l is for her choice after player 1 goes left, and subindex r is for her choice after player 1 goes right. Note that if player 2 would not know the choice of player 1 at a time she makes her own choice, then it would be the Battle of Sexes game, in which each player has just two strategies.

This game has two proper subgames, in each only player 2 is to make a move. The whole game has three Nash equilibria in pure strategies. They are $(left, (up, down))$, $(left, (up, up))$, $(right, (down, down))$. However, only first of them is subgame perfect. Player 2 would prefer the last one, where she gets 3, by threatening player 1 to choose terminal node with zero payoffs if he goes left. But it is not sustainable under the subgame perfection assumption, since if player 1 actually moves left player 2 will have strong incentive to choose the node with payoffs $(3, 1)$ and she has no way to pre-commit herself not to do it.

Theorem 14 *Any finite (i.e., based on a finite tree) game Γ^e in extensive form has at least one subgame perfect equilibrium.*

The proof is by induction in the number of proper subgames the game Γ^e has. If it has no proper subgames, then any Nash equilibrium of the corresponding normal form game will be a subgame perfect equilibrium of Γ^e. Now, consider a subgame $\Gamma^{e\prime}$ of Γ^e which has no its own proper subgames. It has (at least one) Nash equilibrium; pick up one of those. Substitute this whole subgame $\Gamma^{e\prime}$ by a new terminal node for Γ^e, located at the root of this subgame $\Gamma^{e\prime}$. Label this new terminal node with the payoffs from the chosen Nash equilibrium of Γ^e. We thus constructed a new game Γ_1^e which has less proper subgames, and hence has a subgame perfect equilibrium vector of strategies by induction hypothesis. Now, we add to the strategies in this equilibrium for Γ_1^e vector the specification for each player of what to do in $\Gamma^{e\prime}$, namely the prescription to play according to the Nash equilibrium we have picked for $\Gamma^{e\prime}$. It is easy to check that the resulting vector of strategies will be the subgame perfect equilibrium of Γ^e.

Theorem 15 *Any finite game Γ^e in extensive form with perfect information has at least one subgame perfect equilibrium in pure strategies. If for any player all payoffs at all terminal nodes are distinct, then this equilibrium is unique.*

It is easy to see that such subgame perfect equilibrium in pure strategies can be always found by backward induction, starting from the end (by seeing for every node, whose all successors are terminal nodes, what should be the choice there, and then proceeding by induction).

Leader-follower equilibrium

Given a two person game in normal form (S_1, S_2, u_1, u_2), the extensive form game where player i chooses his strategy s_i first, this choice is revealed to player j who then chooses s_j, is called the $i-$Leader,$j-$Follower game. When we speak below of the $i-$Leader,$j-$Follower equilibrium, we always mean its subgame perfect equilibrium, or equilibria.

Comparing the $i-$L,$j-$F equilibrium with the Nash equilibrium (or equilibria) of the initial normal form game, gives useful prediction about commitment tactics in that game. Clearly player i always prefers (sometimes weakly) the $i-$L,$j-$F equilibrium to any of the Nash equilibria. But there are no other restrictions on the comparison of i's $j-$L,$i-$F equilibrium payoff with the two above.

In two-person zero sum games with a value, or in a game with (strictly) dominant strategy, the L-F equilibrium and Nash equilibrium coincide: it does not matter if we choose strategies simultaneously and independently, or sequentially with the first choice being revealed.

In the Battle of the Sexes, in the war of attrition (example 7 chapter 2 and example 6 chapter 3), as well as in the simple Cournot duopoly of example 10 chapter 2, both players prefer to be the leader. In the former two, the leader-follower equilibria coincide with the pure strategy Nash equilibria; in the latter case i's payoff in the $i-$L,$j-$F equilibrium is larger than at the unique Nash equilibrium, whereas j's payoff is lower.

In two-person zero sum games without a value, both players obviously prefer to be follower. The same is true in the following game of timing.

Example 2 *grab the dollar*

This is a symmetrical game of timing with two players. Both functions a and b increase with $a(t) > b(t)$ for all t, and $b(1) > a(0)$. Recall that $a(t)$ is the payoff to the player who cries stop at t. If both stop at $t = 0$, thay both get $a(0)$; if they both stop at $t = 1$, they both get $b(1)$. The normal form game has a unique Nash equilibrium; the Leader-Follower equilibrium favors the Follower, but they both prefer it to the Nash equilibrium of the normal form.

A common difficulty with the interpretation of subgame perfect equilibrium selection is that it involves *imprudent* strategies.

Consider *Kalai's hat game*: a hat passes around the n players; each can put a dollar or nothing in the hat; if all do, they get back \$2 each; if one or more put nothing in the hat, all the money in the hat is lost. There are two Nash equilibria: all put \$1 or nobody does; the former is the s.p. equilibrium, but, unlike the latter, its strategies are imprudent.

The next example is a celebrated paradoxical game.

Example 3 *Selten's chain store paradox*

There are $20 + 1$ players. The incumbent meets successively the 20 *small* potential entrants. At every meeting, the following game takes place: first stage: the small firm chooses to enters or stay out; in the latter case payoffs are $(0, 100)$ to small firm and incumbent; if small firm enters, the incumbent chooses to collude or fight, with corresponding payoffs $(40, 50)$ and $(-10, 0)$ respectively. The only s.p. equilibrium is that all small firms enter, and collusion occurs every time.

Now suppose you are small firm #17 and the incumbent has been challenged 5 times and has fought every time, what do you do? It is certainly imprudent to enter! The other Nash equilibrium where the incumbent is committed to fight every period seems more plausible.

On the other hand, the s.p. equilibrium may display excessive prudence as in the following game.

Example 4 *Rosenthal's centipede game*

This is a multi-stage version of grab the dollar (example 2 above), where the pot starts empty, and grows by 1 cent every period. In odd periods, player 1 can grab half of the pot plus one cent, and leave the rest of the pot to player 2, or do nothing and let the pot grow till next period; in even periods player 2 can grab half of the pot plus one cent, and leave the rest to player 1, or do nothing and let the pot grow till next period. The game lasts for 100 periods. In the last period player 2 gets $51c$ and player 1 gets $49c$.

In the subgame perfect equilibrium, player 1 grabs $1c$ in period 1 and player 2 gets nothing. This is actually the only Nash equilibrium of the game!

2.3 Subgame perfect equilibrium in infinite games

When the number of stages in the game is infinite, the computation of s.p. equilibria becomes more tricky, and can lead to much indeterminacy or to a deterministic prediction. A famous example follows.

Example 5 *Rubinstein's alternating offers bargaining*

The two players divide a dollar by taking turns (starting with player 1) making offers. The first accepted offer is final. No money is handed out until an offer is accepted. Player i's discount rate is $\delta_i, 0 \leq \delta_i \leq 1$: receiving $\$x$ in period k is worth $\$x(\delta_i)^{k-1}$ in period 1. (Alternative interpretation: after each rejected offer, there is a chance $(1-\delta)$ that the game ends with no one getting any money).

Case 1: no impatience, $\delta_1 = \delta_2 = 1$ (or no risk of the game terminating). If the number of periods is finite, whoever makes the last offer acts as the Leader in a Nash demand game, therefore keeps essentially the whole dollar. If the game never stops, infinite number of periods (and disagreement for ever yields zero profit to both players), any division $(x, 1-x)$ of the dollar is a subgame perfect equilibrium outcome. It is achieved by the *inflexible* strategies where player 1 (resp. 2) refuses any offer below x (resp. below $1-x$) and accepts any offer weakly above x (resp. weakly above $1-x$), and the first offer is $(x, 1-x)$.

Case 2: impatient players, $\delta_1 < \delta_2 < 1$

Check first that the inflexible strategies around $(x, 1-x)$ described above, form a Nash equilibrium, but not *not* a subgame perfect equilibrium. Say in his first move player 1 offers $(y, 1-y)$ where $1-x > 1-y > \delta_2(1-x)$. Player 2's inflexible strategy is to say No, however in the subgame starting in period 2 where inflexible strategies are used, player 2 cannot hope any more than $\delta_2(1-x)$, therefore No to $(y, 1-y)$ is not part of any equilibrium strategy in this subgame. The inflexible strategies are not subgame perfect because they contradict the equilibrium rationale in some out of equilibrium subgame.

We show now the equilibrium is unique, and compute the corresponding shares.

Observe first that in any s.p.eq. outcome, agreement takes place immediately. Indeed suppose for instance agreement takes place in period 2 at $(z, \delta - z)$, then player 1 can offer $(z + \frac{1-\delta}{2}, \frac{1+\delta}{2} - z)$ to player 2, a better result for both players, which player 2 should accept under subgame perfection.

Next the set of s.p.eq. outcomes can be shown to be closed, hence compact, so we can talk of the best or worst s.p. share for either agent.

In a s.p.eq. where 1 speaks first, if his offer is rejected we go to a s.p.eq. where 1 speaks second. Hence the best s.p.eq. for 1 when he speaks first is the one followed by the worst s.p.eq. of 2 in the game where 2 speaks first. Let x be 1's share in his best s.p.eq. when he speaks first, and y be player 1's share in his best s.p.eq. when he speaks second. Because the offer $1 - x$ is accepted by 2 in that s.p.eq., we have

$$1 - x = \delta_2(1 - y)$$

Next consider player 2: the worst s.p.eq. for 2 when he speaks first is the one followed by the best s.p.eq. of 1 in the game where 1 speaks first. Because the offer y is accepted by 1 in that s.p.eq., we have

$$y = \delta_1 x$$

We can symmetrically look at the worst s.p.eq. share x' for 1 in the game where he speaks first, and worst s.p.eq. share y' in the game where he speaks second. Check that x', y' satisfies the same system of equations as x, y, implying $x = x'$ and $y = y'$, i.e., the s.p. equilibrium outcome is unique. When player 1 speaks first it is

$$(x, 1 - x) = (\frac{1 - \delta_2}{1 - \delta_1\delta_2}, \frac{\delta_2(1 - \delta_1)}{1 - \delta_1\delta_2})$$

It remains to show that in the game where player 1 speaks first, the following strategies form a s.p.eq.:

player 1 always offers $(x, 1 - x)$, rejects any offer below y, accepts any offer y or more;

player 2 always offers $(y, 1 - y)$, rejects any offer below $1 - x$, accepts any offer $1 - x$ or more;

Our last example involves only two stages but many players. It illustrates the techniques to compute s.p.equilibria in this context.

Example 6 *durable goods monopoly*

A monopolist produces at zero cost a durable good. There are 1000 consumers, with reservation prices for the good uniformly distributed in the interval $[0, 100]$. The common discount rate of the monopolist and consumers is δ. If the monopolist can commit himself to a fixed pricing policy at the beginning of the game, his best choice is a constant price of 50. Consumers are impatient, so the upper half will buy immediately, for a monopolist profit of 25,000 and consumer surplus 12,500. However it is more realistic to assume the monopolist cannot commit ex ante for both periods; in period 2, he wants to cut his price to extract a little more surplus from the consumers who did not buy in the 1st period. But $p_1 = 50, p_2 = 25$ is not an equilibrium, because a consumer who values the object at \$51 prefers to wait for the "sale" rather than buying immediately.

Say p_1 is the price charged in period 1, and all consumers with valuation in $[\varphi(p_1), 100]$ buy in period 1; then in period 2, a regular monopoly situation, the price will be $\frac{\varphi(p_1)}{2}$ and all agents in $[\frac{\varphi(p_1)}{2}, \varphi(p_1)]$ will buy. Equilibrium conditions in period 1:

for consumers

$$\varphi(p_1) - p_1 = \delta(\varphi(p_1) - \frac{\varphi(p_1)}{2}) \Longleftrightarrow \varphi(p_1) = \frac{p_1}{1 - \frac{\delta}{2}}$$

for the monopolist

$$p_1 \text{ maximizes } (100 - \varphi(p_1))p_1 + \delta(\frac{\varphi(p_1)}{2})^2$$

hence

$$p_1 = \frac{(1 - \frac{\delta}{2})^2}{1 - \frac{3\delta}{4}} 50 < 50; \varphi(p_1) = \frac{1 - \frac{\delta}{2}}{1 - \frac{3\delta}{4}} 50 > 50$$

finally both monopoly profit and consumer surplus go up, relative to the non strategic $p_1 = 50, p_2 = 25$.

2.4 Other refinements of Nash equilibrium.

When we represent an extensive form game in the normal form, the normal form could have multiple equilibria which are "behaviorally" the same. For example, assume that player 1 makes move two times. The first time he chooses a or b, and the second time he chooses c or d. This results in four strategies (a, c), (a, d), (b, c), and (b, d). The unique (behaviorally) mixed strategy "play a or b, with probability 1/2 each, at the first move, and play c or d, with probability 1/2 each, at the second move", can be represented in a continuum ways as a mixed strategy in normal form representation, as $p(a, c)+(1/2-p)(a, d)+(1/2-p)(b, c)+p(b, d)$ for any $p \in [0, 1/2]$.

Another way to view Nash equilibrium of an extensive form game is looking at its multiagent representation. Namely, assume that each player i is represented by several agents, one for each of his information sets. All those agents have the same payoffs (same as player i). Each agent acts at most once in the game — if and when the game path goes through the corresponding information set — and at the moment this agent acts he has no additional information compared with what he knew before the game started. Hence, we can regard our game as a game where all players (i.e., all agents) simultaneously and independently choose each a strategy from his strategy set (which is the set of move labels for the information set for which an agent is responsible). This game hence can be viewed as a normal form game.

The Nash equilibria of the original extensive form game can be defined as the Nash equilibria of thus constructed normal form game which is called its multiagent representation. The problem with this definition is that agents are precluded from cooperation. Thus, we get unrealistic equilibria.

Example 7 Consider an extensive form game where agent 1 first chooses a or b. Without knowing his choice, agent 2 then chooses x or w. If agent 1 has chosen b initially, then the game ends there. If agent 1 has chosen a initially, he has now to choose between y and z, without knowing the choice of agent 2. The payoffs are $(3,2)$ for (b,x), $(2,3)$ for (b,w), $(4,1)$ for (a,x,z), $(2,3)$ for (a,x,y), $(0,5)$ for (a,w,z), and $(3,2)$ for (a,w,y). It is easy to check that (b,w,z) is a Nash equilibrium of the multiagent representation of this game, but not an equilibrium of its normal form. The last follows from the fact that player 1's best response to w is not (b,z), but (a,y).

The way to deal with this is to consider as Nash equilibria of extensive form game Γ^e only those equilibria of its multiagent representation which survive as Nash equilibria of the normal representation of initial game Γ^e. These equilibria are called Nash equilibria in behavioral strategies. They always exist for finite extensive form games.

2.4.1 Sequential rationality

Sequential rationality is a generalization of the subgame-perfect equilibrium): the idea that the choice in each information set should be rational (i.e. a best response), given what the player believes about what are the chances for him to be at each particular node from this information set. These beliefs are assumed to be formed by Bayesian update. This idea results in the notion of sequential equilibrium (they always exist for finite extensive form games).

A **sequential equilibrium** is (s,π), a vector of behavioral strategies plus a vector of Bayesian consistent beliefs for all nodes (conditional probabilities that we are at each particular node, given that we are in the information set including this node), such that given those beliefs it is sequentially rational for the players to follow the prescribed strategies. The proper definition includes the way to define the consistency of π for the nodes that have a zero probability to be on the game path under s. It is done by assuming that there exists a sequence of "tremblings" of s, which assign a positive probability to each pure strategy and converge to s, such that the belief about the node with zero probability is the limit of the Bayesian updated beliefs for those tremblings (see below the definition of trembling hand equilibrium).

2.4.2 Trembling hand perfect equilibrium

Trembling hand perfection is the refinement of Nash equilibrium which applies to the normal form games. Consider $\Gamma = (N, (C_i)_{i\in N}, (u_i)_{i\in N})$ with all C_i finite, $S_i = \Delta(C_i)$. A vector of mixed strategies $s \in S$ is a (trembling hand) perfect equilibrium of this game if there exists a sequence of $s^k \in S$, $k = 1, 2, ...$, such that

(1) Any s_i^k is completely mixed strategy, i.e. all pure strategies from C_i belong to its support (are used with positive probability)

(2) $\lim_{k\to\infty} s_i^k(c_i) = s_i(c_i)$ for all $i \in N$, all $c_i \in C_i$ (i.e., s^k converges to s)

(3) $s_i \in \arg\max_{t_i \in S_i} u_i(t_i, s_{-i}^k)$ for all $i \in N$, and all k

I.e., s is a (trembling hand) perfect equilibrium if there exists a sequence of "tremblings" (completely mixed strategies, ones which could end up in using any pure strategy with positive probability, even the most unreasonable one), such that this sequence converges to s, and that each strategy s_i in s is a best response to any of those tremblings made by all players other then i.

The following theorems we will not prove.

Theorem 16 *For any $\Gamma = (N, (C_i)_{i\in N}, (u_i)_{i\in N})$ with all C_i finite there exists at least one (trembling hand) perfect equilibrium.*

Theorem 17 *If Γ^e is an extensive form game with perfect recall, and s is a trembling hand perfect equilibrium of the multiagent representation of Γ^e, then there exists a vector of beliefs π, such that (s, π) is a sequential equilibrium of Γ^e.*

Note that the existence of sequential equilibria follows from these two theorems.

2.5 Problems for Chapter 4

Problem 1 *leader follower equilibrium*

In each case, compare the two leader follower equilibria with the Nash equilibrium (or equilibria) of the normal form game. If the game defined earlier is among n players, simply consider the two player case.

a) variant of the grab the dollar game (example 2 chapter 4) where a and b increase and $b(t) > a(t)$ for all t.

b) in the coordination game example 8 chapter 2

c) in the public good provision game of example 20 chapter 2

d) in the war of attrition with mixed strategies, example 6 chapter 3

e) in the (mixed strategies) lobbying game of example 7 chapter 3.

Problem2 *King Solomon*

King Solomon hears from two mothers A and B who both claim the baby but only one of them is the true mother. Both mothers know who is who, but Solomon does not. However Solomon knows that the baby s worth v_1 to the true mother and v_2 to the false one, with $v_2 < v_1$. He has them play the following game.

Step 1 Mother A is asked to say "mine" or "hers". If she says "hers" mother B gets the baby and the game stops. If she says "mine" we go to Step 2. Mother B must "agree" or "challenge". If she agrees mother A gets the baby and the game stops; if she challenges, mother B pays v and keeps the baby, whereas mother A pays w. These two numbers are chosen so that $v_2 < v < v_1$ and $w > 0$.

Show that in the subgame perfect equilibrium of the game, the true mother gets the baby. What about the money?

Problem 3 *Bertrand duopoly*

Two firms are located town A and B respectively; in each town there is d units of inelastic demand with reservation price p (the same in each town); transportation cost between a and B is t. Thus we have a symmetrical game with strategy set $[0, p]$ and payoff

$$\begin{aligned} u_1(s_1, s_2) &= ds_1 \text{ if } |s_1 - s_2| \leq t \\ &= 2ds_1 \text{ if } s_1 + t < s_2; = 0 \text{ if } s_2 + t < s_1 \end{aligned}$$

(note that when t is exactly the price difference, customers does not travel; the opposite assumption would do just as well).

a) Show the game has no Nash equilibrium if $2t < p$. Compute the Nash equilibrium (or equilibria) if $p \leq 2t$.

b) Compute the Leader-Follower equilibria and show that a firm always prefers to be Follower.

Problem 4 *leader-follower equilibrium*

In this problem we restrict attention to finite two-person games (S_1, S_2, u_1, u_2) in pure strategies, such that the mappings u_1 and u_2 are one-to-one on $S_1 \times S_2$. Therefore the best reply strategies are unique, and so are the $1-$L,$2-$F and $2-$L,$1-$F equilibria. Denote the corresponding payoffs L_i and F_i.

Suppose $L_i = F_i$ for $i = 1, 2$. Show that the $1-$L,$2-$F and $2-$L,$1-$F equilibria coincide, are a Nash equilibrium, Pareto superior to any other Nash equilibrium.

Problem 5 *three way duel (Dixit and Nalebuff)*

Larry, Mo and Curly play a two rounds game. In the 1st round, each has a shot, first Larry then Mo then Curly. Each player, when given a shot, has 3 options: fire at one of the other players, or fire up in the air. After the 1st round, any survivor is given a second shot, again beginning with Larry then Mo then Curly.

For each duelist, best outcome is to be the sole survivor; next is to be one of two survivors; inthird place is the outcome where no one gets killed; dead last is that you get killed.

Larry is a poor shot, with only 30% chance of hitting a person at whom he aims. Mo has 80% accuracy, and Curly has 100% accuracy.

Compute the subgame perfect equilibrium of this game, and the equilibrium probabilities of survival.

Problem 6

Ten pirates (ranked from 10 to 1 from the oldest to the youngest) share 100 gold coins. The oldest first submits an allocation of his choice to a vote. If at least half of the pirates (including the petitioner) approves of this allocation, it is enforced. Otherwise, the oldest pirate walks away with no coin, and the same game is repeated with nine pirates, etc. How would you recommend the players to play? (Find the subgame-perfect Nash equilibrium outcomes)

Problem 7

In an extensive form game, a behavior strategy for player i specifies a probability distribution over alternatives at each information set of player i. Mixed strategy, as always, is a probability distribution over the set of pure strategies. Two strategies of player i are called equivalent if they generate the same payoff for player i for all possible combinations c_{-i} of pure strategies of other players. Prove that in a game of perfect recall, mixed and behavior strategies are equivalent.

More precisely: every mixed strategy is equivalent to the unique behavior strategy it generates, and each behavior strategy is equivalent to every mixed strategy that generates it.

Problem 8 *(difficult!)*

Prove that for a zero-sum game any Nash equilibrium is subgame perfect. More precisely, for any outcome which is the result of a Nash equilibrium strategy profile, there is a subgame perfect equilibrium strategy profile which results in the same outcome (an outcome is a probability distribution over the terminal nodes).

Problem 9 *grab a shrinking dollar*

One dollar is placed in the "pot" in period 1; its value will diminish by a discount of δ at each period (after k periods, it is worth δ^{k-1} to both players). The two players take turns, starting with player 1. When i has the move, she has 2 choices: to stop the game, in which case 40% of the pot goes to i and 60% to player j, or to let player j have the next move. The game goes on until someone stops, or if no one does both players get zero.

a) Show that if δ is small enough, the only Nash equilibrium of the game is that player 1 grabs the dollar immediately. Explain "small enough".

b) Is there any value of δ such that in some Nash equilibrium of the corresponding game, someone grabs the dollar after each player has declined to do so at least once?

c) Show that if δ is large enough, there is a subgame perfect equilibrium where player 1 does not grab the dollar, and player 2 does in the next turn. Explain "large enough".

Problem 10 *bargaining with alternating offers*

In this variant of Rubinstein's model (example 5), the only difference is that after an offer is rejected, the flip of a fair coin decides the player who makes the next offer. Successive draws are independent.

a) Assume first the players have a common discount factor δ. Find the symmetrical subgame perfect equilibrium of the game, and show it is the unique s.p. equilibrium.

b) Now we have 2 different discount factors. Compute similarly the s.p. equilibrium or equilibria.

Note: for both questions you must describe the equilibrium offer and acceptance strategies of both players.

Problem 11 *durable goods monopoly*

In the model of example 6, we now assume the good is infinitely durable and the game lasts for ever. A strategy of the monopolist is a stream of prices $(p_1, p_2, \cdots)$ and his profit is $\sum_{t=1}^{\infty} \delta^t p_t q_t$, where q_t is the quantity sold in period t. A consumer with valuation v gets utility $\delta^t (v - p_t)$ if she buys in period t.

Look for a *linear stationary* s.p. equilibrium: facing price p_t at time t, all consumers with valuation λp_t or above (if any are left) buy, others don't. facing an unserved demand $[0, v]$ at time t, the monopolist charges the price μv. Naturally the two constant λ, μ are such that $\lambda \geq 1, \mu \leq 1$.

write the equilibrium condtions resulting in a system to compute λ, μ. Solve the system numerically for $\delta = 0.9$ and $\delta = 0.5$. Deduce the optimal sequence $(p_1, p_2, \cdots)$ and discuss its rate of convergence. Compute the equilibrium profit and consumer surplus.

Problem 12 *last mover advantage in a first price auction*

In the game of Example 12 chapter 3 with *two* bidders, recall that the unique symmetrical equilibrium has a bid function $x(t) = \frac{t}{2}$, and an expected gain of $\$\frac{100}{6}$ for each player.

Suppose now player 2 has the last mover advantage: he observes player 1's bid before bidding himself. Compute the unique subgame perfect equilibrium of this game, and the corresponding expected gains of the players. Compare to the case of simultaneous bids.

Suppose next that player 2 sees player 1's bid but player 1 is unaware of this (and so he plays as in the case of simultaneous bids). Compute the corresponding expected gains of both players.