

Lecture Notes on Game Theory*

Alfredo Di Tillio

January 19, 2008

Contents

1	Normal Form Games	3
1.1	Basic Notations and Definitions	3
1.2	Dominance and Best Response	4
1.3	Iterated Dominance and Rationalizability	5
1.3.1	Iterated Dominance	5
1.3.2	Correlated Rationalizability	6
1.3.3	An Equivalence Result	6
1.3.4	Independent Rationalizability	9
1.4	Examples	9
1.4.1	Independent vs. Correlated Rationalizability	9
1.4.2	Rationalizability in the Cournot Game	10
2	Equilibrium in Normal Form Games	13
2.1	Nash Equilibrium	13
2.1.1	Existence of Nash Equilibria	13
2.2	Strictly Competitive Games	15
2.3	Normal Form Refinements	18
2.3.1	Trembling-Hand Perfect Equilibrium	18

*PhD Microeconomics III, Spring 2009

2.3.2	Proper Equilibrium	18
2.4	Examples	20
2.4.1	Trembling-Hand Perfection and Weakly Dominated Strategies	20
2.4.2	Trembling-Hand Perfection and Properness	21
3	Extensive Form Games: Basics	23
3.1	Formal Definition	23
3.1.1	Normal Form Representation of an Extensive Form Game	25
3.2	Mixed and Behavior Strategies: Kuhn's Theorem	26
3.2.1	Continuation Strategies, Continuation Outcomes, and Continuation Payoffs	29
3.3	Subgame Perfect Equilibrium	30
3.4	Examples	33
3.4.1	Finding Subgame Perfect Equilibria of Perfect Information Games	33
3.4.2	Trembling Hand Perfect and Subgame Perfect	35
4	Extensive Form Games: Further Topics	37
4.1	Sequential Equilibrium	38
4.2	Alternating Offer Bargaining	39
4.2.1	Finite Horizon	40
4.2.2	Infinite Horizon	41
4.3	Examples	42
4.3.1	Inconsistent Assessments and One-Shot Deviations	42
4.3.2	More on Out of Equilibrium Beliefs	42
5	Repeated Games	44
5.1	Discounting	44
5.2	Two Folk Theorems	45
5.3	Examples	47
5.3.1	Nash and Subgame Perfect Equilibria	47
5.3.2	Repeated Prisoners' Dilemma	50
6	Games with Incomplete Information	52
6.1	Bayesian Games	53

1. Normal Form Games

1.1. Basic Notations and Definitions

A **game in normal form** is a list $\langle N, (S_i, u_i)_{i \in N} \rangle$ comprising the following objects:

- a set of **players** N
- for each player i , a set of **strategies** S_i (we define $S := \times_{i \in N} S_i$)
- for each player i , a **payoff** function $u_i : S \rightarrow \mathbb{R}$

A game in normal form is **finite** if the set of players is finite and the set of strategies of each player is finite.

For any finite set X , we write $\Delta(X)$ to denote the set of all probability distributions over X , that is, the set of all functions $p : X \rightarrow [0, 1]$ such that $\sum_{x \in X} p(x) = 1$.

Let $\langle N, (S_i, u_i)_{i \in N} \rangle$ be a finite game in normal form. A **mixed strategy** of player i in this game is an element of $\Sigma_i := \Delta(S_i)$. Assume without loss of generality (just re-labeling players) that $N = \{1, \dots, I\}$, where $I = |N|$. The set $\Sigma_1 \times \dots \times \Sigma_I$ will be abbreviated as Σ . Moreover, we will write S_{-i} and Σ_{-i} as abbreviations for $S_1 \times \dots \times S_{i-1} \times S_{i+1} \times \dots \times S_I$ and $\Sigma_1 \times \dots \times \Sigma_{i-1} \times \Sigma_{i+1} \times \dots \times \Sigma_I$, respectively.

Remark 1. Note that Σ_{-i} is not the same as $\Delta(S_{-i})$. An element of the latter set is a **belief** of player i over his opponents' strategies, i.e. a probability distribution over the set of all profiles of i 's opponents' strategies, S_{-i} . An element of Σ_{-i} is instead a profile of mixed strategies, one for each of i 's opponents.

Given any $s_i \in S_i$ and $s_{-i} = (s_1, \dots, s_{i-1}, s_{i+1}, \dots, s_I) \in S_{-i}$, we will write (s_i, s_{-i}) as an abbreviation for $(s_1, \dots, s_{i-1}, s_i, s_{i+1}, \dots, s_I)$. For all $i, j \in N$, all $\sigma_i \in \Sigma_i$, and all $s_{-i} \in S_{-i}$, we define

$$u_j(\sigma_i, s_{-i}) := \sum_{s_i \in S_i} \sigma_i(s_i) u_j(s_i, s_{-i}).$$

Then, for all $\mu_i \in \Delta(S_{-i})$, we define

$$u_j(\sigma_i, \mu_i) := \sum_{s_{-i} \in S_{-i}} \mu_i(s_{-i}) u_j(\sigma_i, s_{-i}).$$

One particular case is when μ_i satisfies **independence**, that is, when there exist mixed strategies $\sigma_{-i} = (\sigma_1, \dots, \sigma_{i-1}, \sigma_{i+1}, \dots, \sigma_I) \in \Sigma_{-i}$ such that

$$\mu_i(s_1, \dots, s_{i-1}, s_{i+1}, \dots, s_I) = \sigma_1(s_1) \cdots \sigma_{i-1}(s_{i-1}) \cdot \sigma_{i+1}(s_{i+1}) \cdots \sigma_I(s_I)$$

for all $(s_1, \dots, s_{i-1}, s_{i+1}, \dots, s_I) \in S_{-i}$. In this case, we define

$$u_j(\sigma_1, \dots, \sigma_{i-1}, \sigma_i, \sigma_{i+1}, \dots, \sigma_I) := u_j(\sigma_i, \mu_i),$$

which we may further abbreviate as $u_j(\sigma_i, \sigma_{-i})$. When μ_i can be written this way, we interpret this as saying that i believes his opponents choose their strategies independently.

1.2. Dominance and Best Response

A basic assumption of game theory is that each player maximizes his payoff, given his beliefs about what other players do. The outstanding questions are then:

- (a) What do we mean exactly by “beliefs”?
- (b) Are some beliefs more reasonable than others?

In some cases, a strategy of i does strictly better than another, regardless of the opponents’ choices, hence regardless of how we answer these questions.

Definition 1 (Strict Dominance). A strategy $s_i \in S_i$ is **strictly dominated** if there exists $\sigma_i \in \Sigma_i$ such that $u_i(\sigma_i, s_{-i}) > u_i(s_i, s_{-i})$ for all $s_{-i} \in S_{-i}$. In this case, we say σ_i strictly dominates s_i .

Note that, in some cases, a strategy is strictly dominated by a mixed strategy without being strictly dominated by a **pure strategy** (i.e. a degenerate mixed strategy).

Definition 2 (Weak Dominance). A strategy $s_i \in S_i$ is **weakly dominated** if there exists $\sigma_i \in \Sigma_i$ such that $u_i(\sigma_i, s_{-i}) \geq u_i(s_i, s_{-i})$ for all $s_{-i} \in S_{-i}$, with strict inequality for some s_{-i} .

Definition 3 (Best Response). Let $\mu_i \in \Delta(S_{-i})$. A strategy $s_i \in S_i$ is a **best response** to μ_i if $u_i(s_i, \mu_i) \geq u_i(s'_i, \mu_i)$ for all $s'_i \in S_i$. Given $\sigma_{-i} = (\sigma_1, \dots, \sigma_{i-1}, \sigma_{i+1}, \dots, \sigma_I) \in \Sigma_{-i}$, we say that $s_i \in S_i$ is a best response to σ_{-i} if it is a best response to $\sigma_1 \cdots \sigma_{i-1} \cdot \sigma_{i+1} \cdots \sigma_I$.

If our answer to question (a) above is that a player holds probabilistic beliefs over the opponents' choices, then the rationality requirement that every player i maximizes his expected payoff given his belief about S_{-i} can be stated as follows: every player i must choose a strategy s_i that is a best response to his beliefs μ_i .

Definition 4 (Never Best Response). A strategy $s_i \in S_i$ is a **never best response (NBR)** if there exists no $\mu_i \in \Delta(S_{-i})$ such that s_i is a best response to μ_i .

We are still left with question (b) unanswered. There are two basic ways to approach the problem of finding reasonable beliefs. They are both based on interactive reasoning, i.e. each player's reasoning about other players' choices *and* about other players' beliefs, and they are indeed equivalent in some cases.

1.3. Iterated Dominance and Rationalizability

1.3.1. Iterated Dominance

The basic idea is: a rational player i cannot choose a strictly dominated strategy, nor can he believe player $j \neq i$ would, nor can he believe player j can believe player $k \neq j$ would, and so on. At every step of the reasoning, further restrictions are imposed on what i 's choice can be. How do we formalize this idea?

Starting from the game Γ , define a sequence of games $\Gamma_D^1, \Gamma_D^2, \dots$ as follows: $\Gamma_D^1 = \Gamma$, and then recursively Γ_D^n is the game obtained from Γ_D^{n-1} by deleting all strictly dominated strategies of all players. Clearly, $\Gamma_D^{n+1} = \Gamma_D^n$ for n large enough. Let us denote by Γ_D this game we have converged to. The strategies in Γ that are also strategies in Γ_D are said to survive iterated elimination of strictly dominated strategies. If the strategy set of every player in game Γ_D is a singleton, then the game Γ is said to be dominance solvable.

Note that, in the process of elimination, we are not really assuming that player i has a well defined probabilistic belief μ_i over S_{-i} . All we are requiring is that he considers some of the opponents' strategies impossible. However, we do require that, when arguing that a certain strategy of j will not be chosen, this is because there is a *mixed* strategy of j that does better. The interpretation of the latter assumption may be, in principle, problematic.

1.3.2. Correlated Rationalizability

The basic idea is: a rational player i is Bayesian, i.e. he must have a well defined probabilistic belief μ_i over S_{-i} , thus he cannot choose a NBR strategy, nor can he believe player $j \neq i$ would, nor can he believe player j can believe player $k \neq j$ would, and so on. At every step of the reasoning, further restrictions are imposed on what μ_i can be, and thus on what i 's choice can be. How do we formalize this idea?

Starting from the game Γ , define a sequence of games $\Gamma_R^1, \Gamma_R^2, \dots$ as follows: $\Gamma_R^1 = \Gamma$, and then recursively Γ_R^n is the game obtained from Γ_R^{n-1} by deleting all NBR strategies of all players. Clearly, $\Gamma_R^{n+1} = \Gamma_R^n$ for n large enough. Let us denote by Γ_R this game we have converged to. The strategies in the game Γ_R are the correlated-rationalizable strategies of Γ .

Note that, in the process of elimination, we are not assuming players randomize. However, we are assuming that each player has a well defined probabilistic belief over the opponents' choices.

An equivalent procedure, indeed more in the spirit of the notion of rationalizability as it is intended in the literature, is to eliminate *beliefs*, and *hence* eliminate strategies, at each round. In the first round, for each player i we define

$$S_i^1 = \left\{ s_i \in S_i : s_i \text{ is a best response to some } \mu_i \in \Delta(S_{-i}) \right\}$$

Having defined S_i^n for every player i and some $n \geq 1$, we define recursively S_{-i}^n as an abbreviation for $S_1^n \times \dots \times S_{i-1}^n \times S_{i+1}^n \times \dots \times S_I^n$, and

$$S_i^{n+1} := \left\{ s_i \in S_i : s_i \text{ is a best response to some } \mu_i \in \Delta(S_{-i}) \text{ satisfying } \mu_i(S_{-i}^n) = 1 \right\}.$$

It is clear that $S_i^{n+1} \subseteq S_i^n \subseteq S_i$ for every player i and every $n \geq 1$. It is also rather obvious that S_i^n contains exactly those strategies of Γ that are also strategies in Γ_R^n . Those strategies of player i that belong to S_i^n for every $n \geq 1$ are i 's correlated-rationalizable strategies.

1.3.3. An Equivalence Result

Iterated elimination of strictly dominated strategies and correlated rationalizability are equivalent, as the following result shows.

Proposition 1. *In a finite normal form game Γ , a strategy is NBR if and only if it is strictly dominated. In particular, $\Gamma_D = \Gamma_R$.*

The proof of the proposition is a simple application of the following well known result, a version of Farkas's Lemma.

Lemma 1 (Farkas's Lemma). *Let A be a $m \times n$ matrix and let b be a $1 \times n$ vector. Either there exists a $1 \times m$ vector $x \geq 0$ such that $xA \leq b$, or there exists a $n \times 1$ vector $z \geq 0$ such that $Az \geq 0$ and $bz < 0$, but not both.¹*

Proof. See, for instance, the *Notes on Optimization* on my webpage. □

Proof of Proposition 1. Recall our definition of Γ_D^n and Γ_R^n . The second claim in the proposition easily follows from the first by induction on n , as Γ_D^n and Γ_R^n are, like Γ , finite games in normal form. To prove the first part, fix a strategy $\bar{s}_i \in S_i$, and let A be the matrix whose rows correspond to the elements of S_i , whose columns correspond to the elements of S_{-i} , and whose entry corresponding to row s_i and column s_{-i} is given by the difference $u_i(\bar{s}_i, s_{-i}) - u_i(s_i, s_{-i})$. Also, let b be a $1 \times n$ vector of -1 's. Then \bar{s}_i is NBR if and only if the second alternative in the lemma is false,² hence if and only if the first alternative in the lemma is true. The latter is easily verified to be equivalent to \bar{s}_i being strictly dominated; indeed, \bar{s}_i is strictly dominated if and only if there exists σ_i such that $\sigma_i A \ll 0$, which is equivalent to $xA \ll b$ for $x = a\sigma_i$ and $a > 0$ large enough. □

To see Proposition 1 in action, consider the following game:

	L	R
U	4, 0	3, 0
M	9, 0	0, 0
D	0, 0	9, 0

Player 1, the row player, will play the role of player i in the proposition, and strategy U will play the role of strategy \bar{s}_i . To construct the matrix A used in the proof of the proposition, we first consider only the matrix of player 1's payoffs:

The matrix A is then obtained as follows: for every $s_1 \in \{U, M, D\}$ and $s_2 \in \{L, R\}$, the entry

¹When dealing with two vectors v and w , the inequality $v \geq w$ means that every element of v is greater than or equal to the corresponding element of w . Similarly for the inequality $v \leq w$. The inequality $v \gg w$ instead means that every element of v is strictly greater than the corresponding element of w , and similarly for the inequality $v \ll w$.

²Make sure you know how to give a formal proof of the latter claim.

4	3
9	0
0	9

corresponding to s_1 and s_2 is the difference $u_1(U, s_2) - u_1(s_1, s_2)$. The matrix A is thus

$$A = \begin{array}{|c|c|} \hline 4-4 & 3-3 \\ \hline 4-9 & 3-0 \\ \hline 4-0 & 3-9 \\ \hline \end{array} = \begin{array}{|c|c|} \hline 0 & 0 \\ \hline -5 & 3 \\ \hline 4 & -6 \\ \hline \end{array}$$

Now U is strictly dominated by the mixed strategy that chooses M or D with equal probability; indeed, pre-multiplying A by the row vector $\sigma_1 = [0, 1/2, 1/2]$ we get

$$\sigma_1 A = [-0.5, -1.5].$$

Thus, defining $x = a\sigma_1$ for a large enough — any $a > 2$ will do — we get

$$xA = a\sigma_1 A = [-0.5a, -1.5a] \ll [-1, -1].$$

This means the first alternative in Farkas's Lemma is true, hence that the second is false. To verify that there is indeed no belief $\mu_1 \in \Delta(\{L, R\})$ against which U is a best response, note that this would require

$$4\mu_1(L) + 3\mu_1(R) \geq 9\mu_1(L) + 0\mu_1(R)$$

and

$$4\mu_1(L) + 3\mu_1(R) \geq 0\mu_1(L) + 9\mu_1(R),$$

which (seeing $\mu_1 = [\mu_1(L), \mu_1(R)]$ as a column vector) are together equivalent to

$$A\mu_1 \geq 0$$

and imply the contradiction that $\mu_1(R) \geq (5/3)\mu_1(L)$ and $\mu_1(R) \leq (4/6)\mu_1(L)$.

1.3.4. Independent Rationalizability

Definition 5 (Independent Never Best Response). A strategy $s_i \in S_i$ is an **independent never best response (INBR)** if there exists no $\sigma_{-i} \in \Sigma_{-i}$ such that s_i is a best response to σ_{-i} .

This is closer to the original version of rationalizability introduced by Bernheim and Pearce's papers (Econometrica 52(4), 1984). The basic idea is: a rational player i is Bayesian and believes his opponents' choices are uncorrelated; thus, he must have a well defined probabilistic belief μ_i over S_{-i} satisfying *independence*, and he cannot choose a INBR strategy, nor can he believe player $j \neq i$ would, nor can he believe player j can believe player $k \neq j$ would, and so on. At every step of the reasoning, further restrictions are imposed on what μ_i can be, and thus on what i 's choice can be. These restrictions are stronger than in the correlated case. Thus, iterated elimination of INBR strategies will deliver, in general, smaller sets of rationalizable strategies.

1.4. Examples

1.4.1. Independent vs. Correlated Rationalizability

Independent and correlated rationalizability obviously coincide in two-player games; in these games, a strategy is NBR if and only if it is INBR. In three-player games, the latter is no longer true. Consider the following three-player game, where both player 1 and player 2's payoffs are constant, and player 3 chooses A , B , C , or D . Strategy D is INBR. Indeed, take any $\sigma_1 \in \Sigma_1$ and $\sigma_2 \in \Sigma_2$.

	L	R		L	R		L	R		L	R
a	0, 0, 0	0, 0, 0	a	0, 0, 0	0, 0, 4	a	0, 0, 4	0, 0, 4	a	0, 0, 3	0, 0, 3
b	0, 0, 50	0, 0, 0	b	0, 0, 0	0, 0, 4	b	0, 0, 0	0, 0, 0	b	0, 0, 0	0, 0, 3
	A			B			C			D	

If $\sigma_2(R) > 3/4$, then B does better than D ; if $\sigma_1(a) > 3/4$, then C does better than D ; finally, if $\sigma_2(R) \leq 3/4$ and $\sigma_1(a) \leq 3/4$, then $\sigma_2(L) \geq 1/4$ and $\sigma_1(b) \geq 1/4$, hence $\sigma_{-3}(b, L) \geq 1/16$, hence A gives at least $50/16 > 3$ and therefore does strictly better than D . However, D is not NBR, because it is a best response to the belief μ_3 such that $\mu_3(a, L) = \mu_3(a, R) = \mu_3(b, R) = 1/3$.

Remark 2. Note that D is neither strictly nor weakly dominated by any mixed strategy of player 3. Let us prove this. Since D is not NBR, we know from Proposition 1 that it cannot be strictly dominated. To see that it is not weakly dominated, either, suppose by contradiction that it is weakly

dominated by some $\sigma_3 \in \Sigma_3$. Then we must have $u_3(a, L, \sigma_3) \geq u_3(a, L, D)$, i.e. $4\sigma_3(C) + 3\sigma_3(D) \geq 3$, which implies $\sigma_3(C) \geq (3/4)[1 - \sigma_3(D)]$. Moreover, we must have $u_3(b, R, \sigma_3) \geq u_3(b, R, D)$, i.e. $4\sigma_3(B) + 3\sigma_3(D) \geq 3$, which implies $\sigma_3(B) \geq (3/4)[1 - \sigma_3(D)]$. Thus, we must have $\sigma_3(C) + \sigma_3(B) \geq (6/4)[1 - \sigma_3(D)] \equiv (6/4)[\sigma_3(A) + \sigma_3(B) + \sigma_3(C)]$, which is only possible if $\sigma_3(A) = \sigma_3(B) = \sigma_3(C) = 0$, i.e. if $\sigma_3(D) = 1$, contradicting the initial supposition that σ_3 weakly dominates D (since that supposition obviously implies $\sigma_3(D) < 1$).

1.4.2. Rationalizability in the Cournot Game

The Cournot game is defined as follows: $N = \{1, 2\}$, $S_i = \mathbb{R}_+$, and $u_1(q_1, q_2) = q_1(P - q_1 - q_2)$, $u_2(q_1, q_2) = q_2(P - q_1 - q_2)$ for all $(q_1, q_2) \in S := S_1 \times S_2$, where $P > 0$ is a parameter.

This game is not finite, so checking whether a certain strategy of a firm is NBR requires looking at all probability distributions over the infinite set of possible quantities chosen by the other firm. This is, in principle, a hard task. However, a few simple observations will show that the problem is, in fact, rather easy. This is an example where rationalizability works wonderfully: it pins down a *unique* strategy for each player.

Let Γ denote the Cournot game. Let Γ_R^1 denote the Cournot game after deletion of all strategies that are NBR in Γ . What are such NBR strategies? Let μ_1 be a belief of firm 1 over $S_2 \equiv \mathbb{R}_+$, i.e. a probability measure over the set of player 2's strategies in Γ . A strategy in $S_1 \equiv \mathbb{R}_+$ is a best response to μ_1 for player 1 if and only if it solves the following:³

$$\max_{q_1 \in \mathbb{R}_+} \int_{\mathbb{R}_+} q_1(P - q_1 - q_2) d\mu_1(q_2).$$

By linearity of the integral, and using the fact that $\int d\mu = 1$, the latter problem is equivalent to

$$\max_{q_1 \in \mathbb{R}_+} -q_1^2 + q_1 P - q_1 \int_{\mathbb{R}_+} q_2 d\mu_1(q_2).$$

³Here and later on in this example, the integral sign denotes the Lebesgue integral — see the Notes on Probability on my webpage to find out more on this (and also to find a rigorous definition of probability measure).

In order to apply the Kuhn-Tucker theorem, we rewrite the latter as

$$\begin{aligned} \max_{q_1 \in \mathbb{R}} \quad & -q_1^2 + q_1 P - q_1 \int_{\mathbb{R}_+} q_2 d\mu_1(q_2) \\ \text{subject to} \quad & q_1 \geq 0 \end{aligned}$$

and we conclude, by the Kuhn-Tucker theorem,⁴ that a strategy q_1 is not NBR in Γ if and only if there exists $\mu_1 \in \Delta(\mathbb{R}_+)$ such that either $\int_{\mathbb{R}_+} q_2 d\mu_1(q_2) > P$ and $q_1 = 0$, or $\int_{\mathbb{R}_+} q_2 d\mu_1(q_2) \leq P$ and

$$q_1 = \frac{P - \int_{\mathbb{R}_+} q_2 d\mu_1(q_2)}{2}.$$

As μ_1 varies in $\Delta(\mathbb{R}_+)$, the expectation $\int_{\mathbb{R}_+} q_2 d\mu_1(q_2)$ varies between 0 and $+\infty$, hence the set of NBR strategies of player 1 in game Γ is $(P/2, +\infty)$. By symmetry, this is also the set of NBR strategies of player 2 in this game.

Thus, the strategy set of each player in game Γ^1 is $[0, P/2]$. A strategy q_1 is not NBR in this game if and only if it solves

$$\max_{q_1 \in [0, P/2]} \quad -q_1^2 + q_1 P - q_1 \int_{[0, P/2]} q_2 d\mu_1(q_2).$$

for some $\mu_1 \in \Delta([0, P/2])$. Again applying Kuhn-Tucker (this time with the two constraints $q_1 \geq 0$ and $-q_1 \geq -P/2$) we conclude that a strategy q_1 is not NBR in Γ_1 if and only if

$$q_1 = \frac{P - \int_{[0, P/2]} q_2 d\mu_1(q_2)}{2}$$

for some such μ_1 . As μ_1 varies in $\Delta([0, P/2])$, the expectation $\int_{[0, P/2]} q_2 d\mu_1(q_2)$ varies between 0 and $P/2$, hence the set of NBR strategies of player 1 in game Γ^1 is $[0, P/4)$. By symmetry, this is also the set of NBR strategies of player 2 in this game.

Now let Γ^2 be the game resulting from Γ^1 after deletion of NBR strategies. Thus, the strategy set of each player in game Γ^2 is $[P/4, P/2]$. A strategy q_1 is not NBR in this game if it solves

$$\max_{q_1 \in [P/4, P/2]} \quad -q_1^2 + q_1 P - q_1 \int_{[P/4, P/2]} q_2 d\mu_1(q_2).$$

⁴See my *Notes on Optimization* if you want to check that in this problem the Kuhn-Tucker conditions are indeed necessary and sufficient.

for some $\mu_1 \in \Delta([P/4, P/2])$, that is, again by Kuhn-Tucker, if and only if

$$q_1 = \frac{P - \int_{[P/4, P/2]} q_2 d\mu_1(q_2)}{2}$$

for some such μ_1 . As μ_1 varies in $\Delta([P/4, P/2])$, the expectation $\int_{[P/4, P/2]} q_2 d\mu_1(q_2)$ varies between $P/4$ and $P/2$, hence the set of NBR strategies of player 1 in game Γ^2 is $(3P/8, P/2]$. By symmetry, this is also the set of NBR strategies of player 2 in this game. In game Γ_3 , the strategy set of each player is thus $[P/4, 3P/8]$.

Etc. etc. etc.

Continuing in this fashion, we see that Γ^n is obtained from Γ^{n-1} by removal of the (upper or lower) half of both players' strategy sets in Γ^{n-1} . This process converges to the unique pair $q_1 = q_2 = P/3$.

2. Equilibrium in Normal Form Games

2.1. Nash Equilibrium

In general, correlated (or even independent) rationalizability does not give sharp predictions on how rational players should play and how much should they expect to get. On the other hand, all one needs in order to get the rationalizable outcomes is the assumption of rationality and common knowledge of rationality.

In some cases, a legitimate question is whether there is a strategy profile such that, once each player knows every other player is following the strategies specified in the profile, then he finds it optimal to do so himself.⁵

Definition 6. Let $\langle N, (S_i, u_i)_{i \in N} \rangle$ be a normal form game. A **Nash equilibrium** of this game is a strategy profile $\bar{s} \in S$ such that $u_i(\bar{s}) \geq u_i(s_i, \bar{s}_{-i})$ for every player $i \in N$ and every $s_i \in S_i$.

It is quite easy to see that a Nash equilibrium can only prescribe (independent) rationalizable strategies. If \bar{s} is a Nash equilibrium, then \bar{s}_i survives the first round of elimination of INBR strategies (being a best response to the belief that his opponents will play \bar{s}_{-i} with probability one); since this is true for every player, it remains true at all subsequent rounds as well. Moreover, since a NBR strategy is also an INBR strategy, the strategies in a Nash equilibrium are also correlated-rationalizable; equivalently, they survive iterated elimination of strictly dominated strategies.

2.1.1. Existence of Nash Equilibria

A Nash equilibrium may not exist, e.g. in the game *matching pennies*:

	<i>H</i>	<i>T</i>
<i>H</i>	1, -1	-1, 1
<i>T</i>	-1, 1	1, -1

But if we allow mixed strategies, then it always exists (in finite games).

Definition 7. The **mixed extension** of a finite game $\langle N, (S_i, u_i)_{i \in N} \rangle$ is the game where the set of players is again N , the set of strategies of player i is $\Sigma_i = \Delta(S_i)$, and the payoff to player i from

⁵Here and in what follows, given any $s \in S$ and $i \in N$, we write s_i and s_{-i} to denote the elements of S_i and S_{-i} corresponding to s . Thus, by this definition, $s = (s_i, s_{-i})$ for every $s \in S$ and every $i \in N$.

the strategy profile $\sigma \in \Sigma$ is $u_i(\sigma)$.

The following is the result originally proved by Nash in 1951.

Theorem 1. *The mixed extension of any finite game has a Nash equilibrium.*

The proof of the theorem uses the following result.

Lemma 2 (Brouwer's Fixed Point Theorem). *Let C be a compact convex subset of a Euclidean space. Let $f : C \rightarrow C$ be a continuous function. Then f has a fixed point, i.e. there is a point $x \in C$ for which $f(x) = x$.*

Proof. See, for instance, the book *Fixed Point Theorems with Applications to Economics and Game Theory* by Kim Border. □

Proof of Theorem 1. Let $\langle N, (S_i, u_i)_{i \in N} \rangle$ be a finite normal form game. Suppose without loss of generality that $N = \{1, \dots, I\}$, where $I = |N|$, and $S_i = \{1, \dots, m_i\}$, where $m_i = |S_i|$, for every $i = 1, \dots, I$. Then, for every player i and every $\sigma_i \in \Sigma_i$, write σ_i^n for the probability assigned by σ_i to the n th strategy of player i , where $n = 1, \dots, m_i$. Obviously, Σ_i is compact and convex for each player i , hence so is Σ . For every player i and every $n = 1, \dots, m_i$ define a function $g_i^n : \Sigma \rightarrow \mathbb{R}$ as follows:

$$g_i^n(\sigma) = \max \left\{ 0, u_i(n, \sigma_{-i}) - u_i(\sigma_i, \sigma_{-i}) \right\} \quad \forall \sigma = (\sigma_i, \sigma_{-i}) \in \Sigma.$$

In other words, $g_i^n(\sigma)$ is the gain accruing to player i as a consequence of moving from σ_i to his n th pure strategy, given that other players are choosing σ_{-i} , or it is zero if the latter is actually a loss. Define the function $f : \Sigma \rightarrow \Sigma$ as follows: for every $\sigma = (\sigma_i, \sigma_{-i}) \in \Sigma$, the profile of mixed strategies $f(\sigma) = (f_i(\sigma))_{i=1}^I$ is such that, for every player i , the probability assigned by $f_i(\sigma)$ to the n th strategy of player i is

$$f_i^n(\sigma) = \frac{\sigma_i^n + g_i^n(\sigma)}{1 + \sum_{m=1}^{m_i} g_i^m(\sigma)}.$$

Note that, since $f_i^n(\sigma) \geq 0$ for every n , and moreover $f_i^1(\sigma) + \dots + f_i^{m_i}(\sigma) = 1$, we indeed have $f_i(\sigma) \in \Sigma_i$ and hence $f(\sigma) \in \Sigma$. Moreover, the function f is continuous; since g is the maximum of polynomial functions, it is continuous, so f is a ratio of continuous functions and

thus also continuous. By Brouwer's fixed point theorem, there exists $\bar{\sigma} = (\bar{\sigma}_i, \bar{\sigma}_{-i}) \in \Sigma$ such that $f(\bar{\sigma}) = \bar{\sigma}$.

We now show that $\bar{\sigma}$ is a Nash equilibrium. For every player i and every $n = 1, \dots, m_i$ we have $f_i^n(\bar{\sigma}) = \bar{\sigma}_i^n$, that is,

$$\frac{\bar{\sigma}_i^n + g_i^n(\bar{\sigma})}{1 + \sum_{m=1}^{m_i} g_i^m(\bar{\sigma})} = \bar{\sigma}_i^n$$

and hence

$$g_i^n(\bar{\sigma}) = \bar{\sigma}_i^n \sum_{m=1}^{m_i} g_i^m(\bar{\sigma}).$$

Moreover, for every player i there must exist n such that $\bar{\sigma}_i^n > 0$ and $g_i^n(\bar{\sigma}) = 0$. If this were not true, then, for all m for which $\bar{\sigma}_i^m > 0$, we would have $g_i^m(\bar{\sigma}) > 0$, i.e. $u_i(m, \bar{\sigma}_{-i}) > u_i(\bar{\sigma})$, hence

$$\sum_{m=1}^{m_i} \bar{\sigma}_i^m u_i(m, \bar{\sigma}_{-i}) > u_i(\bar{\sigma})$$

which is a contradiction. Thus,

$$\sum_{m=1}^{m_i} g_i^m(\bar{\sigma}) = 0 \quad \forall i,$$

and therefore, since $g_i^m(\bar{\sigma}) \geq 0$ for all i and all m , also

$$g_i^m(\bar{\sigma}) = 0 \quad \forall i \forall m.$$

This means that $u_i(\bar{\sigma}) \geq u_i(m, \bar{\sigma}_{-i})$ for all i and all $m \in S_i$, hence $u_i(\bar{\sigma}) \geq u_i(\sigma_i, \bar{\sigma}_{-i})$ for all i and all $\sigma_i \in \Sigma_i$. \square

2.2. Strictly Competitive Games

In general, the notion of Nash equilibrium does not give sharp predictions about the behavior of rational players, though certainly sharper than those implied by (either form of) rationalizability.

There is a class of games, however, where the implications of the equilibrium hypothesis are indeed rather precise. This is the class of games whose analysis (pioneered by von Neumann and Morgenstern) gave birth to game theory.

Definition 8. A game $\langle N, (S_i, u_i)_{i \in N} \rangle$ is **strictly competitive** if $N = \{1, 2\}$ and if, moreover, for every $s \in S$ one has $u_1(s) > 0$ if and only if $u_2(s) < 0$.

In the remainder of this section, we will just assume that a strictly competitive game is in fact a **zero-sum** game, that is, we will assume that $u_1(s) + u_2(s) = 0$ for every $s \in S$. The following is one of the most important and classical results about this class of games — it was proved by von Neumann in 1928.⁶

Theorem 2 (The Minimax Theorem). *Let $\langle \{1, 2\}, \{\Sigma_1, \Sigma_2\}, \{u_1, u_2\} \rangle$ be the mixed extension of a finite zero-sum game. A profile $(\bar{\sigma}_1, \bar{\sigma}_2) \in \Sigma_1 \times \Sigma_2$ is a Nash equilibrium if and only if*

$$\bar{\sigma}_1 \in \arg \max_{\sigma_1 \in \Sigma_1} \min_{\sigma_2 \in \Sigma_2} u_1(\sigma_1, \sigma_2) \quad \text{and} \quad \bar{\sigma}_2 \in \arg \min_{\sigma_2 \in \Sigma_2} \max_{\sigma_1 \in \Sigma_1} u_1(\sigma_1, \sigma_2). \quad (1)$$

Moreover, if $(\bar{\sigma}_1, \bar{\sigma}_2)$ is a Nash equilibrium, then

$$u_1(\bar{\sigma}_1, \bar{\sigma}_2) = \min_{\sigma_2 \in \Sigma_2} \max_{\sigma_1 \in \Sigma_1} u_1(\sigma_1, \sigma_2) = \max_{\sigma_1 \in \Sigma_1} \min_{\sigma_2 \in \Sigma_2} u_1(\sigma_1, \sigma_2). \quad (2)$$

Remark 3. Note that, since an equilibrium of the mixed extension of a finite game always exists, the theorem does imply that $v := \min_{\sigma_2 \in \Sigma_2} \max_{\sigma_1 \in \Sigma_1} u_1(\sigma_1, \sigma_2) = \max_{\sigma_1 \in \Sigma_1} \min_{\sigma_2 \in \Sigma_2} u_1(\sigma_1, \sigma_2)$. This number v is called the **value** of the game.

Proof of Theorem 2. Pick any $(\bar{\sigma}_1, \bar{\sigma}_2) \in \Sigma_1 \times \Sigma_2$. Obviously, we have

$$\begin{aligned} u_1(\sigma_1, \bar{\sigma}_2) &\geq \min_{\sigma_2 \in \Sigma_2} u_1(\sigma_1, \sigma_2) & \forall \sigma_1 \in \Sigma_1, \\ u_1(\bar{\sigma}_1, \sigma_2) &\leq \max_{\sigma_1 \in \Sigma_1} u_1(\sigma_1, \sigma_2) & \forall \sigma_2 \in \Sigma_2. \end{aligned}$$

Suppose $(\bar{\sigma}_1, \bar{\sigma}_2)$ is an equilibrium. Then, using the inequalities above,

$$\begin{aligned} u_1(\bar{\sigma}_1, \bar{\sigma}_2) &= \max_{\sigma_1 \in \Sigma_1} u_1(\sigma_1, \bar{\sigma}_2) \geq \max_{\sigma_1 \in \Sigma_1} \min_{\sigma_2 \in \Sigma_2} u_1(\sigma_1, \sigma_2), \\ u_1(\bar{\sigma}_1, \bar{\sigma}_2) &= \min_{\sigma_2 \in \Sigma_2} u_1(\bar{\sigma}_1, \sigma_2) \leq \min_{\sigma_2 \in \Sigma_2} \max_{\sigma_1 \in \Sigma_1} u_1(\sigma_1, \sigma_2), \end{aligned}$$

⁶The statement of Theorem 2 refers to Nash equilibria, and our proof (taken from Myerson's textbook) uses Nash's theorem, that is, Theorem 1 in these notes, which only appeared in 1951. The theorem originally proved by von Neumann, of course, does not refer to Nash equilibria, nor does its proof refer to Nash's theorem; indeed, von Neumann takes the equalities in (2) to mean, by themselves, that $(\bar{\sigma}_1, \bar{\sigma}_2)$ is solution to the game — see Remark 4 below.

and moreover

$$u_1(\bar{\sigma}_1, \bar{\sigma}_2) = \min_{\sigma_2 \in \Sigma_2} u_1(\bar{\sigma}_1, \sigma_2) \leq \max_{\sigma_1 \in \Sigma_1} \min_{\sigma_2 \in \Sigma_2} u_1(\sigma_1, \sigma_2),$$

$$u_1(\bar{\sigma}_1, \bar{\sigma}_2) = \max_{\sigma_1 \in \Sigma_1} u_1(\sigma_1, \bar{\sigma}_2) \geq \min_{\sigma_2 \in \Sigma_2} \max_{\sigma_1 \in \Sigma_1} u_1(\sigma_1, \sigma_2).$$

The latter four inequalities clearly imply (2). Moreover, they imply

$$\min_{\sigma_2 \in \Sigma_2} u_1(\bar{\sigma}_1, \sigma_2) = \max_{\sigma_1 \in \Sigma_1} \min_{\sigma_2 \in \Sigma_2} u_1(\sigma_1, \sigma_2) \quad \text{and} \quad \max_{\sigma_1 \in \Sigma_1} u_1(\sigma_1, \bar{\sigma}_2) = \min_{\sigma_2 \in \Sigma_2} \max_{\sigma_1 \in \Sigma_1} u_1(\sigma_1, \sigma_2)$$

which is the same as (1). Conversely, suppose (1) holds. By Theorem 1, an equilibrium must exist.

Thus, by the first part of this proof, the second equality in (2) holds. Thus,

$$\begin{aligned} u_1(\bar{\sigma}_1, \bar{\sigma}_2) &\geq \max_{\sigma_1 \in \Sigma_1} u_1(\sigma_1, \bar{\sigma}_2) \\ &= \min_{\sigma_2 \in \Sigma_2} u_1(\bar{\sigma}_1, \sigma_2) \quad (\text{by (1) and the second equality in (2)}) \\ &\geq u_1(\bar{\sigma}_1, \bar{\sigma}_2), \end{aligned}$$

which shows that $(\bar{\sigma}_1, \bar{\sigma}_2)$ is a Nash equilibrium. □

Remark 4. The Minimax Theorem states that all Nash equilibria of a finite zero-sum game give player 1 (and hence also player 2) the same payoff, namely, the value of the game. Note that the value is the payoff that player 1 can *guarantee* himself (max min), but also the payoff that player 2 can force player 1 down to (min max). Thus, in a sense, a rational player in a zero-sum game does not need to guess what his opponent will do, in order to play an optimal strategy. This is why, even if the notion of Nash equilibrium did not exist at the time, the conclusion of von Neumann was that the equalities in (2) by themselves constitute an acceptable characterization of a solution to a zero-sum game.

Another implication of the Minimax Theorem is that the equilibria of a finite zero-sum game are *interchangeable*. If $(\bar{\sigma}_1, \bar{\sigma}_2)$ and $(\tilde{\sigma}_1, \tilde{\sigma}_2)$ are both equilibria of the game, then $(\bar{\sigma}_1, \tilde{\sigma}_2)$ and $(\tilde{\sigma}_1, \bar{\sigma}_2)$ are also equilibria of the game (why?) . Once again, this shows that, in a zero-sum game, ★
a rational player 1 (resp. player 2) does not need to guess what exact strategy player 2 (resp. player

1) will choose; all he needs to know is that player 2 (resp. player 1) will choose an optimal strategy, i.e. a strategy $\bar{\sigma}_2$ (resp. $\bar{\sigma}_1$) that satisfies (1).

2.3. Normal Form Refinements

Are there solution concepts that deliver more accurate predictions than Nash equilibrium? The answer is yes, and we will briefly review two of them. Both are based on the idea that, among all Nash equilibria of (the mixed extension of) a game, only those that are robust to small perturbations in the players' choices are reasonable.

2.3.1. Trembling-Hand Perfect Equilibrium

The first refinement we review is due to Selten (International Journal of Game Theory 4, 1975). Let $\langle N, (\Sigma_i, u_i)_{i \in N} \rangle$ be the mixed extension of a finite normal form game. A **completely mixed** strategy profile of this game is a strategy profile $\sigma \in \Sigma$ such that $\sigma_i(s_i) > 0$ for every player $i \in N$ and every $s_i \in S_i$.

Definition 9. Let $\langle N, (\Sigma_i, u_i)_{i \in N} \rangle$ be the mixed extension of a finite normal form game. A mixed strategy profile $\bar{\sigma} \in \Sigma$ is a **trembling-hand perfect equilibrium** (of the original finite normal form game) if there exists a sequence $\sigma^1, \sigma^2, \dots$ of completely mixed strategy profiles converging to $\bar{\sigma}$ and such that, for every $n = 1, 2, \dots$ and every player $i \in N$, the strategy $\bar{\sigma}_i$ is a best response to σ_{-i}^n .

It is quite easy to see that a trembling-hand perfect equilibrium is also a mixed strategy Nash equilibrium. (Would you be able to prove this formally?) Thus, trembling-hand perfection is ★ indeed a refinement of Nash.

2.3.2. Proper Equilibrium

An even more demanding solution concept has been introduced by Myerson (International Journal of Game Theory 7, 1978).

Definition 10. Let $\langle N, (\Sigma_i, u_i)_{i \in N} \rangle$ be the mixed extension of a finite normal form game. A mixed strategy profile $\bar{\sigma} \in \Sigma$ is a **proper equilibrium** (of the original finite normal form game) if there exist a sequence of positive numbers $\varepsilon_1, \varepsilon_2, \dots$ and a sequence of completely mixed strategy

profiles $\sigma^1, \sigma^2, \dots$ converging to $\bar{\sigma}$ such that $\lim_{n \rightarrow \infty} \varepsilon_n = 0$ and, for all $n = 1, 2, \dots$, all $i \in N$ and all $s_i, s'_i \in S_i$,

$$u_i(s'_i, \sigma_{-i}^n) < u_i(s_i, \sigma_{-i}^n) \quad \Rightarrow \quad \sigma_i^n(s'_i) \leq \varepsilon_n \sigma_i^n(s_i).$$

Similarly to a trembling-hand perfect equilibrium, a proper equilibrium must be robust to some small perturbation in the players' behavior. But, in addition, this perturbation must be “rational”, in the sense that strategies that do not do well along the sequence must be played with much smaller probability than those which do better, i.e. players are more likely to tremble on strategies that are not too bad for them.

One can show that a proper equilibrium must be trembling-hand perfect. (Giving a formal proof is a little hard, but definitely doable.) Thus, the following result establishes existence of trembling-hand perfect equilibria as well. ★★★

Proposition 2. *Every finite normal form game has a proper equilibrium.*

The proof uses the following result.

Lemma 3 (Kakutani's Fixed Point Theorem). *Let C be a compact, convex subset of a Euclidean space. Let $\Phi : C \rightrightarrows C$ be a correspondence having closed graph and such that $\Phi(x)$ is nonempty and convex for every $x \in C$. Then Φ has a fixed point, i.e. there exists $x \in C$ for which $x \in \Phi(x)$.*

Proof. See, for instance, Corollary 15.3 in Border's book. □

Proof of Proposition 2. Choose $0 < \varepsilon < 1$ and, for every $i \in N$, define

$$R_i^\varepsilon = \{\sigma_i \in \Delta(S_i) : \sigma_i(s_i) \geq \varepsilon^{|S_i|} / |S_i| \quad \forall s_i \in S_i\}.$$

Clearly, R_i^ε is compact and convex for every $i \in N$, hence so is $R^\varepsilon := \times_{i \in N} R_i^\varepsilon$. Now define $\Phi^\varepsilon : R^\varepsilon \rightrightarrows R^\varepsilon$ as follows: for every $\sigma \in R^\varepsilon$, $\Phi^\varepsilon(\sigma) = (\Phi_i^\varepsilon(\sigma))_{i \in N}$, where

$$\Phi_i^\varepsilon(\sigma) = \{\sigma'_i \in R_i^\varepsilon : \text{for all } s_i, s'_i \in S_i, \quad u_i(s'_i, \sigma_{-i}) < u_i(s_i, \sigma_{-i}) \Rightarrow \sigma'_i(s'_i) \leq \varepsilon \sigma'_i(s_i)\}$$

for every $i \in N$. This correspondence has a closed graph and is convex valued. (Can you show ★★

why this is true?) Moreover, it is nonempty valued, i.e. $\Phi^\varepsilon(\sigma) \neq \emptyset$ for all $\sigma \in R^\varepsilon$. To show this, take any $\sigma \in R^\varepsilon$ and define, for every $i \in N$ and every $s_i \in S_i$,

$$k(s_i, \sigma) = \left| \{s'_i \in S_i : u_i(s_i, \sigma_{-i}) < u_i(s'_i, \sigma_{-i})\} \right|$$

and

$$\tilde{\sigma}_i(s_i) = \frac{\varepsilon^{k(s_i, \sigma)}}{\sum_{s'_i \in S_i} \varepsilon^{k(s'_i, \sigma)}}.$$

Clearly, if $s_i, s'_i \in S_i$ and $u_i(s_i, \sigma_{-i}) < u_i(s'_i, \sigma_{-i})$, then $\tilde{\sigma}_i(s_i) \leq \varepsilon \tilde{\sigma}_i(s'_i)$. Thus, $\tilde{\sigma} \in \Phi^\varepsilon(\sigma)$.

We conclude by Kakutani's theorem that Φ^ε has a fixed point. (Why does this imply existence ★★ of a proper equilibrium? Hint: Every bounded sequence in a Euclidean space has a convergent subsequence; moreover, $\times_{i \in N} \Delta(S_i)$ is closed ...) \square

2.4. Examples

2.4.1. Trembling-Hand Perfection and Weakly Dominated Strategies

A trembling-hand perfect equilibrium cannot prescribe that a weakly dominated strategy be played with positive probability. Consider the following example:

	<i>L</i>	<i>R</i>
<i>U</i>	2, 0	1, 1
<i>D</i>	1, 1	1, 0

This game has one pure strategy Nash equilibrium, (U, R) , and a *continuum* of mixed strategy equilibria (player 2 chooses R , player 1 chooses U with probability p , where $1/2 \leq p < 1$). However, none of these mixed equilibria is trembling-hand perfect. Since a proper equilibrium exists, we conclude that the unique proper equilibrium (and also the unique trembling-hand perfect equilibrium) is (U, R) .

To see why the mixed equilibria are not trembling-hand perfect, let $\sigma^n = (\sigma_1^n, \sigma_2^n)$ be any sequence of completely mixed profiles, and let $\bar{\sigma} = (\bar{\sigma}_1, \bar{\sigma}_2)$ be such that $1/2 \leq \bar{\sigma}_1(U) < 1$ and

$\bar{\sigma}_2(R) = 1$. Since $\bar{\sigma}_1(D) > 0$ and $\sigma_2^n(L) > 0$ for every $n = 1, 2, \dots$, we have

$$\begin{aligned} u_1(\bar{\sigma}_1, \sigma_2^n) &= \bar{\sigma}_1(U)[2\sigma_2^n(L) + \sigma_2^n(R)] + \bar{\sigma}_1(D) \\ &< 2\sigma_2^n(L) + \sigma_2^n(R) \\ &= u_1(U, \sigma_2^n). \end{aligned}$$

Thus, $\bar{\sigma}_1$ is not a best response to σ_2^n for any n . Therefore, $\bar{\sigma}$ is not trembling-hand perfect.

To verify that (U, R) is a trembling-hand perfect equilibrium, consider the sequence of completely mixed strategy profiles $\sigma^1, \sigma^2, \dots$ such that, for every $n = 1, 2, \dots$,

$$\sigma_1^n(U) = 1 - (1/2)^n \quad \text{and} \quad \sigma_2^n(R) = 1 - (1/2)^n$$

Clearly, σ^n converges to (U, R) . Moreover, U and R are best responses to σ_2^n and σ_1^n , respectively, for every n . Thus, (U, R) is trembling-hand perfect.

2.4.2. Trembling-Hand Perfection and Properness

The following example shows that a trembling-hand perfect equilibrium need not be proper.

	L	M	R
U	1, 1	0, 0	-9, -9
M	0, 0	0, 0	-7, -7
D	-9, -9	-7, -7	-7, -7

The strategy profile (M, M) is trembling-hand perfect. (You should have no problem in proving the latter.) However, it is not proper. Indeed, take any sequence of positive numbers $\varepsilon_1, \varepsilon_2, \dots$ converging to 0 and any sequence of strategy profiles $\sigma^1, \sigma^2, \dots$ converging to (M, M) such that ★★

$$u_i(s'_i, \sigma_{-i}^n) < u_i(s_i, \sigma_{-i}^n) \quad \Rightarrow \quad \sigma_i^n(s'_i) \leq \varepsilon_n \sigma_i^n(s_i)$$

for all $n = 1, 2, \dots$, all $i = 1, 2$, and all $s_i, s'_i \in S_i$. Since σ^n converges to (M, M) , we have $u_1(D, \sigma_2^n) < u_1(U, \sigma_2^n)$ and $u_2(R, \sigma_1^n) < u_2(L, \sigma_1^n)$ for all n large enough. (Why?) Thus, ★

$$\sigma_1^n(D) \leq \varepsilon_n \sigma_1^n(U) \quad \text{and} \quad \sigma_2^n(R) \leq \varepsilon_n \sigma_2^n(L) \quad (3)$$

for all n large enough. Now, against σ_2^n , the strategies U and M give respectively $u_1(U, \sigma_2^n) = \sigma_2^n(L) - 9\sigma_2^n(R)$ and $u_1(M, \sigma_2^n) = -7\sigma_2^n(R)$. By (3), then,

$$\begin{aligned} u_1(U, \sigma_2^n) - u_1(M, \sigma_2^n) &= \sigma_2^n(L) - 2\sigma_2^n(R) \\ &\geq \sigma_2^n(L) - 2\varepsilon_n \sigma_2^n(L) \\ &= (1 - 2\varepsilon_n) \sigma_2^n(L). \end{aligned}$$

Since σ^n is completely mixed, $\sigma_2^n(L) > 0$ for every n . Moreover, since ε_n converges to zero, $1 - 2\varepsilon_n > 0$ for all n large enough. We conclude that $u_1(U, \sigma_2^n) - u_1(M, \sigma_2^n) > 0$ for all n large enough, hence that $\sigma_1^n(M) \leq \varepsilon_n \sigma_1^n(U)$ for all n large enough, hence (why?) that $\sigma_1^n(M)$ converges to zero. This contradicts the initial supposition that σ^n converges to (M, M) . ★

3. Extensive Form Games: Basics

A game in normal form specifies what the players' strategies are, and what each player's payoff is, as a function of the possible profiles of strategies chosen by all players. A game in **extensive form** is a more detailed description of a strategic situation, and allows to model *dynamic* choice.

3.1. Formal Definition

An extensive form specifies the physical order of play (who moves when, and what choices are available), the information available to a player when making a choice, and the payoff to each player as a function of the moves selected by all players.

Definition 11. A **finite extensive form game** is a list $\langle N, X, p, A, \alpha, H, \iota, \sigma_0, (u_i)_{i \in N} \rangle$ comprising the following objects:

- a finite set of **players** $N = \{0, 1, \dots, I\}$; player 0 is called **Nature**;
- a finite set of **nodes** X and a function $p : X \rightarrow X \cup \emptyset$, indicating **immediate precedence**; we define $p_1 = p$ and recursively $p_{n+1} = p \circ p_n$ for all $n \geq 1$; if $x, x' \in X$ and $x = p^n(x')$ for some $n \geq 1$, then we say that x is a **predecessor** of x' and that x' is a **successor** of x ; we require p to satisfy the following properties:

- there exists a unique **initial node**, that is, a unique $x_0 \in X$ such that $p(x_0) = \emptyset$;
- for every $x \in X$, the set of predecessors of x and the set of successors of x are disjoint; in other words, there exist no $n \geq 1$ and $x' \in X$ for which $x' = p_n(x')$;

the nodes in the set $Z := \{x \in X : p^{-1}(x) = \emptyset\}$ are called **terminal nodes**;

- a finite set of possible **actions** A and a function $\alpha : X \setminus \{x_0\} \rightarrow A$ specifying, for each non-initial node, the action that leads to that node; it is assumed that $\alpha(x') \neq \alpha(x'')$ for every $x \in X$ and every two distinct $x', x'' \in p^{-1}(x)$; the set of actions available at any $x \in X$ is defined as $A(x) := \{a \in A : \exists x' \in p^{-1}(x) \text{ s.t. } \alpha(x') = a\}$;
- a partition H of $X \setminus Z$ into **information sets**, and a function $\iota : H \rightarrow N$ specifying who moves at each information set, such that $A(x) = A(x')$ for every $h \in H$ and every $x, x' \in h$; this means that the actions available at node x in information set h are the same as those

available at any other node in h , hence we can write $A(h)$ to denote the set of such actions; we define $H_i := \iota^{-1}(i)$ for each player $i = 0, 1, \dots, I$; for every $x \in X \setminus Z$ we write $H(x)$ to denote the member of H containing x ; we assume that, if Nature is an active player, then it moves only at the beginning of the game — i.e. either $H_0 = \emptyset$ or $H_0 = H(x_0)$ — and, in the latter case, it does so by choosing randomly an element of $A(x_0)$ according to the probability distribution $\sigma_0 \in \Delta(A(x_0))$;

- for each player $i = 1, \dots, I$, a **payoff** function $u_i : Z \rightarrow \mathbb{R}$.

Here is an example of an extensive form game with $I = 2$ and no moves by Nature. Player 1 chooses L or R and then player 2, without knowing what the choice of player 1 has been, chooses A or B . The fact that 2 does not know whether 1 chose L or R is reflected by the fact that the two successors of the initial node are in the same information set, denoted by a dotted line. Note that at the two nodes player 2 has the same actions available.

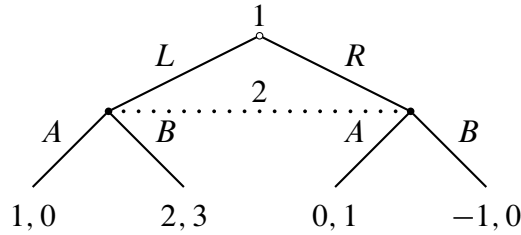


Figure 1: An extensive form game.

Here is another example: Player 1 chooses L or R and then player 2, having observed the choice of player 1, chooses A or B ; if he chooses B , then it is 1's turn again, and he has to choose between ℓ and r .

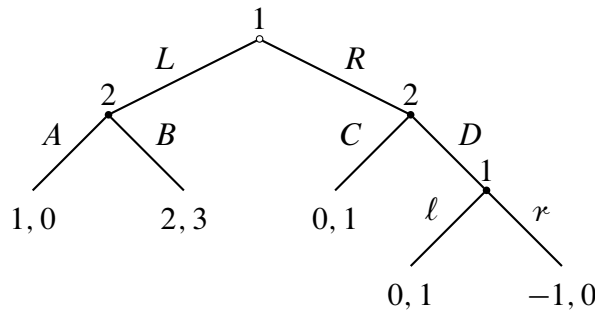


Figure 2: An extensive form game where player 1 moves twice.

A common property that we will assume throughout is **perfect recall**. Roughly, this means that no player ever forgets what he once knew. Formally, the property can be stated as follows: for every

player i and every three nonterminal nodes x, y, w satisfying $H(x) \subseteq H_i$, $H(y) = H(w) \subseteq H_i$, and $x = p_n(y)$ for some $n \geq 1$, there must exist a nonterminal node x' such that $H(x') = H(x)$, $x' = p_m(w)$ for some $m \geq 1$, and $\alpha(p_{n-1}(y)) = \alpha(p_{m-1}(w))$.

3.1.1. Normal Form Representation of an Extensive Form Game

A **strategy** for player $i = 1, \dots, I$ is a function $s_i : H_i \rightarrow A$ such that $s_i(h) \in A(h)$ for all $h \in H_i$. Thus, the set of all strategies of player i is the cartesian product

$$S_i = \prod_{h \in H_i} A(h).$$

If Nature is not active, a pure strategy profile induces a unique terminal node. To see this, suppose $H_0 = \emptyset$ and take any $s \in S$. Since actions leading to distinct immediate successors of a nonterminal node are labeled differently by α , there exists a unique $x_1 \in p^{-1}(x_0)$ such that $\alpha(x_1) = s_{\iota(H(x_0))}(H(x_0))$, and then a unique $x_2 \in p^{-1}(x_1)$ such that $\alpha(x_2) = s_{\iota(H(x_1))}(H(x_1))$, and so on. Thus we have a sequence x_0, x_1, \dots, x_n such that $x_k = p(x_{k+1})$ for every $k = 0, \dots, n-1$, and $x_n \in Z$. Each of the nodes x_0, \dots, x_n is said to be **reached** under s . The terminal node x_n is the **outcome induced** by s and is denoted by $\phi(s)$. The function $\phi : S \rightarrow Z$ thus constructed is the **outcome function** of the game, and the associated game in normal form is then $\langle N, (S_i, U_i) \rangle$, where $U_i : S \rightarrow \mathbb{R}$ is defined as $U_i(s) := u_i(\phi(s))$ for all $s \in S$.

Remark 5. If Nature is active, a profile of pure strategies s induces a *probability distribution* over the terminal nodes. Every move $a \in A(x_0)$ by Nature results in a node $x_1^a \in p^{-1}(x_0)$, and given this x_1^a there exists a unique $x_2^a \in p^{-1}(x_1^a)$ such that $\alpha(x_2^a) = s_{\iota(H(x_1^a))}(H(x_1^a))$, then a unique $x_3^a \in p^{-1}(x_2^a)$ such that $\alpha(x_3^a) = s_{\iota(H(x_2^a))}(H(x_2^a))$, and so on. Thus each $a \in A(x_0)$, together with s , uniquely identifies a sequence $x_0, x_1^a, \dots, x_{n(a)}^a$ such that $x_0 = p(x_1^a)$ and $x_k^a = p(x_{k+1}^a)$ for every $k = 1, \dots, n(a)-1$, and $x_{n(a)}^a \in Z$. Each of the nodes in the set

$$\{x_0\} \cup \{x \in X : x = x_k^a \text{ for some } a \in A(x_0) \text{ such that } \sigma_0(a) > 0 \text{ and some } 1 \leq k \leq n(a)\}$$

is said to be reached under s . The random outcome induced by s is then the probability distribution $\phi(s) \in \Delta(Z)$ such that $\phi(s)(x_{n(a)}^a) = \sigma_0(a)$ for every $a \in A(x_0)$. The function $\phi : S \rightarrow \Delta(Z)$ thus constructed is the outcome function of the game, and the associated game in normal form is

$\langle N, (S_i, U_i) \rangle$, where $U_i : S \rightarrow \mathbb{R}$ is defined as $U_i(s) := \sum_{z \in Z} \phi(s)(z) u_i(z)$ for all $s \in S$.

3.2. Mixed and Behavior Strategies: Kuhn's Theorem

A **mixed strategy** for player $i \in \{1, \dots, I\}$ is a probability distribution $\sigma_i \in \Delta(S_i)$. Note that, if Nature is active, then the elements of $A(x_0)$ can be seen as the “pure strategies” of Nature, and σ_0 can be seen as an exogenously given “mixed strategy” of Nature.

If Nature is not active, the random outcome induced by a mixed strategy profile $\sigma \in \Sigma$ is simply the probability distribution $\hat{\phi}(\sigma) \in \Delta(Z)$ which assigns to terminal node z the probability

$$\hat{\phi}(\sigma)(z) = \sum_{s \in \phi^{-1}(z)} \sigma(s), \quad (4)$$

where ϕ is the function constructed in 3.1.1 above, and $\sigma(s)$ is the probability that s is chosen when the players randomize according to σ . Given any $p \in [0, 1]$, we say a node $x \in X$ is **reached with probability p under σ** if, denoting by $S(h)$ the set of strategy profiles under which x is reached, we have $\sum_{s \in S(h)} \sigma(s) = p$. (If $p = 0$, we may just say that x is not reached under σ .) The expected payoff of player i under σ , which we denote by $U_i(\sigma)$ with slight abuse of notation, is

$$U_i(\sigma) := \sum_{z \in Z} u_i(z) \hat{\phi}(\sigma)(z). \quad (5)$$

Two mixed strategies σ_i and σ'_i of player $i \in \{1, \dots, I\}$ are said to be **realization equivalent** if $\hat{\phi}(\sigma_i, \sigma_{-i}) = \hat{\phi}(\sigma'_i, \sigma_{-i})$ for all $\sigma_{-i} \in \Sigma_{-i}$.

Remark 6. If Nature is active, the random outcome induced by a mixed strategy profile $\sigma \in \Sigma$ is the probability distribution $\hat{\phi}(\sigma) \in \Delta(Z)$ which assigns to terminal node z the probability

$$\hat{\phi}(\sigma)(z) = \sum_{s \in S} \sigma(s) \phi(s)(z),$$

where $\phi : S \rightarrow \Delta(Z)$ is the function constructed in Remark 5. As before, x_0 is reached with probability one, and each $x_1 \in p^{-1}(x_0)$ is reached with probability $\sigma_0(\alpha(x_1))$, under σ . For any other node $x \in X$ that is neither the initial node nor an immediate successor of it, given any $p \in [0, 1]$ we say x is **reached with probability p under σ** if, letting x' be the unique predecessor of x in $p^{-1}(x_0)$, and denoting by $S(x)$ the set of strategy profiles under which x is reached in the

sense of Remark 5, we have $\sigma_0(\alpha(x')) \sum_{s \in S(h)} \sigma(s) = p$. (If $p = 0$, we may just say that x is not reached under σ .) Player i 's expected payoff under σ is again defined as in (5), and realization equivalence is defined exactly as before.

Mixed strategies are often complicated to visualize, and it is more convenient to work with a simpler kind of randomization by player i , as in the following definition.

Definition 12. A **behavior strategy** for player $i \in \{1, \dots, I\}$ specifies a probability distribution $b_i(h) \in \Delta(A(h))$ for each $h \in H_i$. The **mixed representation** of a behavior strategy b_i is the mixed strategy of i that assigns to each $s_i \in S_i$ the probability

$$\prod_{h \in H_i} b_i(h)(s_i(h))$$

The probabilities of reaching the various nodes in the game and the induced distributions on terminal nodes and payoffs corresponding to a behavior strategy profile $b = (b_1, \dots, b_I)$ are simply the ones given by the profile of mixed representations of b_1, \dots, b_I . Again with slight abuse of notation, we denote the induced distribution on outcomes and player i 's expected payoff by $\hat{\phi}(b)$ and $U_i(b)$, respectively. Analogously, we say a behavior strategy b_i is realization equivalent to a mixed strategy $\sigma_i \in \Sigma_i$ if the mixed representation of b_i is realization equivalent to σ_i .

As an example, consider the game in Figure 2 and the behavior strategy of player 1 that chooses L with probability $1/4$ and ℓ with probability $1/3$. The mixed representation of this behavior strategy is the mixed strategy that puts probabilities $1/12$, $1/6$, $1/4$, and $1/2$ on $L\ell$, Lr , $R\ell$, and Rr , respectively.

Behavior strategies are much easier to visualize than mixed strategies. One must simply specify, for each information set h of player i , the probabilities with which i will choose the various actions available at h . Thus, the set of all behavior strategies of player i is the cartesian product

$$B_i := \prod_{h \in H_i} \Delta(A(h)).$$

While behavior strategies are simpler objects than mixed strategies, they are also less general. Indeed, mixed strategies allow for correlation among different information sets of the same player,

whereas mixed representations of behavior strategies do not. For instance, in Figure 2 there is no behavior strategy for player 1 whose mixed representation is the mixed strategy that puts equal probability on $L\ell$, Lr , and Rr . (Make sure you are able to prove this claim.) However, for the purposes of computing distributions on outcomes, and hence expected payoffs, this difference has no consequences; Kuhn's theorem is the formal statement of the latter claim. ★

Theorem 3 (Kuhn, 1953). *In a finite extensive form game with perfect recall, every mixed strategy has a realization equivalent behavior strategy.*

Proof. A formal proof can be found in various textbooks (see, for instance, Theorem 4.1 in Myerson's book). For any $\sigma_i \in \Sigma_i$, the realization equivalent behavior strategy b_i whose existence is established in the theorem can be described as follows. Say an information set $h \in H_i$ is **reachable** under a mixed strategy σ'_i if there exists σ_{-i} such that some node in h (and hence all nodes in h , by perfect recall) is reached with positive probability under (σ'_i, σ_{-i}) . For every information set $h \in H_i$ and every $a \in A(h)$, let $S_i(h)$ be the set of all $s_i \in S_i$ such that h is reachable under s_i , let $S_i(a)$ be the set of all $s_i \in S_i$ such that $s_i(h) = a$, and let $S_i(h, a)$ be the set of all $s_i \in S_i(h)$ such that $s_i(h) = a$. Then the probability assigned by b_i to action a is

$$b_i(h)(a) = \begin{cases} \frac{\sum_{s_i \in S_i(h)} \sigma_i(s_i)}{\sum_{s_i \in S_i(h)} \sigma_i(s_i)} & \text{if } h \text{ is reachable under } \sigma_i, \\ \sum_{s_i \in S_i(a)} \sigma_i(s_i) & \text{otherwise.} \end{cases}$$

In other words, the probability with which b_i chooses a at h is either the *conditional* probability of choosing a under σ_i , given that h is reachable, or, if h is not reachable under σ_i , it is defined arbitrarily. □

As an illustration, you should verify that, in the game of Figure 2, the mixed strategy that puts equal probability on $L\ell$, Lr , and Rr is realization equivalent to the behavior strategy that chooses L with probability $2/3$ and ℓ with probability zero. ★

3.2.1. Continuation Strategies, Continuation Outcomes, and Continuation Payoffs

Take any nonterminal node $x \in X \setminus (Z \cup \{x_0\})$ and write H^x to denote the set of information sets containing either x or some successor of x , that is,

$$H^x = H(x) \cup \{h \in H : h = H(x') \text{ for some } n \geq 1 \text{ and some } x' \in X \text{ such that } x = p_n(x')\}.$$

The **continuation strategy** induced by a strategy profile $s = (s_1, \dots, s_I)$ from node x is the profile of functions s_1^x, \dots, s_I^x such that s_i^x is the restriction of s_i to $H_i \cap H^x$ for every player i . (If x is the initial node of a subgame — see Definition 13 below — then these functions will be well defined strategies for the subgame.) Since actions leading to distinct immediate successors of a nonterminal node are labeled differently by α , there exists a unique $x_1 \in p^{-1}(x)$ such that $\alpha(x_1) = s_{\iota(H(x))}(H(x))$, a unique $x_2 \in p^{-1}(x_1)$ such that $\alpha(x_2) = s_{\iota(H(x_1))}(H(x_1))$, and so on. Thus we have a sequence x, x_1, \dots, x_n such that $x = p(x_1)$, $x_k = p(x_{k+1})$ for every $k = 1, \dots, n-1$, and $x_n \in Z$. The terminal node x_n is the **continuation outcome** associated to x and s and is denoted by $\phi(s|x)$. The **continuation payoff** to any player i associated to x and s is the payoff to player i from the continuation outcome, that is, $U_i(s|x) := U_i(\phi(s|x))$.

It is not obvious how to define a continuation strategy associated to x and a *mixed* strategy σ , especially if x is not reached under σ . However, an extension of our definitions of continuation strategy, outcome, and payoffs to *behavior* strategies does make sense. Indeed, contrary to a mixed strategy profile, a behavior strategy profile (even one under which x is not reached) tells us explicitly what player $\iota(H(x))$ would do at node x . Just like we did for pure strategies, we define the **continuation behavior strategy** induced by a behavior strategy profile $b = (b_1, \dots, b_I)$ from a nonterminal node x as the profile of functions b_1^x, \dots, b_I^x such that b_i^x is the restriction of b_i to $H_i \cap H^x$ for every player i . (The above comment about subgames also applies here.) The random continuation outcome — i.e. the element of $\Delta(Z)$ — associated to x and b will be denoted by $\hat{\phi}(b|x)$, and the associated expected continuation payoff to any player i will be $U_i(b|x) := \sum_{z \in Z} \hat{\phi}(b|x)(z) u_i(z)$. The probability distribution $\hat{\phi}(b|x) \in \Delta(Z)$ is determined in the obvious way, as follows. Let Z^x denote the set of terminal nodes that are successors of x . For every $z \in Z^x$ there exists a unique sequence $x_1^z, \dots, x_{n(z)}^z$ such that $x = p(x_1^z)$, $x_k^z = p(x_{k+1}^z)$ for every $k = 1, \dots, n(z)-1$, and $x_{n(z)}^z = z$. Then, for every terminal node $z \in Z$, we set

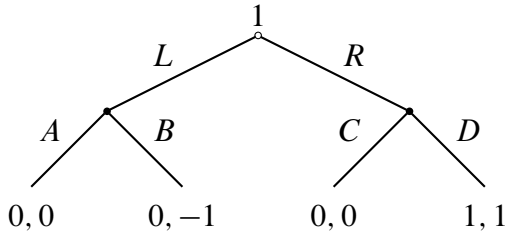
$\hat{\phi}(b|x)(z) = 0$ if $z \notin Z^x$ and

$$\hat{\phi}(b|x)(z) = b_{\iota(H(x))}(\alpha(x_1^z)) \cdot \prod_{k=1}^{n(z)-1} b_{\iota(H(x_k^z))}(\alpha(x_{k+1}^z))$$

if $z \in Z^x$.

3.3. Subgame Perfect Equilibrium

The notion of Nash equilibrium very often leads to unreasonable predictions in extensive form games. Consider the following extensive form game and its associated game in normal form.



	AC	AD	BC	BD
L	0, 0	0, 0	0, -1	0, -1
R	0, 0	1, 1	0, 0	1, 1

The profile (L, AC) is a Nash equilibrium, yet it is unreasonable. Indeed, player 2 is supposed to choose C if player 1 chooses R (which he does not), whereas D would give more. Similarly, (R, BD) is a Nash equilibrium, but it is unreasonable because it requires player 2 to play B after L , whereas A would give more.

The idea of subgame perfect equilibrium (Selten, 1965) is that behavior in parts of the game that can be regarded as games in themselves should agree with Nash equilibria of them.

Definition 13. Let $\langle N, X, p, A, \alpha, H, \iota, \sigma_0, (u_i)_{i \in N} \rangle$ be a game in extensive form. A **subgame** of this game is any extensive form game $\langle N, \hat{X}, \hat{p}, A, \hat{\alpha}, \hat{H}, \hat{\iota}, \sigma_0, (\hat{u}_i)_{i \in N} \rangle$ such that:

- $\hat{X} \subseteq X$ and $\hat{p} : \hat{X} \rightarrow \hat{X} \cup \emptyset$ is the restriction of p to \hat{X} ;
the set of terminal nodes of the subgame is $\hat{Z} := \{\hat{x} \in \hat{X} : \hat{p}^{-1}(\hat{x}) = \emptyset\}$; the initial node is denoted \hat{x}_0 ; it is assumed that \hat{X} contains all successors of \hat{x}_0 , i.e. for every node $x \in X$ satisfying $x \in \hat{X}$ we must have $p^{-1}(x) \subseteq \hat{X}$; note that this implies $\hat{Z} \subseteq Z$;
- $\hat{\alpha} : \hat{X} \setminus \{\hat{x}_0\} \rightarrow A$ is the restriction of α to $\hat{X} \setminus \{\hat{x}_0\}$;
- $\hat{H} \subseteq H$, and $\hat{\iota} : \hat{H} \rightarrow N$ is the restriction of ι to \hat{H} ; note that $\hat{H} \subseteq H$ implies that if a node of the game is also a node of the subgame, then all nodes in the same information set

are also in the subgame, i.e. $x \in \hat{X}$ implies $\hat{H}(x) = H(x) \subseteq \hat{X}$;

- $\hat{u}_i : \hat{Z} \rightarrow \mathbb{R}$ is the restriction of u_i to \hat{Z} for each player i .

Obviously, every extensive form game is a subgame of itself. The game in Figure 1 is the *only* subgame of itself. The game in Figure 2 instead has three more subgames besides itself, depicted in Figure 3.

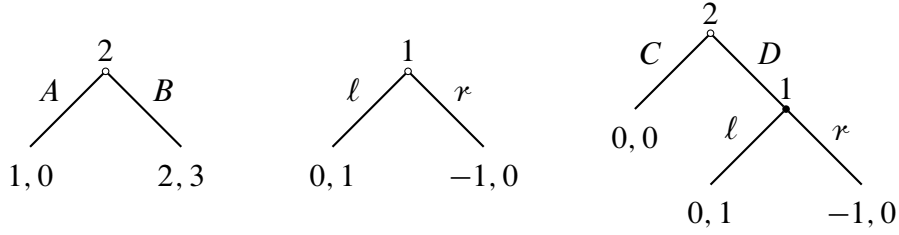


Figure 3: The subgames (other than itself) of the game in Figure 2

Note that, given a pure (or behavior) strategy profile, the continuation strategy (or continuation behavior strategy) from the initial node of a subgame is a well defined strategy profile (or behavior strategy profile) for the subgame.

Definition 14. A **subgame perfect equilibrium** is a strategy profile such that, for every subgame, the continuation strategy profile from the initial node of the subgame is a Nash equilibrium of the subgame. A **subgame perfect equilibrium in behavior strategies** is a behavior strategy profile such that, for every subgame, the mixed extension of the continuation behavior strategy profile from the initial node of the subgame is a mixed strategy Nash equilibrium of the subgame.

Every finite extensive form game has at least one subgame perfect equilibrium in behavior strategies. In games with **perfect information**, i.e. games where every information set contains exactly one node, at least one pure strategy subgame perfect equilibrium must exist. (We do not prove these claims formally, but we will illustrate why they are true with a few examples — see 3.4.1 below.) In fact, in such games checking whether a behavior strategy profile is subgame perfect is rather easy: the following result indeed shows that all one needs to check is that, at each node, the player moving at that node has no incentive to deviate at that node *only*. (Since the proposition deals with perfect information games, in its statement and proof we identify each nonterminal node with the corresponding information set.)

Proposition 3 (One-Shot Deviation Principle). *In a perfect information extensive form game, a behavior strategy profile b is subgame perfect if and only if $U_i(b_i, b_{-i}|x) \geq U_i(b'_i, b_{-i}|x)$ for every $i \in \{1, \dots, I\}$, every $x \in H_i$, and every b'_i such that $b'_i(x') = b_i(x')$ for all $x' \in H_i \setminus \{x\}$.*

Proof. Necessity is obvious. To prove sufficiency, assume b is not subgame perfect. Then the set

$$D = \{x \in X \setminus Z : U_i(b'_i, b_{-i}|x) > U_i(b|x) \text{ for some } i \in \{1, \dots, I\} \text{ and some } b'_i \in B_i\}$$

is nonempty. Since X is finite, D is finite, hence there exists $\bar{x} \in D$ such that no successor of \bar{x} belongs to D . Since $\bar{x} \in D$, there exists $j \in \{1, \dots, I\}$ and $b'_j \in B_j$ such that

$$U_j(b|\bar{x}) < U_j(b'_j, b_{-j}|\bar{x}). \quad (6)$$

Since no successor of \bar{x} belongs to D , we have

$$U_j(b'_j, b_{-j}|x) \leq U_j(b|x) \quad \text{for all } x \in p^{-1}(\bar{x}). \quad (7)$$

These inequalities imply that $j = \iota(\bar{x})$. Indeed, $j \neq \iota(\bar{x})$ would imply, by (7) and by the definition of continuation payoff,

$$\begin{aligned} U_j(b'_j, b_{-j}|\bar{x}) &= \sum_{x \in p^{-1}(\bar{x})} b_{\iota(\bar{x})}(\bar{x})(\alpha(x)) U_j(b'_j, b_{-j}|x) \\ &\leq \sum_{x \in p^{-1}(\bar{x})} b_{\iota(\bar{x})}(\bar{x})(\alpha(x)) U_j(b|x) \\ &= U_j(b|\bar{x}), \end{aligned}$$

which contradicts (6). Now define $b''_j \in B_j$ as follows: $b''_j(\bar{x}) = b'_j(\bar{x})$ and $b''_j(x) = b_j(x)$ for every other $x \in H_j \setminus \{\bar{x}\}$. Then, again by (7) and by the definition of continuation payoff, and

using the fact that $j = \iota(\bar{x})$, we have

$$\begin{aligned} U_j(b'_j, b_{-j}|\bar{x}) &= \sum_{x \in p^{-1}(\bar{x})} b'_j(\bar{x})(\alpha(x)) U_j(b'_j, b_{-j}|x) \\ &\leq \sum_{x \in p^{-1}(\bar{x})} b'_j(\bar{x})(\alpha(x)) U_j(b|x) \\ &= U_j(b''_j, b_{-j}|\bar{x}). \end{aligned}$$

The latter, together with (6), gives $U_j(b|\bar{x}) < U_j(b''_j, b_{-j}|\bar{x})$. Since b''_j differs from b_j only at \bar{x} , we are done. \square

Note that the analogous property does *not* hold for Nash equilibria. In the game of Figure 4, the profile $(L\ell, B)$ is not Nash, and yet no one-shot profitable deviation exists, because in order to profitably deviate player 1 must change his prescribed action at *both* of his information sets (i.e. choose Rr instead of $L\ell$).

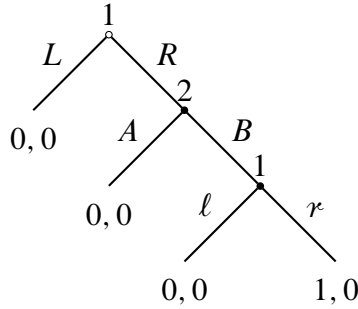


Figure 4: The One-Shot Deviation property does not hold for Nash equilibria.

3.4. Examples

3.4.1. Finding Subgame Perfect Equilibria of Perfect Information Games

For perfect information extensive form games, finding the subgame perfect equilibria is very easy, using the **backward induction** method. Consider the variant of the game in Figure 2 depicted in Figure 5. For each node whose successors are all terminal nodes, we highlight the action leading to the highest payoff for the player moving at that node with a thick black line, if there is a unique such action; if there are more than one optimal actions, then we highlight each of them with a different color, other than black. Then we proceed analogously going up in the game tree, highlighting the

actions leading to the highest payoffs for the player whose turn it is, using possibly different colors accordingly.

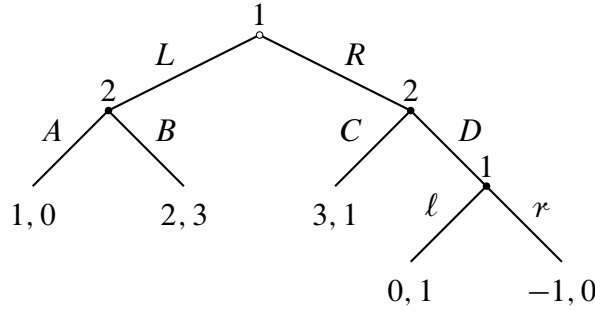


Figure 5: A variant of the game in Figure 2.

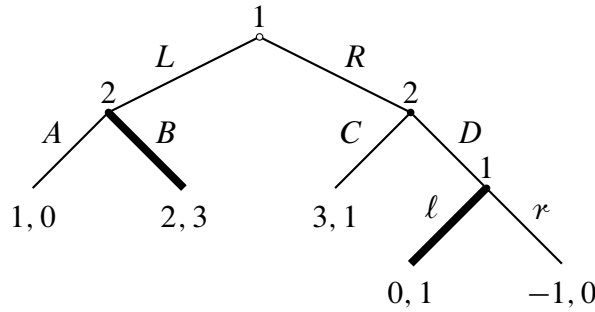


Figure 6: Backward Induction in the game of Figure 5: Step I.

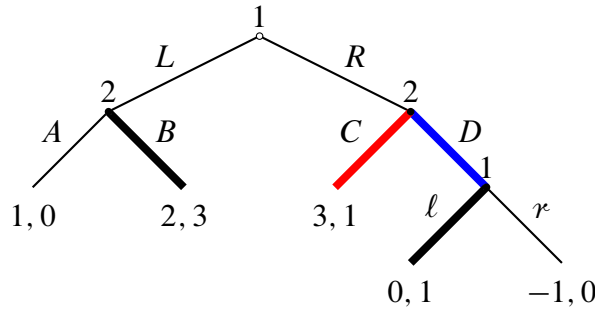


Figure 7: Backward Induction in the game of Figure 5: Step II.

This allows us to find all pure strategy subgame perfect equilibria of the game. Indeed, a strategy profile is a subgame perfect equilibrium if and only if it is described by highlighted branches that are either black or of the same color.

In the game of Figure 5 there are then two pure strategy subgame perfect equilibria, $(L\ell, BD)$ and $(R\ell, BC)$. There are also many (a continuum of) mixed strategy equilibria. Indeed, every (σ_1, σ_2) such that $\sigma_1(L\ell) = 1$, $\sigma_2(BC) < 2/3$, and $\sigma_2(BD) = 1 - \sigma_2(BC)$ is subgame perfect. Moreover, every (σ_1, σ_2) such that $\sigma_1(R\ell) = 1$, $\sigma_2(BC) > 2/3$, and $\sigma_2(BD) = 1 - \sigma_2(BC)$ is

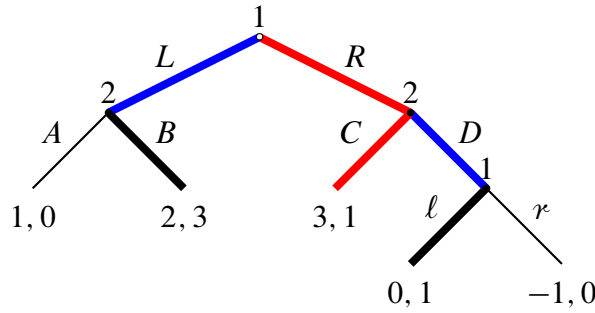


Figure 8: Backward Induction in the game of Figure 5: Step III.

subgame perfect. Finally, every (σ_1, σ_2) such that $\sigma_2(BC) = 2/3$ and $\sigma_2(BD) = 1 - \sigma_2(BC)$ is subgame perfect. (Make sure you understand why all these mixed strategy profiles are subgame perfect equilibria.) ★

Obviously, if the payoffs at the various terminal nodes are such that no tie ever occurs (e.g. if they are all different), then there will be a unique subgame perfect equilibrium, and it will be in pure strategies.

3.4.2. Trembling Hand Perfect and Subgame Perfect

Given an extensive form game, one can construct the **agent normal form** representation by treating different information sets as different players. Figure 9 depicts an extensive form game and its normal form representation. Figure 10 depicts its agent normal form representation.

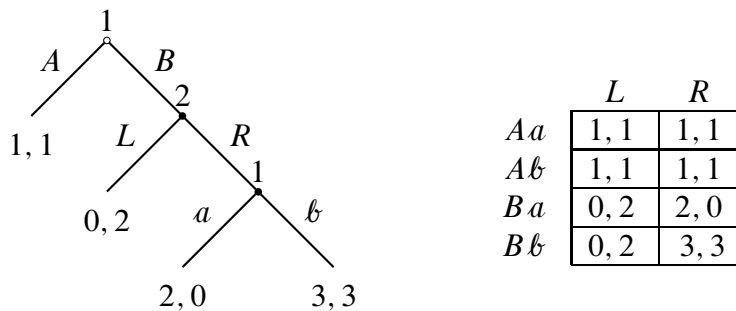


Figure 9: An extensive form game and its normal form representation.

		L	R
A		1, 1, 1	1, 1, 1
B		0, 2, 0	2, 0, 2
		a	b

Figure 10: The agent normal form representation of the extensive form game in Figure 9.

This example shows that trembling-hand perfection in the normal form does not imply subgame perfection. The only subgame perfect equilibrium is (Bb, R) , and yet (Aa, L) is trembling-hand perfect. The example also shows that this problem disappears when we consider the agent normal form; in the latter, a is weakly dominated by b , hence (A, a, L) cannot be trembling-hand perfect in the agent normal form. Indeed, strategies that are trembling-hand perfect in the agent normal form (also called **extensive form perfect**) are, together with some system of beliefs μ , necessarily sequential and hence subgame perfect (the definition of sequential equilibrium will be given shortly).

4. Extensive Form Games: Further Topics

While subgame perfection is good at ruling out unreasonable equilibria in perfect information games, the same cannot be said about imperfect information games. The game in Figure 11 provides an example. The profile (L, B) is a subgame perfect equilibrium, since it is Nash and the game has no other subgames than itself.

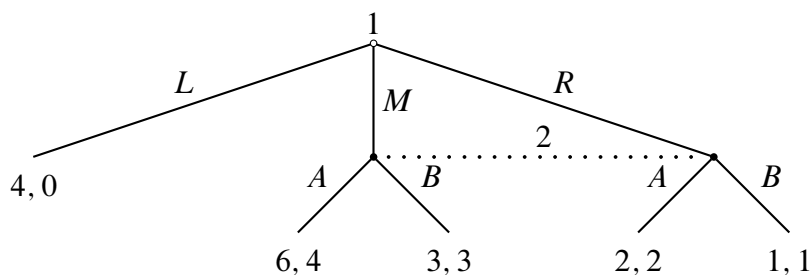


Figure 11: The profile (L, B) is subgame perfect, but not really reasonable.

The problem is that player 2 is choosing strategy B which, should player 2's information set be actually reached, would give strictly less than A regardless of how the information set is reached, i.e. regardless of whether player 1 has chosen M or R .

In the game of Figure 11, A does always better than B conditional on player 2's information set being reached. Now consider the game in Figure 12. In this game, whether A or B is best depends on player 2's belief about ℓ and r .

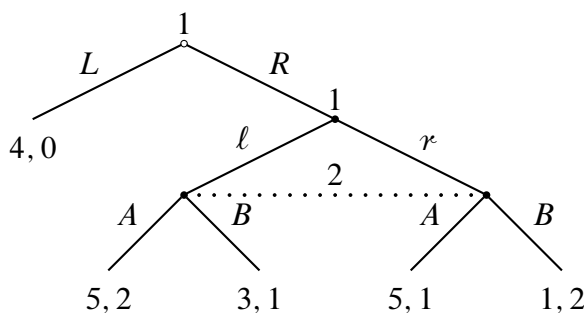


Figure 12: The profile $(L\ell, B)$ is not even subgame perfect, yet it is “weak perfect Bayesian”.

Presumably, these beliefs should be computed using equilibrium strategies, as the conditional probability of player 1 choosing ℓ , conditional on his not choosing L . But what if player 1's equilibrium strategy prescribes L with probability one? In this case, the notion of “weak perfect Bayesian equilibrium” (see Mas-Colell, Whinston, and Green) allows any belief by player 2.

4.1. Sequential Equilibrium

The notion of sequential equilibrium (Kreps and Wilson, 1982) extends the idea of subgame perfection to games with imperfect information in a meaningful way, fixing many problems with subgame perfection (and the various versions of “perfect Bayesian equilibrium”) like the ones illustrated above.

An **assessment** is a pair (μ, b) where b is a behavior strategy and μ is a **system of beliefs**, i.e. an assignment of a belief $\mu(h) \in \Delta(h)$ for each information set $h \in H$. Given an assessment (μ, b) , at each information set h player $i(h)$ can compute his expected continuation payoff, namely,

$$U_{i(h)}(b|h, \mu) := \sum_{x \in h} \mu(h)(x) U_{i(h)}(b|x).$$

This is simply the expected value of the continuation payoffs associated to b and the various nodes in h , when the probabilities of these nodes are given by μ .

An assessment (μ, b) is **sequentially rational** if $U_{i(h)}(b|h, \mu) \geq U_{i(h)}(b'_{i(h)}, b_{-i(h)}|h, \mu)$ for every $h \in H$ and every $b'_{i(h)} \in B_{i(h)}$. We say (μ, b) is **consistent** if there exists a sequence of assessments $(\mu^1, b^1), (\mu^2, b^2), \dots$ converging to (μ, b) and such that, for every $n = 1, 2, \dots$, the profile b^n is completely mixed and μ^n is obtained from σ_0 and b^n using Bayes' rule.

Definition 15. A **sequential equilibrium** is a consistent and sequentially rational assessment.

Let us go back to the examples above. In the game of Figure 11, the profile (L, B) violates sequential rationality; indeed, there is no system of beliefs that can make B sequentially rational. In the game of Figure 12, the profile $(L\ell, B)$ satisfies sequential rationality according to any system of beliefs μ for which r is at least as likely as ℓ . However, such μ cannot be consistent with $(L\ell, B)$. Indeed, one cannot converge to such μ with a sequence $(\mu^1, b^1), (\mu^2, b^2), \dots$ such that the probability given by b^n to ℓ converges to 1 and μ^n is derived from b^n using Bayes rule.

For sequential equilibria, the following analogue to the one-shot deviation principle holds.

Proposition 4 (One-Shot Deviation Principle for Sequential Equilibria). *A consistent assessment (μ, b) is sequentially rational if and only if it is myopically rational, that is, if and only if one has $U_i(b|h, \mu) \geq U_i(b'_i, b_{-i}|h, \mu)$ for every $i \in \{1, \dots, I\}$, every $h \in H_i$, and every $b'_i \in B_i$ such that $b'_i(h') = b_i(h')$ for all $h' \in H_i \setminus \{h\}$.*

Proof. See Exercise 227.1 in Osborne and Rubinstein's *A Course in Game Theory*. \square

4.2. Alternating Offer Bargaining

Here we deal with the following game, first analyzed by Rubinstein in 1982. Two agents, 1 and 2, must decide how to share the interval $[0, 1]$. A proposal is a division $(x, 1 - x)$ with $0 \leq x \leq 1$, where x represents the share to player 1. The agents take turns to make proposals. Player 1 makes proposals, and player 2 either accepts or rejects, at *odd* times $t = 1, 3, 5, \dots$. Player 2 makes proposals, and player 1 accepts or rejects, at *even* times $t = 2, 4, 6, \dots$. If at time t a proposal $(x, 1 - x)$ is accepted, the game ends, and payoffs are $\delta_1^{t-1}x$ to player 1 and $\delta_2^{t-1}(1 - x)$ to player 2. (The discount factors δ_1 and δ_2 are assumed to be strictly between zero and one.) If the proposal is rejected, the game proceeds to period $t + 1$. Perpetual disagreement yields a payoff of zero to both players. The extensive form is illustrated in Figure 13.

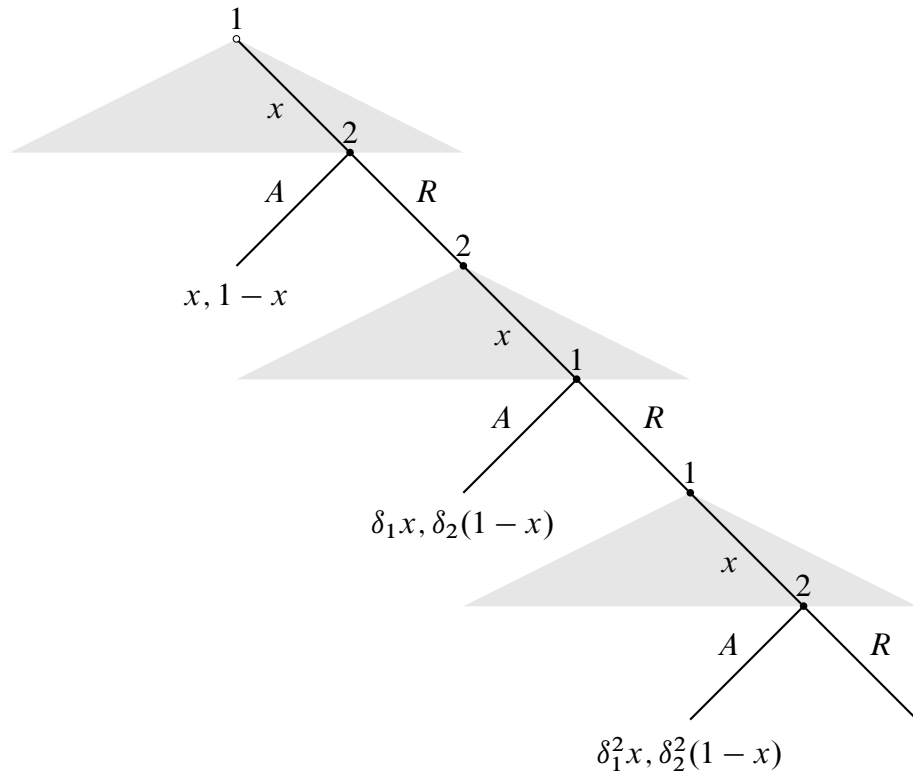


Figure 13: Rubinstein's Bargaining Game.

While the game has infinite horizon and infinite action sets at those nodes where a player proposes, it is a game with perfect information, and it is not too difficult to apply the idea of subgame perfection to it.

4.2.1. Finite Horizon

We begin analyzing the finite horizon case; disagreement for T periods ends the game. This game can be solved using backward induction.

First assume T is odd. Then in period T player 1 will propose $(1, 0)$ and player 2 will accept. Thus, at $T - 1$, player 2 will propose $(\delta_1, 1 - \delta_1)$ and player 1 will accept, and so on. Working backwards, let us denote the period t proposal by $(x_{T-t+1}, 1 - x_{T-t+1})$. Thus, $(x_1, 1 - x_1)$ is the proposal at time T , $(x_2, 1 - x_2)$ is the proposal at time $T - 1$, and so on. Then we have $x_1 = 1$ and

$$x_k = \begin{cases} \delta_1 x_{k-1} & \text{if } k \text{ is even} \\ 1 - \delta_2(1 - x_{k-1}) & \text{if } k \text{ is odd} \end{cases} \quad (8)$$

for all $k = 2, \dots, T$. The unique subgame perfect equilibrium is: at each odd t , player 1 proposes $(x_{T-t+1}, 1 - x_{T-t+1})$ and player 2 accepts $(x, 1 - x)$ if and only if $x \leq x_{T-t+1}$; at each even t , player 2 proposes $(x_{T-t+1}, 1 - x_{T-t+1})$ and player 1 accepts $(x, 1 - x)$ if and only if $x \geq x_{T-t+1}$. You should verify that, for each $k \geq 3$ odd, we have \star

$$x_k = (1 - \delta_2) \sum_{n=0}^{(k-3)/2} (\delta_1 \delta_2)^n + (\delta_1 \delta_2)^{(k-1)/2}$$

You should also verify that $x_{k+2} < x_k$ for every $k \geq 1$ odd.

Now assume T is even. Then in period T player 2 will propose $(0, 1)$ and player 1 will accept. Thus, at $T - 1$, player 1 will propose $(1 - \delta_2, \delta_2)$ and player 2 will accept, and so on. Working backwards, let us denote the period t proposal by $(y_{T-t+1}, 1 - y_{T-t+1})$. Thus, $(y_1, 1 - y_1)$ is the proposal at time T , $(y_2, 1 - y_2)$ is the proposal at time $T - 1$, and so on. Then we have $y_1 = 0$ and

$$y_k = \begin{cases} \delta_1 y_{k-1} & \text{if } k \text{ is odd} \\ 1 - \delta_2(1 - y_{k-1}) & \text{if } k \text{ is even} \end{cases} \quad (9)$$

for all $k = 2, \dots, T$. The unique subgame perfect equilibrium is: at each odd t , player 1 proposes $(y_{T-t+1}, 1 - y_{T-t+1})$ and player 2 accepts $(y, 1 - y)$ if and only if $y \leq y_{T-t+1}$; at each even t , player 2 proposes $(y_{T-t+1}, 1 - y_{T-t+1})$ and player 1 accepts $(y, 1 - y)$ if and only if $y \geq y_{T-t+1}$. You should verify that, for each $k \geq 2$ even, we have \star

$$y_k = (1 - \delta_2) \sum_{n=0}^{(k-2)/2} (\delta_1 \delta_2)^n$$

You should also verify that $y_{k+2} > y_k$ for every $k \geq 2$ even.

4.2.2. Infinite Horizon

We will first establish existence of and describe a *stationary* subgame perfect equilibrium, i.e. a subgame perfect equilibrium in which the proposal made by player 1 (resp. player 2) at all odd (resp. even) periods is the same, i.e. the same regardless of what happened before, and player 2's (resp. player 1's) acceptance or rejection decision at each odd (resp. even) period is the same, i.e. only dependent on the proposal received in that period. Then we will show that this is in fact the only equilibrium.

Suppose that there exist $x^* \in [0, 1]$ and $z^* \in [0, 1]$ with the property that

$$1 - x^* = \delta_2(1 - z^*) \quad \text{and} \quad z^* = \delta_1 x^*. \quad (10)$$

Then you can check that the following strategy profile is indeed a stationary subgame perfect equilibrium: at each odd t , player 1 proposes $(x^*, 1 - x^*)$ and player 2 accepts $(x, 1 - x)$ if and only if $x \leq x^*$; at each even t , player 2 proposes $(z^*, 1 - z^*)$ and player 1 accepts $(z, 1 - z)$ if and only if $z \geq z^*$. Using (10), we get

$$x^* = \frac{1 - \delta_2}{1 - \delta_1 \delta_2} = \lim_{k \rightarrow \infty} x_k = \lim_{k \rightarrow \infty} y_k \quad z^* = \frac{\delta_1(1 - \delta_2)}{1 - \delta_1 \delta_2}.$$

In principle, non-stationary SPE might exist. However, we have the following.

Lemma 4. *In every subgame perfect equilibrium, player 1's payoff is at most x_k for every odd $k \geq 1$, and at least y_k for every even $k \geq 2$.*

Proof. Pick any SPE with payoffs u_1 and u_2 to player 1 and player 2, respectively. First we prove that $u_1 \leq x_k$ for every odd k . Since $x_1 = 1$, this is obvious for $k = 1$. Now assume it is true for some k odd. Then $u_1 + u_2 \leq 1$ and $u_2 \geq \delta_2(1 - \delta_1 x_k)$. (Why?) Thus, ★

$$u_1 \leq 1 - u_2 \leq 1 - \delta_2(1 - \delta_1 x_k) = x_{k+2}.$$

Now we prove $u_1 \geq y_k$ for every even k . Since $y_2 = 1 - \delta_2$, this is obvious for $k = 2$. (Why?) Suppose it is true for some even $k \geq 2$. Once again $u_1 + u_2 \leq 1$ and, moreover, ★
 $u_1 \geq 1 - \delta_2(1 - \delta_1 y_k) = 1 - \delta_2(1 - \delta_1 y_k)$ by (9), because (by the induction hypothesis) player 2 gets a continuation payoff no greater than $1 - \delta_1 y_k$ in every subgame starting with a proposal by player 2, so he must accept any proposal $(x, 1 - x)$ such that $1 - x > \delta_2(1 - \delta_1 y_k)$. □

Since the sequences x_k and y_k have the same limit, namely x^* , an immediate consequence of the lemma is that in every subgame perfect equilibrium player 1's payoff is equal to x^* . This implies that player 2's payoff in every subgame perfect equilibrium is $1 - x^*$ (why?) and, moreover, ★
 this uniqueness of equilibrium payoffs also implies that equilibrium *strategies* are unique (why?) ★★
 and thus coincide with the stationary ones we described above.

4.3. Examples

4.3.1. Inconsistent Assessments and One-Shot Deviations

An inconsistent assessment can fail to be sequentially rational even if no one-shot profitable deviation exists. In the game in Figure 14, the assessment (μ, b) such that $b_1(L) = b_1(r) = b_2(A) = 1$ and $\mu(y) = 1$ is inconsistent; however. There is no profitable one-shot deviation, and yet sequential rationality is violated, since player 1 should choose L at his first information set and then ℓ at his second information set.

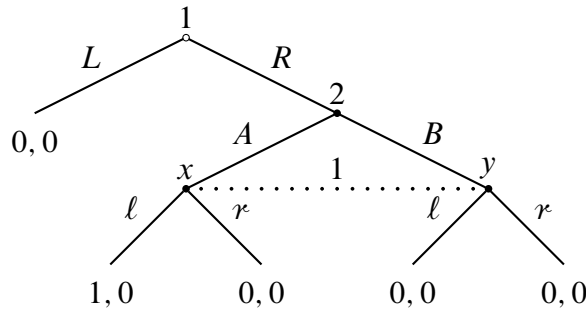


Figure 14: An assessment (μ, b) with $b_2(A) = 1$ and $\mu(y) = 1$ is inconsistent.

4.3.2. More on Out of Equilibrium Beliefs

The strategy profile $(L\ell, B)$ in the game in Figure 12 is not even subgame perfect, yet it is a “weak perfect Bayesian equilibrium”. One could think of fixing this problem by adding the obvious

requirement that μ must put probability zero on \mathcal{I} . In other words, one could think of amending the definition of weak perfect Bayesian equilibrium by requiring that Bayes rule be used to determine beliefs from behavior strategies not just “whenever possible”, but also at information sets reached with probability zero.

The one just described is the idea behind **perfect Bayesian equilibrium**, a notion that has been defined for a particular class of games (multi-stage games with observable actions and possibly incomplete information, among which signaling games) and is used often in the applied literature. But even in this class of games, it can happen that a PBE is not a sequential equilibrium. See e.g. Osborne and Rubinstein for more on this.

5. Repeated Games

A repeated game consists in the infinite repetition of a normal form game $\langle N, (A_i, u_i)_{i \in N} \rangle$ called the **stage game**. Note that we write A_i instead of S_i to denote the stage game strategy set of player i . This is because we reserve the symbol S_i to denote i 's set of strategies in the repeated game. A strategy of the stage game is instead called an **action**. Here we assume that $N = \{1, 2\}$ and A_1, A_2 and hence $A := A_1 \times A_2$ are finite. In addition, we assume **perfect monitoring**, i.e. we assume that at each stage the two players choose their actions simultaneously and, once actions have been chosen, they are observed before the game moves to the following stage.

A **history** is either the empty set or a finite sequence of action profiles, representing what has happened so far in the game. The empty set \emptyset is the initial history, i.e. nothing has happened yet. Thus, the set of initial histories is $\mathcal{H}^0 := \{\emptyset\}$ and, for all $t \geq 2$, the set of period t histories is $\mathcal{H}^t := A^{t-1}$. The set of all histories is

$$\mathcal{H} = \bigcup_{t \geq 1} \mathcal{H}^t$$

A **strategy** for player i in the repeated game is a function $s_i : \mathcal{H} \rightarrow A_i$. In other words, a strategy for i specifies an action $s_i(\emptyset)$ to take at the first stage and, for each $t \geq 2$ and for each possible period t history h^t , an action $s_i(h^t)$ to take in stage t . The set of all strategies for player i is denoted S_i and as usual we define $S := S_1 \times S_2$. For the sake of brevity, given any $s = (s_1, s_2) \in S$ and any history $h \in \mathcal{H}$, we write $s(h)$ instead of $(s_1(h), s_2(h))$. An **outcome** in the repeated game is an infinite sequence in A , i.e. the set of outcomes is

$$A^\infty := \prod_{t=1}^{\infty} A$$

Note that an outcome is different from a history; the latter has finite length, the former has infinite length. A strategy profile s induces a sequence of histories and an outcome in the obvious way. It also induces continuation strategies and continuation streams of payoffs starting from any history.

5.1. Discounting

Subscripts indicate players, superscripts indicate periods. An outcome is an infinite sequence $(a^1, a^2, \dots) \in A^\infty$, so it is not obvious how to compute player i 's payoff in the repeated game. For

instance, since A is finite, if $u_i(a^t) > 0$ for all t , then just taking the sum $u_i(a^1) + u_i(a^2) + \dots$ will give $+\infty$ regardless of what the sequence (a^1, a^2, \dots) actually is.

There are several possibilities to fix this problem, and one of the most commonly used in economics is **discounting**. It is assumed that both players evaluate streams of stage payoffs by taking the average discounted sum according to a common discount factor $\delta \in (0, 1)$. Thus, player i 's payoff from the outcome (a^1, a^2, \dots) will be given by

$$U_i(a^1, a^2, \dots) := (1 - \delta)(u_i(a^1) + u_i(a^2) + \dots) = (1 - \delta) \sum_{t=1}^{\infty} \delta^{t-1} u_i(a^t)$$

Note that this is an *average* discounted sum, i.e. we multiply the discounted sum by $(1 - \delta)$. We do this so that payoffs in the repeated game have the same unit of measurement as the stage game payoffs. For instance, if $u_i(a^t) = 2$ for every t , then $U_i(a^1, a^2, \dots) = 2$, i.e. the payoff from an infinite stream of 2's will give exactly a payoff of 2 in the repeated game. For any strategy profile s , we write $U_i(s)$ to denote the utility to player i corresponding to the outcome induced by s .

Using the discounting criterion to evaluate payoffs is convenient because the one-shot deviation principle applies. (This is because the sets A_1 and A_2 of stage game actions are finite. If they were infinite, for the one-shot deviation principle to apply we should require u_1 and u_2 to be bounded.) In other words, a strategy profile $s = (s_1, s_2)$ is a subgame perfect equilibrium if and only if $U_i(s|h) \geq U_i(s'_i, s_{-i}|h)$ for every $h \in \mathcal{H}$, every $i = 1, 2$, and every $s'_i \in S_i$ for which $s'_i(h') = s_i(h')$ whenever $h' \neq h$.

5.2. Two Folk Theorems

The strategy profile where both players choose a stage-game NE in every period, regardless of history, is a SPE of the repeated game. This is what the following result establishes.

Proposition 5. *Suppose the stage game $\langle N, (A_i, u_i)_{i \in N} \rangle$ has a pure strategy Nash equilibrium $a^* \in A$. Then the strategy profile such that $s(h) = a^*$ for all $h \in \mathcal{H}$ is a subgame perfect equilibrium of the repeated game.*

Proof. A useful exercise for you. □ ★

A profile of payoffs is **feasible** if there is an action profile in the stage game that gives exactly

those payoffs. The set of feasible payoffs is then

$$F = \{(v_1, v_2) \in \mathbb{R}^2 : \text{there exists } a \in A \text{ such that } v_1 = u_1(a) \text{ and } v_2 = u_2(a)\}.$$

Player i 's **pure action minimax utility** is defined as

$$\underline{v}_i = \min_{a_{-i} \in A_{-i}} \max_{a_i \in A_i} u_i(a_i, a_{-i})$$

This is the minimum payoff that i can get, assuming that player i maximizes his payoff knowing the action the other player is choosing. A profile of payoffs (v_1, v_2) is **strictly individually rational** if $v_i > \underline{v}_i$ for all $i = 1, 2$. By the name *Folk Theorem* we mean, roughly, a result of the following type: for every profile of feasible and (strictly) individually rational payoffs, there exists a discount factor $\underline{\delta}$ such that, for all $\delta \geq \underline{\delta}$, there exists a Nash (or subgame perfect) equilibrium of the δ -discounted repeated game that gives the players exactly those payoffs. The following is one simple version of the Folk Theorem.

Theorem 4. *There exists $\underline{\delta}$ such that, for every $\delta \geq \underline{\delta}$ and every $a \in A$ such that $u_i(a) > \underline{v}_i \forall i$, there exists a Nash equilibrium s of the repeated game such that $U_i(s) = u_i(a) \forall i$.*

Proof. Fix any $a \in A$ such that $u_i(a) > \underline{v}_i$ for all $i = 1, 2$. For every $i = 1, 2$, choose

$$a_{-i}^* \in \arg \min_{a_{-i} \in A_{-i}} \max_{a_i \in A_i} u_i(a_i, a_{-i})$$

and

$$\hat{a}_i \in \arg \max_{a_i \in A_i} u_i(a_i, a_{-i}^*).$$

The idea is play a at the first stage and then, in subsequent stages, play a if a was played in every previous stage, and a^* otherwise. Thus, consider the following strategy profile: $s(\emptyset) = a$ and, for all $t \geq 2$,

$$s(a^1, \dots, a^{t-1}) = \begin{cases} a & \text{if } a^1 = \dots = a^{t-1} = a, \\ a^* & \text{otherwise} \end{cases}$$

Player i has no profitable deviation if

$$(1 - \delta) \max_{a'_i \in A_i \setminus \{a_i\}} u_i(a'_i, a_{-i}) + \delta u_i(\hat{a}_i, a_{-i}^*) \leq u_i(a).$$

Since $u_i(\hat{a}_i, a_{-i}^*) = \underline{v}_i$, the left-hand side of the latter inequality tends to \underline{v}_i as $\delta \rightarrow 1$. Thus, there exists $\delta(a)$ such that, for all $\delta \geq \delta(a)$, the inequality is indeed satisfied. Having done this for every a such that $u_i(a) > \underline{v}_i$ for all $i = 1, 2$, letting $\underline{\delta} := \max \{\delta(a) : a \in A, u_i(a) > \underline{v}_i\}$ we are done. \square

The strategy profile constructed in the proof of Theorem 4 need not be subgame perfect. While player i cannot gain by deviating because player $-i$ is using a_{-i}^* , it might happen that player $-i$ does not find it profitable to use a_{-i}^* in the punishment phase. In other words, once something different than a is observed, play switches to a^* being played forever after, and this continuation play need not be Nash. The following result, another version of the Folk Theorem, is based precisely on this observation.

Theorem 5. Choose a Nash equilibrium $a^{*,i} \in A$ of the stage game for every player i . There exists $\underline{\delta}$ such that, for every $\delta \geq \underline{\delta}$ and every $a \in A$ such that $u_i(a) > u_i(a^{*,i}) \forall i$, there exists a subgame perfect equilibrium s of the repeated game such that $U_i(s) = u_i(a) \forall i$.

Proof. The idea is play a at the first stage and then, in subsequent stages, play a if a was played in every previous stage, $a^{*,i}$ if the first deviation from the path (a, a, a, \dots) occurred because of a deviation by player i , and any stage game NE otherwise. The rest of the proof is an exercise for you. $\star\star$ \square

5.3. Examples

5.3.1. Nash and Subgame Perfect Equilibria

Consider the infinite repetition of the stage game *Chicken* depicted in Figure 15.

	A	C
A	-3, -3	6, 1
C	1, 6	4, 4

Figure 15: Chicken

The stage game has two pure strategy NE, namely (C, A) and (A, C) , as well as a mixed strategy NE in which each player chooses A with probability $1/3$. The latter gives an expected payoff of 3 to each player. In the repeated game, the outcome $(C, C), (C, C), (C, C), \dots$, which gives a payoff of 4 to each player, is a NE outcome. An equilibrium that results in (C, C) being played in every stage is as follows: $s(\emptyset) = (C, C)$ and, for all $t \geq 2$,

$$s(a^1, \dots, a^{t-1}) = \begin{cases} (C, C) & \text{if } a^{t-1} = (C, C), \\ (A, C) & \text{if } a^{t-1} = (C, A), \\ (C, A) & \text{if } a^{t-1} = (A, C), \\ (A, A) & \text{if } a^{t-1} = (A, A) \end{cases}$$

The strategy adopted by each player in this profile is called *tit-for-tat*. A player chooses the cooperative action C in the first stage; in each subsequent stage $t = 2, 3, \dots$, he chooses whatever the opponent has chosen in stage $t - 1$. Such a strategy has **finite memory**, because at any stage $t \geq 2$ all one needs to know in order to decide is what happened at $t - 1$. (Note that, in fact, $s(a^1, \dots, a^{t-1})$ does not depend on a^1, \dots, a^{t-2} .)

This strategy profile is a NE of the repeated game if $\delta \geq 2/3$. To prove this,⁷ let U_C (resp. U_A) denote the highest continuation payoff that player 1 can achieve, given that player 2 plays tit-for-tat, at a history where player 2 is supposed to choose C (resp. A).⁸ All we need to show is that $U_C \leq 4$ if $\delta \geq 2/3$. First notice that

$$U_C = \max \{4(1 - \delta) + \delta U_C, 6(1 - \delta) + \delta U_A\}, \quad (11)$$

$$U_A = \max \{-3(1 - \delta) + \delta U_A, (1 - \delta) + \delta U_C\}. \quad (12)$$

By (11), we have $U_C = 4$ if $U_C = 4(1 - \delta) + \delta U_C$. So suppose

$$U_C = 6(1 - \delta) + \delta U_A. \quad (13)$$

⁷Remember that, in order to prove that a strategy profile is Nash, we cannot use the one shot deviation principle.

⁸Note that U_C (resp. U_A) is well defined; this is because player 2 is using tit-for-tat, so the highest continuation payoff that player 1 can get at *any* two histories where player 2 is supposed C (resp. A) is the same.

Then (12) becomes

$$U_A = \max \left\{ -3(1 - \delta) + U_C - 6(1 - \delta), (1 - \delta) + \delta U_C \right\} = (1 - \delta) + \delta U_C.$$

But then the first expression becomes

$$U_C = \max \left\{ 4(1 - \delta) + U_A - (1 - \delta), 6(1 - \delta) + \delta U_A \right\}. \quad (14)$$

Now (13) and (14) together imply $U_A \leq 3$ (why?), hence, again by (13), also $U_C \leq 4$ if $\delta \geq 2/3$. ★

While the strategy profile where each player plays tit-for-tat is Nash for $\delta \geq 2/3$, it is not subgame perfect for any δ . Pick any $t \geq 2$ and any history (a^1, \dots, a^{t-1}) such that $a^{t-1} = AA$. After this history, players are supposed to choose (A, A) in stage t . Both players get a continuation payoff of -3 , whereas each could get at least 1 by choosing C forever after.

One can still achieve the outcome $((C, C), (C, C), \dots)$ as a subgame perfect equilibrium outcome, and actually in more than one way — i.e. there are more than one SPE of the game resulting in that outcome. Consider the following strategy profile: $s(\emptyset) = (C, C)$ and, for all $t \geq 2$,

$$s(a^1, \dots, a^{t-1}) = \begin{cases} (C, C) & \text{if } a^1 = \dots = a^{t-1} = (C, C), \\ (A, C) & \text{if } |1 \leq k < t : a^k = (C, A)| > |1 \leq k < t : a^k = (A, C)|, \\ (C, A) & \text{if } |1 \leq k < t : a^k = (C, A)| < |1 \leq k < t : a^k = (A, C)|, \\ (C, C) & \text{otherwise} \end{cases}$$

The players follow *getting-even* strategies. They choose C in the first stage and continue to do so provided that no player has played A more often than the other player; if in stages $1, \dots, t-1$ a player has played A more often than the other player, then he is “punished” in stage t with the less preferred stage game NE. Such a strategy has **infinite memory**, because the players need to keep count of how many times each player has played A in order to decide what to do.

The getting-even strategy profile is subgame perfect, and in order to verify this we can use the one shot deviation principle. Let h be a history where the players are supposed to choose (C, C) . This means the players are already even, so deviating only at h gives a continuation payoff equal to $(1 - \delta)(6 + \delta) + 4\delta^2$, which is no greater than 4 for $\delta \geq 2/3$. Next, pick any $t \geq 2$, any

$1 \leq m \leq t - 1$, and any history (a^1, \dots, a^{t-1}) such that, in this history, player i has played A more often (m times more) than player $-i$, i.e. such that

$$|1 \leq k < t : a_i^k = A| = |1 \leq k < t : a_{-i}^k = A| + m.$$

After such a history, player i has no profitable one shot deviation, since

$$(1 - \delta)(1 + \dots + \delta^{m-1}) + 4\delta^m > (1 - \delta)(-3 + \delta + \dots + \delta^m) + 4\delta^{m+1}.$$

Moreover, player $-i$ has no profitable one shot deviation, either, because

$$(1 - \delta)(6 + \dots + 6\delta^{m-1}) + 4\delta^m > (1 - \delta)(4 + 6\delta + \dots + 6\delta^m) + 4\delta^{m+1}.$$

5.3.2. Repeated Prisoners' Dilemma

The game *Prisoners' Dilemma* is depicted in Figure 16.

	D	C
D	1, 1	3, 0
C	0, 3	2, 2

Figure 16: Prisoners' Dilemma

The unique NE of this game (indeed the unique rationalizable strategy profile of this game) is (D, D) . The infinitely repeated game, however, has many subgame perfect equilibria. In particular, the outcome $((C, C), (C, C), \dots)$ is a SPE outcome for δ large enough. In the following strategy profile, each player adopts the *grim-trigger* strategy; at every stage, players cooperate if and only if they have both cooperated in the past: $s(\emptyset) = (C, C)$ and, for all $t \geq 2$,

$$s(a^1, \dots, a^{t-1}) = \begin{cases} (C, C) & \text{if } a^1 = \dots = a^{t-1} = (C, C), \\ (D, D) & \text{otherwise} \end{cases}$$

Checking that this profile is a SPE for δ large enough is quite easy, again because we can use the one-shot deviation principle. At histories where players are supposed to play (C, C) , no profitable deviation is available as long as

$$2 \geq (1 - \delta)3 + \delta$$

i.e. $\delta \geq 1/2$. At histories where players are supposed to play (D, D) , no profitable deviation is available, regardless of the value of δ . (Why?) We conclude that the grim-trigger strategy profile \star above is a SPE if and only if $\delta \geq 1/2$.

The grim-trigger strategy profile is **symmetric**: it prescribes the same action to both players, both on and off the equilibrium path. An example of an asymmetric, one-period-memory strategy profile is the following: $s(\emptyset) = (C, D)$ and, for all $t \geq 2$,

$$s(a^1, \dots, a^{t-1}) = \begin{cases} (D, C) & \text{if } a^{t-1} = (C, D), \\ (C, D) & \text{if } a^{t-1} = (D, C), \\ (D, D) & \text{otherwise} \end{cases}$$

If the players follow this profile, then player 1's payoff is $(1 - \delta)(0 + 3\delta + 0\delta^2 + 3\delta^3 + \dots) = 3\delta/(1 + \delta)$ and player 2's payoff is $(1 - \delta)(3 + 0\delta + 3\delta^2 + 0\delta^3 + \dots) = 3/(1 + \delta)$. At histories where players are supposed to play (C, D) , no profitable deviation is available to player 1 as long as

$$(1 - \delta)(0 + 3\delta + 0\delta^2 + 3\delta^3 + \dots) \geq (1 - \delta)(1 + 1\delta + 1\delta^2 + 1\delta^3 + \dots)$$

i.e. $\delta \geq 1/2$, and no profitable deviation is available to player 2 as long as

$$(1 - \delta)(3 + 0\delta + 3\delta^2 + 0\delta^3 + \dots) \geq (1 - \delta)(2 + 1\delta + 1\delta^2 + 1\delta^3 + \dots)$$

which is true for any δ . Symmetrically, at histories where players are supposed to play (D, C) , no profitable deviation is available to either player as long as $\delta \geq 1/2$. Finally, at histories where players are supposed to play (D, D) , no profitable deviation is available, provided that

$$1 \geq (1 - \delta)(0 + 3\delta + 0\delta^2 + 3\delta^3 + \dots)$$

i.e. $\delta \leq 1/2$. Thus, this strategy profile is a SPE if and only if $\delta = 1/2$.

6. Games with Incomplete Information

Most interesting strategic scenarios involve *asymmetric information*, i.e. different players having different opinions about the parameters of the game being played, or different opinions about each other's opinions, and so on. Let us begin with a simple example. Suppose player 1 is unsure whether the game he is playing with player 2 is the one on the left or the one on the right of Figure 17, while player 2 instead knows. Player 1's payoff does not depend on the true game that is being

	L	R		L	R
A	2, 1	0, 0	A	2, 0	0, 1
B	0, 1	2, 0	B	0, 0	2, 1
	ℓ			r	

Figure 17: An incomplete information situation.

played, even though the right action to choose would be A in the first case and B in the second. Indeed, player 2 knows the true game and will choose L in the first case and R in the second. So what should player 1 choose? In order to analyze the situation, it seems that we must first specify a belief for player 1 over the possible games that he thinks might be true, i.e. over the set $\{\ell, r\}$.

Now look again at Figure 17 and consider the opposite scenario, i.e. player 2 does not know which game is the true one, while player 1 knows. Player 2's payoff does not depend on player 1's choice; whether he should choose L or R depends on his belief over $\{\ell, r\}$. But, as before, player 1's payoff depends on what player 2 chooses, which in this case depends on which game player 2 believes is more likely to be the true one. Thus, in order to analyze player 1's choice, we must first specify a belief for player 1 over the possible beliefs over $\{\ell, r\}$ that player 1 thinks player 2 might have. Indeed, if player 1 believes that player 2 believes that ℓ is more likely, then he should choose A , whereas he should choose B otherwise.

Things can get even more complicated. Consider the two games in Figure 18. In this situation,

	A_2	B_2		A_2	B_2
A_1	1, 1	0, 0	A_1	0, 0	1, 1
B_1	0, 0	1, 1	B_1	1, 1	0, 0
	ℓ			r	

Figure 18: An incomplete information situation.

the players would like to choose the same action if ℓ is the true game, and choose different actions

if \mathcal{r} is the true game. (You can think of the first game as a situation where conformism is socially good, and of the second as a situation where it is bad.) Suppose that player 1 does not know which game is the true one, whereas player 2 knows. What should they do? This depends on each player's beliefs about ℓ and \mathcal{r} , on each player's beliefs about the other player's beliefs about ℓ and \mathcal{r} , on each player's beliefs about the other player's beliefs about the first player's beliefs, and so on. Indeed, player 1 should choose A_1 if sufficiently confident in either ℓ and A_2 or \mathcal{r} and B_2 . But whether A_2 or B_2 is chosen depends on whether player 2 believes ℓ and A_1 or \mathcal{r} and B_1 are more likely than ℓ and B_1 or \mathcal{r} and A_1 , and so on.

6.1. Bayesian Games

Harsanyi (Management Science 1967/68) showed how to describe and analyze potentially complicated asymmetric information situations like the ones above in simpler terms. Two main contributions are present in Harsanyi's papers. First, he introduced the notion of **type**; this allows us to replace the infinite sequence of each player's potentially relevant beliefs (about the parameters of the game, about the opponents' beliefs about the parameters of the game, and so on) with a *single* probability distribution over a well defined space of uncertainty. Second, he introduced the notion of **Bayesian game** and explained how to define and compute equilibria of these games using traditional game theoretic analysis.⁹

Definition 16. A **Bayesian game** is a list $\langle \Theta, N, (A_i, T_i, \pi_i, u_i)_{i \in N} \rangle$ comprising the following objects:

- a finite set of **states of nature** Θ
- a finite set of **players** N
- for each player i , a set of **actions** A_i (define A and A_{-i} as usual)
- for each player i , a set of **types** T_i (define $T := \times_{i \in N} T_i$ and $T_{-i} := \times_{j \in N \setminus \{i\}} T_j$)
- for each player i , a function $\pi_i : T_i \rightarrow \Delta(\Theta \times T_{-i})$
- for each player i , a **payoff** function $u_i : \Theta \times A \rightarrow \mathbb{R}$

⁹Actually, Harsanyi referred to the object defined in Definition 16 as an *incomplete information game*, and reserved the terminology *Bayesian game* for those games satisfying, in addition, the requirement of *consistency* given in Definition 17 below.

A **strategy** of player i is a function $s_i : T_i \rightarrow A_i$. The set of all strategies of player i is denoted S_i , while S and S_{-i} are defined as usual. A **Bayesian equilibrium** is a strategy profile $s \in S$ such that, for every i and every $t_i \in T_i$,

$$s_i(t_i) \in \arg \max_{a_i \in A_i} \sum_{\theta \in \Theta} \sum_{t_{-i} \in T_{-i}} \pi_i(t_i)[t_{-i}] u_i(\theta, a_i, s_{-i}(t_{-i})). \quad (15)$$

(Of course, you need integrals if the sets of types are not finite.)

Let us go back to the two examples based on Figure 17 and discussed above. In the first example, player 1 does not know what the true game is, while player 2 does. In order to fully describe what each player believes about the true game and about the opponent's beliefs, we add the following: player 1 believes ℓ and r are equally likely, player 2 knows this, player 1 knows that player 2 knows, and so on. This can be modeled as follows:

$$\Theta = \{\ell, r\} \quad N = \{1, 2\} \quad A_1 = \{A, B\} \quad A_2 = \{L, R\} \quad T_1 = \{t_1\} \quad T_2 = \{t'_2, t''_2\}$$

$$p_1(t_1)[\ell, t'_2] = p_1(t_1)[r, t''_2] = 1/2 \quad p_2(t'_2)[\ell, t_1] = 1 \quad p_2(t''_2)[r, t_1] = 1$$

$$u_1(\theta, a) = \begin{cases} 2 & \text{if } a = AL, \\ 2 & \text{if } a = BR, \\ 0 & \text{otherwise} \end{cases} \quad u_2(\theta, a) = \begin{cases} 1 & \text{if } a_2 = L \text{ and } \theta = \ell, \\ 1 & \text{if } a_2 = R \text{ and } \theta = r, \\ 0 & \text{otherwise} \end{cases}$$

Note that, for every player, we must specify as many types as his possible states of mind are, where by “state of mind” of a player we mean his beliefs about the game, his beliefs about the opponent's beliefs, and so on, and by “possible” we mean “possible from the point of view of his opponent”. In this example, player 2 has two possible states of mind (hence two types), namely, (I) the state of mind according to which (i) ℓ is the true game, (ii) player 1 thinks ℓ and r are equally likely, and (iii) both (i) and (ii) are known, known to be known, and so on; and (II) the state of mind according to which (i) r is the true game, (ii) player 1 thinks ℓ and r are equally likely, and (iii) both (i) and (ii) are known, known to be known, and so on. Player 1 has instead only one possible state of mind (hence only one type), namely, the state of mind according to which (i) ℓ and r are equally likely, (ii) player 2 knows the true game, and (iii) both (i) and (ii) are known, known to be known, and so

on.

Now consider the second example based on Figure 17, where player 1 knows and player 2 does not. Suppose player 2 thinks ℓ and κ are equally likely but that player 1, while knowing whether ℓ or κ is the true game, does not know player 2's beliefs. For example, assume that he believes with probability ε that player 2 is actually sure that κ is the true game. Assume that there is no further uncertainty i.e. assume that player 2 knows player 1 believes with probability ε that player 2 is sure κ is the true game, that player 1 knows that player 2 knows, and so on. We can model this as follows:

$$\Theta = \{\ell, \kappa\} \quad N = \{1, 2\} \quad A_1 = \{A, B\} \quad A_2 = \{L, R\} \quad T_1 = \{t'_1, t''_1\} \quad T_2 = \{t'_2, t''_2\}$$

$$p_1(t'_1)[\ell, t'_2] = 1 - \varepsilon \quad p_1(t'_1)[\ell, t''_2] = \varepsilon \quad p_1(t''_1)[\kappa, t'_2] = 1 - \varepsilon \quad p_1(t''_1)[\kappa, t''_2] = \varepsilon$$

$$p_2(t'_2)[\ell, t'_1] = p_2(t'_2)[\kappa, t'_1] = 1/2 \quad p_2(t''_2)[\kappa, t'_1] = 1$$

$$u_1(\theta, a) = \begin{cases} 2 & \text{if } a = AL, \\ 2 & \text{if } a = BR, \\ 0 & \text{otherwise} \end{cases} \quad u_2(\theta, a) = \begin{cases} 1 & \text{if } a_2 = L \text{ and } \theta = \ell, \\ 1 & \text{if } a_2 = R \text{ and } \theta = \kappa, \\ 0 & \text{otherwise} \end{cases}$$

Contrary to the first example, here even the player who is uninformed about the true game (that is, player 2) has more than one type. This is because player 1 knows the true game but does not know player 2's state of mind, being uncertain between two possible ones. Accordingly, we need two types for player 2.

Definition 17. A Bayesian game $\langle \Theta, N, (A_i, T_i, \pi_i, u_i)_{i \in N} \rangle$ is said to be **consistent** if there exists a probability distribution $q \in \Delta(\Theta \times T)$ such that, for every player i and every $t_i \in T_i$, one has

$$\pi(t_i)[\theta, t_{-i}] = \frac{q(\theta, t_i, t_{-i})}{\sum_{\theta' \in \Theta} \sum_{t'_{-i} \in T_{-i}} q(\theta', t_i, t'_{-i})} \quad (16)$$

whenever the denominator is nonzero. If such probability distribution q exists, it is called a **common prior**.

A consistent Bayesian game can be interpreted as a situation where Nature moves first, choosing a state of nature $\theta \in \Theta$ and a profile of types $t \in T$ according to some common prior q , and

then each player i learns her type and chooses an action $a_i \in A_i$. In this case we can define an associated normal form game just like the ones discussed at the beginning of the course, namely, the list $\langle N, (S_i, u_i^*)_{i \in N} \rangle$ where S_i is the set of functions from T_i into A_i and $u_i^* : S \rightarrow \mathbb{R}$ is the payoff function such that $u_i^*(s_1, \dots, s_I)$ is the expectation of $u_i(\theta, s_1(t_1), \dots, s_I(t_I))$ computed according to q .

Most models in economics assume that a common prior exists, i.e. they deal with consistent games only. In fact, in applications one only uses one of the associated normal form games. This is justified by Harsanyi's theorem below. Before stating and proving the theorem, consider once again the two examples based on Figure 17. It is immediate to see that in this case a common prior exists and is in fact unique: $q(\ell, t_1, t_2') = q(r, t_1, t_2'') = 1/2$. In the second example, however, no common prior exists. Indeed, suppose by contradiction that a common prior q exists. Then, since $\pi_2(t_2'')[\ell, t_1''] = 0$ and $\pi_2(t_2'')[r, t_1''] = 0$, we must have $q(\ell, t_1'', t_2'') = 0$ and $q(r, t_1'', t_2'') = 0$. (Why?) But then $q(\ell, t_1'', t_2')$ and $q(r, t_1'', t_2')$ should also be zero (if either were positive, then t_1'' ★ would have positive probability under q , so $\pi_1(t_1'')[\ell, t_2']$ and $\pi_1(t_1'')[r, t_2']$ should be zero, not ε). But then $\pi_2(t_2')[\ell, t_1']$ and $\pi_2(t_2')[r, t_1']$ should be zero, not $1/2$, a contradiction.

Theorem 6 (Harsanyi, 1967). *Let $\langle \Theta, N, (A_i, T_i, \pi_i, u_i)_{i \in N} \rangle$ be a consistent Bayesian game. Then every Bayesian equilibrium of this game is a Nash equilibrium of every associated normal form game. Moreover, if q is a common prior and $\langle N, (S_i, u_i^*)_{i \in N} \rangle$ is the corresponding associated normal form game, then every Nash equilibrium $s \in S$ of this game is q -almost surely a Bayesian equilibrium of the Bayesian game, in the sense that it satisfies (15) for every player i and every $t_i \in T_i$ occurring with positive probability under q .*

Proof. Let q be a common prior and let $s \in S$. Suppose s is a Bayesian equilibrium. Then it satisfies (15), and since q satisfies (16), taking expected values we see immediately that s is a Nash equilibrium of the associated normal form corresponding to q . Now suppose s is not a Bayesian equilibrium and that, in fact, it violates (15) for some i and some \hat{t}_i occurring with positive probability under q . Choose $\hat{a}_i \in A_i$ such that

$$\hat{a}_i \in \arg \max_{a_i \in A_i} \sum_{\theta \in \Theta} \sum_{t_{-i} \in T_{-i}} \pi_i(\hat{t}_i)[\theta, t_{-i}] u_i(\theta, a_i, s_{i-1}(t_{i-1})). \quad (17)$$

Define $s'_i : T_i \rightarrow A_i$ as $s'_i(t_i) = s_i(t_i)$ if $t_i \neq \hat{t}_i$ and $s'_i(\hat{t}_i) = \hat{a}_i$. Then it is easy to see that

$u_i^*(s'_i, s_{-i}) > u_i^*(s)$ and that, therefore, s is not a Nash equilibrium of the associated normal form corresponding to q . □