

Economics 2010a Game Theory Section Notes

Kevin He

December 1, 2015

Contents

	Page
1. Welcome to Game Theory.....	1
2. Nash Equilibrium and Its Extensions.....	9
3. Miscellaneous Topics; Bayesian Games.....	18
4. Auctions.....	29
5. Subgame Perfection and Bargaining.....	36
6. Repeated Games.....	44
7. Equilibrium Refinements and Signaling.....	52
8. References.....	60

"But I don't want to go among mad people," Alice remarked.

"Oh, you can't help that," said the Cat: "we're all mad here. I'm mad. You're mad."

"How do you know I'm mad?" said Alice.

"You must be," said the Cat, "or you wouldn't have come here."

— Alice in Wonderland, on mutual knowledge of irrationality

(1) Course outline; (2) Normal-form games; (3) Extensive-form games; (4) Extensive-form strategies

TF: Kevin He (he02@fas.harvard.edu)

1 Course Outline via a Taxonomy of Games

1.1 A taxonomy of games. The second half of Economics 2010a is organized around **several types of games**, paying particular attention to (i) relevant **solution concepts** in different settings, and (ii) some key **economic applications** belonging to these settings. To understand the course outline, it might be helpful to first introduce some binary classification schemes that give rise to these game types. Unfortunately, rigorous definitions of the following terminologies are not feasible without first laying down some background, so at this point we will instead appeal to hopefully familiar games to illustrate the classifications.

A game may have...

- Simultaneous moves (eg. rock-paper-scissors) or sequential moves (eg. checkers)
- Complete information (eg. chess) or incomplete information (eg. Hearthstone)
- Chance moves (eg. Backgammon) or no chance moves (eg. Reversi)
- Finite horizon (eg. tic-tac-toe) or infinite horizon (eg. Gomoku on an infinite board)
- Zero-sum payoff structure (eg. poker) or non-zero-sum payoff structure (eg. the usual model of prisoner's dilemma)

1.2 Course outline. Roughly, the course can be divided into 4 units. Each unit is focused on one type of game, studying first its solution concepts then some important examples and applications.

GAME TYPE 1: simultaneous move games with complete information

- theory: Nash equilibrium and its extensions, rationalizability
- application: Nash implementation

GAME TYPE 2: simultaneous move games with incomplete information

- theory: Bayesian Nash equilibrium
- application: Auctions

GAME TYPE 3: sequential move games with complete information

- theory: Subgame-perfect equilibrium

¹Figure 1 is from Haluk Ergin's game theory class at Berkeley (Economics 201A, Fall 2011), my first introduction to this field. Figures for Example 10 and Example 11 are adapted from Maschler, Solan, and Zamir (2013): Game Theory [1].

- application: Bargaining games, repeated games

GAME TYPE 4: **sequential move** games with **incomplete information**

- theory: Perfect Bayesian equilibrium, sequential equilibrium, trembling-hand perfect equilibrium, strategically stable equilibrium
- application: Reputation, signaling games

1.3 About sections. Sections are optional. We will review lecture material and work out some additional examples. Please interrupt to ask questions. The use of the plural first-person pronoun “we” in these section notes does not indicate royal lineage or pregnancy – rather, it suggests the notes form a conversation between the writer and the audience.

2 Normal-Form Games

2.1 Interpreting the payoff matrix. Here is the familiar **payoff matrix** representation of a two-player game.

	<i>L</i>	<i>R</i>
<i>T</i>	1,1	0,0
<i>B</i>	0,0	2,2

Player 1 (**P1**) chooses a row (**Top** or **Bottom**) while player 2 (**P2**) chooses a column (**Left** or **Right**). Each cell contains the payoffs to the two players when the corresponding pair of strategies is played. The first number in the cell is the payoff to P1 while the second number is the payoff to P2. (By the way, this game is sometimes called the “game of assurance”).

Two important things to keep in mind:

(1) In a normal-form game, players choose their strategies **simultaneously**. That is, P2 cannot observe which row P1 picks when choosing his column.

(2) The terminology “payoff matrix” is slightly misleading. The numbers that appear in a payoff matrix are actually **Bernoulli utilities**, not monetary payoffs. To spell out this point in painstaking details: the set of possible outcomes of the game is $X := \{TL, TR, BL, BR\}$. Each player j has a preference \succsim_j over $\Delta(X)$, the set of distributions on this 4 point set. Assume \succsim_j satisfies the vNM axioms of independence and continuity. Then, running \succsim_j through the vNM representation theorem, we find that \succsim_j is represented by a utility function $U_j : \Delta(X) \rightarrow \mathbb{R}$ with the functional form $U_j(\mathbf{p}) = p_{TL} \cdot u_j(TL) + p_{TR} \cdot u_j(TR) + p_{BL} \cdot u_j(BL) + p_{BR} \cdot u_j(BR)$. We then enter $u_j(TL), u_j(TR), u_j(BL), u_j(BR)$ into the payoff matrix cells, which happen to be 1, 0, 0, 2.

In particular, in computing the expected utility of each player under a mixed strategy profile, we simply take a weighted average of the matrix entries – there is no need to apply a “utility function” to the entries before taking the average as they are already denominated in utils. Furthermore, it is important to remember that this kind of linearity does not imply risk-neutrality of the players, but is rather a property of the vNM representation.²

²In fact, mixed strategies in game theory provided one of the motivations for von Neumann and Morgenstern’s work on their representation theorem for preference over lotteries. Von Neumann’s 1928 theorem on the equality between maxmin and minmax values in mixed strategies for zero-sum games assumed players choose the mixed strategy giving the highest expected value. But why should players choose between mixed strategies based on expected payoff rather than median payoff, mean payoff minus variance of payoff, or say the 4th moment of payoff? The vNM representation theorem rationalizes players maximizing expected payoff through a pair of conditions on their preference over lotteries.

2.2 General definition of a normal-form game. The payoff matrix representation of a game is convenient, but it is not sufficiently general. In particular, it seems unclear how we can represent games in which players have **infinitely** many possible strategies, such as a **Cournot duopoly**, in a finite payoff matrix. We therefore require a more general definition.

Definition 1. A **normal form game** $\mathcal{G} = \langle \mathcal{N}, (S_k)_{k \in \mathcal{N}}, (u_k)_{k \in \mathcal{N}} \rangle$ consists of:

- A finite collection of players³ $\mathcal{N} = \{1, 2, \dots, N\}$
- A set of pure strategies S_j for each $j \in \mathcal{N}$
- A (Bernoulli) utility function $u_j : \times_{k=1}^N S_k \rightarrow \mathbb{R}$ for each $j \in \mathcal{N}$

To interpret, the pure strategy set S_j is the set of actions that player j can take in the game. When each player chooses an action simultaneously from their own pure strategy set, we get a **strategy profile** $(s_1, s_2, \dots, s_N) \in \times_{k=1}^N S_k$. Players derive payoffs by applying their respective utility functions to the strategy profile.

The payoff matrix representation of a game is a **specialization** of this definition. In a payoff matrix for 2 players, the elements of S_1 and S_2 are written as the names of the rows and columns, while the values of u_1 and u_2 at different members of $S_1 \times S_2$ are written in the cells. If $S_1 = \{s_1^A, s_1^B\}$ and $S_2 = \{s_2^A, s_2^B\}$, then the game $\mathcal{G} = \langle \{1, 2\}, (A_1, A_2), (u_1, u_2) \rangle$ can be written in a payoff matrix:

	s_2^A	s_2^B
s_1^A	$u_1(s_1^A, s_2^A), u_2(s_1^A, s_2^A)$	$u_1(s_1^A, s_2^B), u_2(s_1^A, s_2^B)$
s_1^B	$u_1(s_1^B, s_2^A), u_2(s_1^B, s_2^A)$	$u_1(s_1^B, s_2^B), u_2(s_1^B, s_2^B)$

Conversely, the game of assurance can be converted into the standard definition by taking $\mathcal{N} = \{1, 2\}$, $S_1 = \{T, B\}$, $S_2 = \{L, R\}$, $u_1(T, L) = 1$, $u_1(B, R) = 2$, $u_1(T, R) = u_1(B, L) = 0$, $u_2(T, L) = 1$, $u_2(B, R) = 2$, $u_2(T, R) = u_2(B, L) = 0$,

The general definition allows us to write down games with infinite strategy sets. In a **duopoly** setting where firms choose own production quantity, their choices are not taken from a finite set of possible quantities, but are in principle allowed to be any positive real number. So, consider a game with $S_1 = S_2 = [0, \infty)$,

$$u_1(s_1, s_2) = p(s_1 + s_2) \cdot s_1 - C(s_1)$$

$$u_2(s_1, s_2) = p(s_1 + s_2) \cdot s_2 - C(s_2)$$

where $p(\cdot)$ and $C(\cdot)$ are inverse demand function and cost function, respectively. Interpreting s_1 and s_2 as the quantity choices of firm 1 and firm 2, this is Cournot competition phrased as a normal form game.

³By convention, players 1, 3, 5, ... are female while players 2, 4, 6, ... are male.

3 Extensive-Form Games

3.1 Definition of an extensive-form game. The rich framework of extensive-form games can incorporate sequential moves, incomplete and perhaps asymmetric information, randomization devices such as dice and coins, etc. It is one of the most powerful modeling tools of game theory, allowing researchers to formally study a wide range of economic interactions. Due to this richness, however, the general definition of an extensive-form game is somewhat cumbersome. Roughly speaking, an extensive-form game is a **tree** endowed with some **additional structures**. These additional structures formalize the **rules of the game**: the timing and order of play, the information of different players, randomization devices relevant to the game, outcomes and players' preferences over these outcomes, etc.

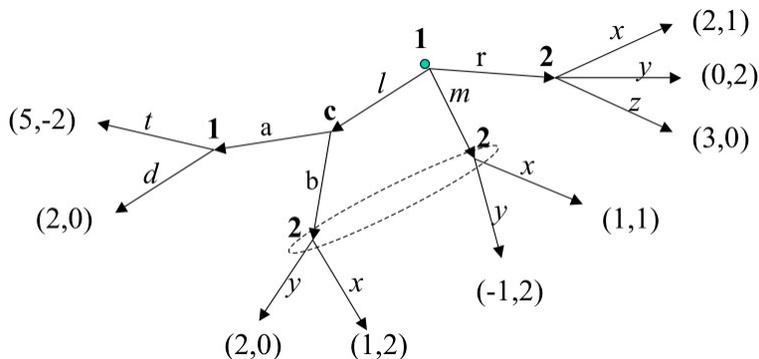
Definition 2. A finite-horizon extensive-form game Γ has the following components:

- A finite-depth **tree** with vertices V and terminal vertices $Z \subseteq V$.
- A set of **players** $\mathcal{N} = \{1, 2, \dots, N\}$.
- A **player function** $J : V \setminus Z \rightarrow \mathcal{N} \cup \{c\}$.
- A set of available **moves** $M_{j,v}$ for each $v \in J^{-1}(j)$, $j \in \mathcal{N}$. Each move in $M_{j,v}$ is associated with a unique child of v in the tree.
- A **probability distribution** $f(\cdot|v)$ over v 's children for each $v \in J^{-1}(c)$.
- A (Bernoulli) **utility function** $u_j : Z \rightarrow \mathbb{R}$ for each $j \in \mathcal{N}$.
- An **information partition** \mathcal{I}_j of $J^{-1}(j)$ for each $j \in \mathcal{N}$, whose elements are **information sets** $I_j \in \mathcal{I}_j$. It is required that $v, v' \in I_j \Rightarrow M_{j,v} = M_{j,v'}$.

The game tree captures all possible states of the game. When players reach a terminal vertex $z \in Z$ of the game tree, the game ends and each player j receives utility $u_j(z)$. The player function J indicates who moves at each non-terminal vertex. The move might belong to an actual player $j \in \mathcal{N}$, or to **chance**, “ c ”. Note that $J^{-1}(j)$ refers to the set of all vertices where player j has the move. If a player j moves at vertex v , she gets to pick an element from the set $M_{j,v}$ and play proceeds along the corresponding edge. If chance moves, then play proceeds along a random edge chosen according to $f(\cdot|v)$.

An information set I_j of player j refers to a set of vertices that player j **cannot distinguish between**.⁴ It might be useful to imagine the players conducting the game in a lab, mediated by

⁴The use of an information partition to model a player's knowledge predates extensive-form games with incomplete information. Such models usually specify a set of states of the world, Ω , then some partition \mathcal{I}_j on Ω . When a state of the world $\omega \in \Omega$ realizes, player j is told the information set $I_j \in \mathcal{I}_j$ containing ω , but not the exact identity of ω . So then, **finer partitions** correspond to **better information**. To take an example, suppose 3 people are standing facing the same direction. An observer places a hat of one of two colors (say color 0 and color 1) on each of the 3 people. These 3 people cannot see their own hat color or the hat color of those standing behind them. Then the states of the world are $\Omega = \{000, 001, 010, 011, 100, 101, 110, 111\}$. The person in the front of the line has no information, so her information partition contains just one information set with all the states, $\mathcal{I}_1 = \{ \{000, 001, 010, 011, 100, 101, 110, 111\} \}$. The second person in line sees only the hat color of the first person, so that $\mathcal{I}_2 = \{ \{000, 010, 100, 110\}, \{001, 011, 101, 111\} \}$. Finally, the last person sees the hats of persons 1 and 2, so that $\mathcal{I}_3 = \{ \{000, 100\}, \{001, 101\}, \{010, 110\}, \{011, 111\} \}$. In the context extensive form games, one might think of $J^{-1}(j)$ as the relevant “states of the world” for j 's decision-making and the fineness of her information partition \mathcal{I}_j reflects the extent to which she can distinguish between these states.



Chance move distributions: $f(a|l) = f(b|l) = 0.5$.
Information partitions: $\mathcal{I}_1 = \{ \emptyset, (l, a) \}$, $\mathcal{I}_2 = \{ (l, r), \{(l, b), (m)\} \}$.

Figure 1: An extensive-form game with incomplete information and chance moves.

a computer. At each vertex $v \in V \setminus Z$, the computer finds the player $J(v)$ who has the move and informs her that the game has arrived at the information set $I_{J(v)} \ni v$. In the event that this $I_{J(v)}$ is a singleton, player $J(v)$ knows exactly her location in the game tree. Else, she knows only that she is at one of the vertices in $I_{J(v)}$, but she does not know for sure which one.⁵ The requirement that two vertices in the same information set must have the same sets of moves is to prevent a player from gaining additional information by simply examining the set of moves available to her, which would defeat the idea that the player supposedly cannot distinguish between any of the vertices in the same information set. For convenience, we also write M_{j,I_j} for the common move set for all vertices $v \in I_j$. There are two conventions for indicating an information set I_j in a game tree diagrams. Either all of the vertices in I_j are connected using dashed lines, or all of the vertices are encircled in an oval.

Example 3. Figure 1 illustrates all the pieces of the general definition of an extensive-form game. For convenience, let's name each vertex with the sequence of moves leading to it (and name the root as \emptyset). The set of players is $\mathcal{N} = \{1, 2\}$. The player function $J(v)$ is shown on each $v \in V \setminus Z$ in Figure 1, while the payoff pair $(u_1(z), u_2(z))$ is shown on each $z \in Z$. The set of moves $M_{j,v}$ at vertex v is shown on the corresponding edges. P1 moves at two vertices, $J^{-1}(1) = \{ \emptyset, (l, a) \}$. Her information partition contains only singleton sets, meaning she always knows where in the game tree she is when called upon to move. P2 moves at three vertices, $J^{-1}(2) = \{ (r), (l, b), (m) \}$. However, P2 cannot distinguish between (m) and (l, b) , though he can distinguish (r) from the other two vertices. As such, his information partition contains two information sets, one containing just (r) , the other containing the two vertices (m) and (l, b) . As required by the definition, $M_{2,(m)} = M_{2,(l,b)} = \{x, y\}$, so that P2 cannot figure out whether he is at (m) or (l, b) by looking at the set of available moves. \blacklozenge

3.2 Using information sets to convert a normal-form game into extensive-form. Every finite normal form game $\mathcal{G} = \langle \mathcal{N}, (S_j)_{j=1}^N, (u_j)_{j=1}^N \rangle$ may be **converted** into an extensive form game of incomplete information with $1 + \sum_{L=1}^N \prod_{j=1}^L |S_j|$ vertices. Construct a game tree with $N + 1$ levels, so that all the vertices at level L belong to a single information set for player L , for $1 \leq L \leq N$. Level 1 contains the root. The root has $|S_1|$ children, corresponding to the actions in S_1 . These children form the level 2 vertices. Each of these level 2 vertices has $|S_2|$ children, corresponding to the actions in S_2 , and so forth. Each terminal vertex z in level $N + 1$ corresponds to some action profile $(s_1^z, s_2^z, \dots, s_N^z)$

⁵She might, however, be able to form a belief as to the likelihood of being at each vertex in $I_{J(v)}$, based on her knowledge of other players' strategies and the chance move distributions.

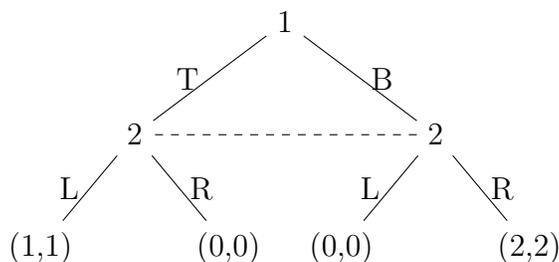


Figure 2: Game of assurance in extensive form.

in the normal form game \mathcal{G} and is assigned utility $u_j(s_1^z, s_2^z, \dots, s_N^z)$ for player j in the extensive form game. Figure 2 illustrates such a conversion using the game of assurance discussed earlier.

4 Strategies in Extensive-Form Games

4.1 Pure strategy in extensive form games. How would you write a program to play an extensive-form game as player j ? Whenever it is player j 's turn, the program should take the **information set as an input** and return **one of the feasible moves as an output**. As the programmer does not *a priori* know the strategies that other players will use, the program must encode a **complete contingency plan** for playing the game so that it returns a legal move at every vertex of the game tree where j might be called upon to play. This motivates the definition of a pure strategy in an extensive-form game.

Definition 4. In an extensive-form game, a **pure strategy** for player j is a function $s_j : \mathcal{I}_j \rightarrow \cup_{v \in J^{-1}(j)} M_{j,v}$, so that $s_j(I_j) \in M_{j,I_j}$ for each $I_j \in \mathcal{I}_j$. Write S_j for the set of all pure strategies of player j .

That is, a pure strategy for player j returns a legal move at every information set of j .

Example 5. In Figure 1, one of the strategies of P1 is $s_1(\emptyset) = m$, $s_1(l, a) = d$. Even though playing m at the root means the vertex (l, a) will never be reached, P1's strategy must still specify what she **would have done** at (l, a) . This is because some solution concepts we will study later in the course will require us to examine parts of the game tree which are **unreached** when the game is played. One of the strategies of P2 is $s_2(\{(l, b), (m)\}) = y$, $s_2(r) = z$. In every pure strategy P2 must play the same action at both (l, b) and (m) , as pure strategies are functions of information sets, not individual vertices. In total, P1 has 6 different pure strategies in the game and P2 has 6 different pure strategies. ◆

4.2 Two definitions of randomization. There are at least two natural notions of "randomizing" in an extensive-form game: (1) Player j could enumerate the set of all possible pure strategies, S_j , then choose an element of S_j at random; (2) Player j could pick a randomization over M_{j,I_j} for each of her information sets $I_j \in \mathcal{I}_j$. These two notions of randomization lead to two different classes of strategies that incorporate stochastic elements:

Definition 6. A **mixed strategy** for player j is an element $\sigma_j \in \Delta(S_j)$.

Definition 7. A **behavioral strategy** for player j is a collection of distributions $\{b_{I_j}\}_{I_j \in \mathcal{I}_j}$, where $b_{I_j} \in \Delta(M_{j,I_j})$.

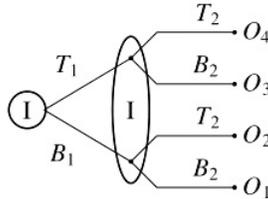
Strictly speaking, mixed strategies and behavioral strategies form two **distinct classes of objects**. We may, however, talk about the equivalence between a mixed strategy and a behavioral strategy in the following way:

Definition 8. Say a mixed strategy σ_j and a behavioral strategy $\{b_{I_j}\}$ are **equivalent** if they generate the same distribution over terminal vertices regardless of the strategies used by opponents, which may be mixed or behavioral.

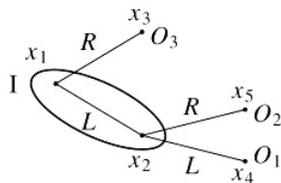
Example 9. In Figure 1, a behavioral strategy for P1 is: $b_\emptyset^*(l) = 0.5, b_\emptyset^*(m) = 0, b_\emptyset^*(r) = 0.5, b_{(l,a)}^*(t) = 0.7, b_{(l,a)}^*(d) = 0.3$. That is, P1 decides that she will play m and r each with 50% probability at the root of the game. If she ever reaches the vertex (l, a) , she will play t with 70% probability, d with 30% probability. But now, consider the following 4 pure strategies: $s_1^{(1)}(\emptyset) = l, s_1^{(1)}(l, a) = t$; $s_1^{(2)}(\emptyset) = l, s_1^{(2)}(l, a) = d$; $s_1^{(3)}(\emptyset) = r, s_1^{(3)}(l, a) = t$; $s_1^{(4)}(\emptyset) = r, s_1^{(4)}(l, a) = d$ and construct the mixed strategy σ_j^* so that $\sigma_j^*(s_1^{(1)}) = 0.35, \sigma_j^*(s_1^{(2)}) = 0.15, \sigma_j^*(s_1^{(3)}) = 0.35, \sigma_j^*(s_1^{(4)}) = 0.15$. Then the behavioral strategy b^* is equivalent to the mixed strategy σ_j^* . \blacklozenge

It is often “nicer” to work with behavioral strategies than mixed strategies, for at least two reasons. One, behavioral strategies are **easier to write down** and usually involve fewer parameters than mixed strategies. Two, it feels **more natural** for a player to randomize at each decision node than to choose a “grand plan” at the start of the game. In general, however, neither the set of mixed strategies nor the set of behavioral strategies is a “subset” of the other, as we now demonstrate.

Example 10. (A mixed strategy without an equivalent behavioral strategy) Consider an **absent-minded city driver** who must make turns at two consecutive intersections. Upon encountering the second intersection, however, she does not remember whether she turned left (T) or right (B) at the first intersection. The mixed strategy σ_1 putting probability 50% on each of the two pure strategies T_1T_2 and B_1B_2 generates the outcome O_1 50% of the time and the outcome O_4 50% of the time. However, this outcome distribution cannot be obtained using any behavioral strategy. That is, if the driver chooses some probability of turning left at the first intersection and some probability of turning left at the second intersection, and furthermore these two randomizations are independent, then she can never generate the outcome distribution of 50% O_1 , 50% O_4 .



Example 11. (A behavioral strategy without an equivalent mixed strategy) Consider an **absent-minded highway driver** who wants to take the second highway exit. Starting from the root of the tree, x_1 , he wants to keep left (L) at the first highway exit but keep right (R) at the second highway exit. Upon encountering each highway exit, however, he does not remember if he has already encountered an exit before. The driver has only two pure strategies: always L or always R. It is easy to see no mixed strategy can ever achieve the outcome O_2 . However, the behavioral strategy of taking L and R each with 50% probability each time he arrives at his information set gets the outcome O_2 with 25% probability. \blacklozenge



These two examples are pathological in the sense that the drivers “forget” some information that they knew before. The city driver forgets what action she took at the previous information set. The highway driver forgets what information sets he has encountered. The definition of perfect recall rules out these two pathologies.

Definition 12. An extensive-form game has **perfect recall** if whenever $v, v' \in I_j$, the two paths leading from the root to v and v' pass through the same sequence of information sets and take the same sequence of actions at these information sets.

In particular, the city driver game fails perfect recall since taking two different actions from the root vertex lead to two vertices in the same information set. The highway driver game fails perfect recall since vertices x_1 and x_2 are in the same information set, yet the path from root to x_1 is empty while the path from root to x_2 passes through one information set.

Kuhn’s theorem states that in a game with perfect recall, it is without loss to analyze only behavioral strategies. Its proof is beyond the scope of this course.

Theorem 13. (Kuhn 1957) *In a finite extensive-form game with perfect recall, (i) every mixed strategy has an equivalent behavioral strategy, and (ii) every behavioral strategy has an equivalent mixed strategy.*

(1) Strat. in normal-form games; (2) Nash eqm; (3) Solving for NE; (4) Extensions of NE

TF: Kevin He (he02@fas.harvard.edu)

1 Strategies in Normal Form Games

1.1 Recurring notations. The following notations are common in game theory but usually go unexplained. If X_1, X_2, \dots, X_N are sets with typical elements $x_1 \in X_1, x_2 \in X_2, \dots$, then:

- X_{-i} means $\times_{1 \leq k \leq N, k \neq i} X_k$
- X is sometimes understood to mean $\times_{k=1}^N X_k$.
- (x_i) refers to a vector⁶ (x_1, x_2, \dots, x_N) . So (x_i) is an element in $\times_{k=1}^N X_k$. The parentheses are used to distinguish it from x_i , which is an element of X_i .
- x_{-i} is an element in X_{-i} , i.e. $\times_{1 \leq k \leq N, k \neq i} X_k$. Confusingly, usually no parentheses are used around x_{-i} .

To see an example of these notations, suppose we are studying a three player game

$$\mathcal{G} = \langle \{1, 2, 3\}, (S_1, S_2, S_3), (u_1, u_2, u_3) \rangle$$

Then “ s_{-2} ” usually refers to a vector containing strategies from P1 and P3, but not P2. It is an element of $S_1 \times S_3$, also written as S_{-2} .

1.2 Mixed strategies in normal-form games. A player who uses a mixed strategy in a game **intentionally introduces randomness** into her play. Instead of picking a deterministic action as in a pure strategy, a mixed strategy user tosses a coin to determine what action to play. Game theorists are interested in mixed strategies for at least two reasons: (i) mixed strategies correspond to **how humans play certain games**, such as rock-paper-scissors; (ii) the space of mixed strategies represents a “**convexification**” of the action set S_i and convexity is required for many existence results.

Definition 14. Suppose $\mathcal{G} = \langle \mathcal{N}, (S_k)_{k \in \mathcal{N}}, (u_k)_{k \in \mathcal{N}} \rangle$ is a normal-form game where each S_k is finite⁷. Then a **mixed strategy** σ_i is a member of $\Delta(S_i)$.

Sometimes the mixed strategy putting probability p_1 on action $s_1^{(1)}$ and probability $1 - p_1$ on action $s_1^{(2)}$ is written as $p_1 s_1^{(1)} \oplus (1 - p_1) s_1^{(2)}$. The “ \oplus ” notation (in lieu of “ $+$ ”) is especially useful when $s_1^{(1)}, s_1^{(2)}$ are real numbers, as to avoid confusing the mixed strategy with an arithmetic expression to be simplified.

Two remarks:

- When two or more players play mixed strategies, their randomizations are assumed to be **independent**.

⁶Sometimes also called a “profile”.

⁷We can also define mixed strategies when the set of actions S_k is infinite. However, we would need to first equip S_k with a sigma-algebra, then define the mixed strategy as a measure on this sigma-algebra.

- Technically, **pure strategies also count as mixed strategies** – they are simply degenerate distributions on the action set. The term “**strictly mixed**” is usually used for a mixed strategy that puts strictly positive probability on every action.

When a profile of mixed strategies $(\sigma_k)_{k=1}^N$ is played, the assumption on independent mixing, together with previous week’s discussion on payoff matrix entries as Bernoulli utilities in a vNM representation, implies that player i gets utility:

$$\sum_{(s'_k) \in \times_k S_k} u_i(s'_1, \dots, s'_N) \cdot \sigma_1(s'_1) \cdot \dots \cdot \sigma_N(s'_N)$$

We will abuse notation and write $u_i(\sigma_i, \sigma_{-i})$ for this utility, extending the domain of u_i into mixed strategies.

1.3 What does it mean to “solve” a game? A detour into combinatorial game theory. Why are economists interested in Nash equilibrium, or solution concepts in general? As a slight aside, you may want to know that there actually exist two areas of research that go by the name of “game theory”. The full names of these two areas are “**combinatorial game theory**” and “**equilibrium game theory**”. Despite the similarity in name, these two versions of game theory have quite different research agendas. The most salient difference is that combinatorial game theory studies well-known board games like chess where there exists (theoretically) a “winning strategy” for one player. Combinatorial game theorists aim to find these winning strategies, thereby solving the game. On the other hand, no “winning strategies” (usually called “**dominant strategies**” in our lingo) exist for most games studied by equilibrium game theorists⁸. In the Battle of the Sexes, for example, due to the simultaneous-move condition, there is no one strategy that is optimal for P1 regardless of how P2 plays, in contrast to the existence of such optimal strategies in, say, tic-tac-toe.

If a game has a dominant strategy for one of the players, then it is straight-forward to predict its outcome under optimal play. The player with the dominant strategy will employ this strategy and the other player will do the best they can to minimize their losses. However, predicting outcome in a game without dominant strategies requires the analyst to make assumptions. These assumptions are usually called **equilibrium assumptions** and give “equilibrium game theory” its name. One of the most common equilibrium assumptions in normal-form games with complete information is the Nash equilibrium, which we now study.

2 Nash Equilibrium

2.1 Defining Nash equilibrium. A Nash equilibrium⁹ is a strategy profile where no player can improve upon her own payoff through a **unilateral** deviation, taking as given the actions of others. This leads to the usual definition of pure and mixed Nash equilibria.

Definition 15. In a normal-form game $\mathcal{G} = \langle \mathcal{N}, (S_k)_{k \in \mathcal{N}}, (u_k)_{k \in \mathcal{N}} \rangle$, a **Nash equilibrium in pure strategies** is a pure strategy profile $(s_k^*)_{k \in \mathcal{N}}$ such that for every player i , $u_i(s_i^*, s_{-i}^*) \geq u_i(s'_i, s_{-i}^*)$ for all $s'_i \in S_i$.

Definition 16. In a normal-form game $\mathcal{G} = \langle \mathcal{N}, (S_k)_{k \in \mathcal{N}}, (u_k)_{k \in \mathcal{N}} \rangle$, a **Nash equilibrium in mixed strategies** is a mixed strategy profile $(\sigma_k^*)_{k \in \mathcal{N}}$ such that for every player i , $u_i(\sigma_i^*, \sigma_{-i}^*) \geq u_i(s'_i, \sigma_{-i}^*)$ for all $s'_i \in S_i$.

⁸The one-shot prisoner’s dilemma is an exception here.

⁹John Nash called this equilibrium concept “equilibrium point” but later researchers referred to it as “Nash equilibrium”. We will see a similar situation next week.

In the definition of a mixed Nash equilibrium, we required no profitable unilateral deviation to any **pure** strategy, s'_i . It would be equivalent to require no profitable unilateral deviation to any **mixed** strategy, due to the following fact.

Fact 17. For any fixed σ_{-i} , the map $\sigma_i \mapsto u_i(\sigma_i, \sigma_{-i})$ is affine, in the sense that $u_i(\sigma_i, \sigma_{-i}) = \sum_{s_i \in S_i} \sigma_i(s_i) \cdot u_i(s_i, \sigma_{-i})$.

That is, the payoff to playing σ_i against opponents' mixed strategy profile σ_{-i} is some **weighted average** of the $|S_i|$ numbers $(u_i(s_i, \sigma_{-i}))_{s_i \in S_i}$, where the weights are given by the probabilities that σ_i assigns to these different actions. So, if there is some profitable mixed strategy deviation σ'_i from a strategy profile $(\sigma_i^*, \sigma_{-i}^*)$, then it must be the case that for at least one $s'_i \in S_i$ with $\sigma'_i(s'_i) > 0$, we have $u_i(s'_i, \sigma_{-i}^*) > u_i(\sigma_i^*, \sigma_{-i}^*)$.

Example 18. Consider the game of assurance,

	L	R
T	1,1	0,0
B	0,0	2,2

We readily verify that both (T, L) and (B, R) are pure-strategy Nash equilibria. Note one of these two NEs Pareto dominates the other. In general, NEs **need not be Pareto efficient**. This is because the definition of NE only accounts for the absence of profitable **unilateral** deviations. Indeed, starting from the strategy profile (T, L) , if P1 and P2 can contract on simultaneously changing their strategies, then they would both be better off. However, these sorts of simultaneous deviations by a “coalition” are not allowed.

But wait, there's more! Suppose P1 plays $\frac{2}{3}T \oplus \frac{1}{3}B$. Suppose P2 plays $\frac{2}{3}L \oplus \frac{1}{3}R$. This strategy profile is also a mixed NE. When P1 is playing $\frac{2}{3}T \oplus \frac{1}{3}B$, P2 gets an expected payoff of $\frac{2}{3}$ from playing L and an expected payoff of $\frac{2}{3}$ from playing R . Therefore, P2 has no profitable unilateral deviation because every strategy he could play, pure or mixed, would give the same payoff of $\frac{2}{3}$. Similarly, P2's mixed strategy $\frac{2}{3}L \oplus \frac{1}{3}R$ means P1 gets an expected payoff of $\frac{2}{3}$ whether she plays T or B , so P1 does not have a profitable deviation either. \blacklozenge

2.2 Nash equilibrium as a fixed-point of the best response correspondence. Nash equilibrium embodies the idea of stability. To make this point clear, it is useful to introduce an equivalent view of the Nash equilibrium through the lens of best response correspondence.

Definition 19. The **individual pure best-response correspondence** for player i is $BR_i : S_{-i} \rightrightarrows S_i$ ¹⁰ where

$$BR_i(s_{-i}) := \arg \max_{\hat{s}_i \in S_i} u_i(\hat{s}_i, s_{-i})$$

The **pure best-response correspondence** involves putting the N pure best-response correspondences into a vector: $BR : S \rightrightarrows S$ where $BR(s) := (BR_1(s_{-1}) \dots BR_N(s_{-N}))$.

Analogously, the **individual mixed best-response correspondence** for player i is $\overline{BR}_i : \prod_{k \neq i} \Delta(S_k) \rightrightarrows \Delta(S_i)$ where

$$\overline{BR}_i(\sigma_{-i}) := \arg \max_{\hat{\sigma}_i \in \Delta(S_i)} u_i(\hat{\sigma}_i, \sigma_{-i})$$

The **mixed best-response correspondence** involves putting the N mixed best-response correspondences into a vector: $\overline{BR} : S \rightrightarrows S$ where $\overline{BR}(\sigma) := (\overline{BR}_1(\sigma_{-1}) \dots \overline{BR}_N(\sigma_{-N}))$.

¹⁰The notation $f : A \rightrightarrows B$ is equivalent to $f : A \rightarrow 2^B$.

To interpret, the individual best-response correspondences return the argmax of each player's utility function when opponents play some known strategy profile. Depending on others' strategies, the player may have multiple maximizers, all yielding the same utility. As a result, we must allow the best responses to be **correspondences** rather than functions. Then, it is easy to see that:

Proposition 20. *A pure strategy profile is a pure NE iff it is a fixed point of BR. A mixed strategy profile is a mixed NE iff it is a fixed point of \overline{BR} .*

Fixed points of the best response correspondences reflect **stability** of NE strategy profiles, in the sense that even if player i knew what others were going to play, she still would not find it beneficial to change her actions. This rules out cases where a player plays in a certain way only because she held the **wrong expectations** about other players' strategies. We might expect such outcomes to arise initially when inexperienced players participate in the game, but we would also expect such outcomes to vanish as players learn to adjust their strategies to maximize their payoffs over time. That is to say, we expect non-NE strategy profiles to be unstable.

2.3 Important properties of NE. Here are two important properties for computing NEs:

Property 1: Iterated elimination of strictly dominated strategies does not change the set of NEs. In a game $\mathcal{G}^{(1)}$, call a strategy $s_i^\circ \in S_i$ **strictly dominated** if there exists some mixed strategy $\sigma_i \in \Delta(S_i)$ such that $u_i(\sigma_i, s_{-i}) > u_i(s_i^\circ, s_{-i})$ for **every** $s_{-i} \in S_{-i}$. We can remove some or all of each player's strictly dominated strategies to arrive at a new game $\mathcal{G}^{(2)}$, which will always have the same set of NEs as $\mathcal{G}^{(1)}$. Furthermore, this procedure can be repeated, removing some of each player's strictly dominated strategies in $\mathcal{G}^{(t)}$ to arrive at $\mathcal{G}^{(t+1)}$. All of the games $\mathcal{G}^{(1)}, \mathcal{G}^{(2)}, \mathcal{G}^{(3)}, \dots$ will have the same set of NEs, but solving for the NEs of the later games is probably easier than solving the NEs of $\mathcal{G}^{(1)}$.

Property 2, the indifference condition in mixed NEs. In Example 18, we saw that each action that player i plays with strictly positive probability yields the same expected payoff against the mixed strategy profile of the opponent. Turns out this is a general phenomenon.

Proposition 21. *Suppose (σ_i^*) is a mixed Nash equilibrium. Then for any $s_i \in S_i$ such that $\sigma_i^*(s_i) > 0$, we have $u_i(s_i, \sigma_{-i}^*) = u_i(\sigma_i^*, \sigma_{-i}^*)$.*

Proof. Suppose we may find $s_i \in S_i$ so that $\sigma_i^*(s_i) > 0$ but $u_i(s_i, \sigma_{-i}^*) \neq u_i(\sigma_i^*, \sigma_{-i}^*)$. In the event that $u_i(s_i, \sigma_{-i}^*) > u_i(\sigma_i^*, \sigma_{-i}^*)$, we contradict the optimality of σ_i^* in the maximization problem $\arg \max u_i(\hat{\sigma}_i, \sigma_{-i}^*)$, for we should have just picked $\hat{\sigma}_i = s_i$, the degenerate distribution on pure strategy s_i . In the event that $u_i(s_i, \sigma_{-i}^*) < u_i(\sigma_i^*, \sigma_{-i}^*)$, we enumerate $S_i = \{s_i^{(1)}, \dots, s_i^{(r)}\}$ and use the Fact 17 to expand:

$$u_i(\sigma_i^*, \sigma_{-i}^*) = \sum_{k=1}^r \sigma_i^*(s_i^{(k)}) \cdot u_i(s_i^{(k)}, \sigma_{-i}^*)$$

The term $u_i(s_i, \sigma_{-i}^*)$ appears in the summation on the right with a strictly positive weight, so if $u_i(s_i, \sigma_{-i}^*) < u_i(\sigma_i^*, \sigma_{-i}^*)$ then there must exist another $s_i' \in S_i$ such that $u_i(s_i', \sigma_{-i}^*) > u_i(\sigma_i^*, \sigma_{-i}^*)$. But now we have again contradicted the fact that σ_i^* is a mixed best response to σ_{-i}^* . \square

3 Solving for Nash Equilibria

The following steps may be helpful in solving for NEs of two-player games.

1. Use **iterated elimination** of strictly dominated strategies to simplify the problem.
2. Find all the **pure-strategy Nash equilibria** by considering all cells in the payoff matrix.
3. Look for a **mixed** Nash equilibrium where one player is playing a pure strategy while the other is strictly mixing.
4. Look for a **mixed** Nash equilibrium where **both** players are strictly mixing.

Example 22. (December 2013 Final Exam) Find all NEs, pure and mixed, in the following payoff matrix.

	L	R	Y
T	2, 2	-1, 2	0, 0
B	-1, -1	0, 1	1, -2
X	0, 0	-2, 1	0, 2

Solution:

Step 1: Strategy X for P1 is strictly dominated by $\frac{1}{2}T \oplus \frac{1}{2}B$. Indeed, $u_1(X, L) = 0 < 0.5 = u_1(\frac{1}{2}T \oplus \frac{1}{2}B, L)$, $u_1(X, R) = -2 < -0.5 = u_1(\frac{1}{2}T \oplus \frac{1}{2}B, R)$, and $u_1(X, Y) = 0 < 0.5 = u_1(\frac{1}{2}T \oplus \frac{1}{2}B, Y)$. But having eliminated X for P1, strategy Y for P2 is strictly dominated by R : $u_2(T, Y) = 0 < 2 = u_2(T, R)$, $u_2(B, Y) = -2 < 1 = u_2(B, R)$. Hence we can restrict attention to the smaller, 2x2 game in the upper left corner.

Step 2: (T, L) is a pure Nash equilibrium as no player has a profitable unilateral deviation. (The deviation $L \rightarrow R$ does not **strictly** improve the payoff of P2, so it doesn't break the equilibrium.) At (T, R) , P1 deviates $T \rightarrow B$, so it is not a pure strategy Nash equilibrium. At (B, L) , P2 deviates $L \rightarrow R$. At (B, R) , no player has a profitable unilateral deviation, so it is a pure-strategy Nash equilibrium. In summary, the game has two pure-strategy Nash equilibria: (T, L) and (B, R) .

Step 3: Now we look for mixed Nash equilibria where one player is using a pure strategy while the other is using a strictly mixed strategy. As discussed before, if a player strictly mixes between two pure strategies, then they must be getting the **same payoff** from playing either of these two pure strategies.

Using this **indifference condition**, we quickly realize it cannot be the case that P2 is playing a pure strategy while P1 strictly mixes. Indeed, if P2 plays L then $u_1(T, L) > u_1(B, L)$. If P2 plays R then $u_1(B, R) > u_1(T, R)$.

Similarly, if P1 is playing B , then the indifference condition cannot be sustained for P2 since $u_2(R, B) > u_2(L, B)$.

Now suppose P1 plays T . Then $u_2(T, L) = u_2(T, R)$. This indifference condition ensures that any strictly mixed strategy of P2 $pL \oplus (1 - p)R$ for $p \in (0, 1)$ is a mixed best response to P1's strategy. However, to ensure this is a mixed Nash equilibrium, we must **also make check** P1 does not have any profitable unilateral deviation. This requires:

$$u_1(T, pL \oplus (1 - p)R) \geq u_1(B, pL \oplus (1 - p)R)$$

that is to say,

$$\begin{aligned}
2p + (-1) \cdot (1 - p) &\geq (-1) \cdot p + 0 \cdot (1 - p) \\
4p &\geq 1 \\
p &\geq \frac{1}{4}
\end{aligned}$$

Therefore, $(T, pL \oplus (1 - p)R)$ is a mixed Nash equilibrium where P2 strictly mixes when $p \in [\frac{1}{4}, 1)$.

Step 4: There are no mixed Nash equilibria where both players are strictly mixing. To see this, notice if $\sigma_1^*(B) > 0$, then

$$u_2(\sigma_1^*, L) = 2 \cdot (1 - \sigma_1^*(B)) + (-1) \cdot (\sigma_1^*(B)) < 2 \cdot (1 - \sigma_1^*(B)) + (1) \cdot (\sigma_1^*(B)) = u_2(\sigma_1^*, R)$$

So it cannot be the case that P2 is also strictly mixing, since P2 is not indifferent between L and R . In total, the game has two pure Nash equilibria, (T, L) and (B, R) , as well as infinitely many mixed Nash equilibria, $(T, pL \oplus (1 - p)R)$ for $p \in [\frac{1}{4}, 1)$. \blacklozenge

Sometimes, iterated elimination of strictly dominated strategy simplifies the game so much that the solution is immediate after **Step 1**. The following example illustrates.

Example 23. (Guess two-thirds the average, also sometimes called the beauty-contest game¹¹) Consider a game of 2 players $\mathcal{G}^{(1)}$ where $S_1 = S_2 = [0, 100]$, $u_i(s_i, s_{-i}) = -\left(s_i - \frac{2}{3} \cdot \frac{s_i + s_{-i}}{2}\right)^2$. That is, each player wants to play an action as close to two-thirds the average of the two actions as possible. We claim that for each player i , every action in $(50, 100]$ is strictly dominated by the action 50. To see this, for any opponent action $s_{-i} \in [0, 100]$, we have $\frac{2}{3} \cdot \frac{50 + s_{-i}}{2} \leq 50$, so the guess 50 is already too high:

$$50 - \frac{2}{3} \cdot \frac{50 + s_{-i}}{2} \geq 0$$

At the same time, $\frac{d}{ds_i} \left[s_i - \frac{2}{3} \cdot \frac{s_i + s_{-i}}{2} \right] = \frac{2}{3} > 0$. Hence we conclude playing any $s_i > 50$ exacerbates the error relative to playing 50,

$$s_i - \frac{2}{3} \cdot \frac{s_i + s_{-i}}{2} > 50 - \frac{2}{3} \cdot \frac{50 + s_{-i}}{2} \geq 0$$

so then $-\left(s_i - \frac{2}{3} \cdot \frac{s_i + s_{-i}}{2}\right)^2 < -\left(50 - \frac{2}{3} \cdot \frac{50 + s_{-i}}{2}\right)^2$ for all $s_i \in (50, 100]$ and we have the claimed strict dominance.

This means we may delete the set of actions $(50, 100]$ from each S_i to arrive at a new game $\mathcal{G}^{(2)}$ where each player is restricted to using only $[0, 50]$. The game $\mathcal{G}^{(2)}$ will have the same set of NEs as the original game. But the same logic may be applied again to show that in $\mathcal{G}^{(2)}$, for each player any action in $(25, 50]$ is strictly dominated by the action 25. We may continue in this way iteratively to arrive at a sequence of games $(\mathcal{G}^{(k)})_{k \geq 1}$, so that in the game $\mathcal{G}^{(k+1)}$, player i 's action set is $\left[0, \left(\frac{1}{2}\right)^k \cdot 100\right]$. All of the games $\mathcal{G}^{(1)}, \mathcal{G}^{(2)}, \mathcal{G}^{(3)}, \dots$ have the same NEs. This means any NE of $\mathcal{G}^{(1)}$ must involve each player playing an action in

$$\bigcap_{k=1}^{\infty} \left[0, \left(\frac{1}{2}\right)^k \cdot 100\right] = \{0\}$$

Hence, $(0, 0)$ is the unique NE. \blacklozenge

¹¹The name ‘‘beauty-contest game’’ comes from a newspaper game where readers pick the 6 faces they consider the most beautiful from a set of 100 portraits. The readers who pick the six most popular choices won a prize.

4 Extensions of Nash Equilibrium

4.1 Correlated equilibrium. Let's begin with the definition of a correlated equilibrium in a normal-form game.

Definition 24. In a normal form game $\mathcal{G} = \langle \mathcal{N}, (S_k)_{k \in \mathcal{N}}, (u_k)_{k \in \mathcal{N}} \rangle$, a **correlated equilibrium** (CE) consists of:

- A **finite set of signals** Ω_i for each $i \in \mathcal{N}$. Write $\Omega := \times_{k \in \mathcal{N}} \Omega_k$.
- A **joint distribution** $p \in \Delta(\Omega)$, so that the marginal distributions satisfy $p_i(\omega_i) > 0$ for each $\omega_i \in \Omega_i$.¹²
- A **signal-dependent strategy** $s_i^* : \Omega_i \rightarrow S_i$ for each $i \in \mathcal{N}$

such that for every $i \in \mathcal{N}$, $\omega_i \in \Omega_i$, $\hat{s}_i \in S_i$,

$$\sum_{\omega_{-i}} p(\omega_{-i} | \omega_i) \cdot u_i(s_i^*(\omega_i), s_{-i}^*(\omega_{-i})) \geq \sum_{\omega_{-i}} p(\omega_{-i} | \omega_i) \cdot u_i(\hat{s}_i, s_{-i}^*(\omega_{-i}))$$

A correlated equilibrium envisions the following situation. At the start of the game, an N -dimensional vector of signals ω realizes according to the distribution p . Player i observes only the i -th dimension of the signal, ω_i , and plays an action $s_i^*(\omega_i)$ as a function of the signal she sees. Whereas a pure Nash equilibrium has each player playing one action and requires that no player has a profitable unilateral deviation, in a correlated equilibrium each player may take **different actions depending on her signal**. Correlated equilibrium requires that no player can strictly improve her expected payoffs after seeing any of her signals. More precisely, seeing the signal ω_i leads her to have some belief over the signals that **others must have seen**, formalized by the conditional distribution $p(\cdot | \omega_i) \in \Delta(\Omega_{-i})$. Since she knows how these opponent signals translate into opponent actions through s_{-i}^* , she can compute the expected payoffs of taking different actions after seeing signal ω_i . She finds it optimal to play the action $s_i^*(\omega_i)$ instead of deviating to any other $\hat{s}_i \in S_i$ after seeing signal ω_i .

We make two remarks about correlated equilibria.

(1) The signal space and its associated joint distribution, (Ω, p) , are **not part of the game** \mathcal{G} , but part of the equilibrium. That is, a correlated equilibrium constructs an information structure under which a particular outcome can arise.

(2) There is no institution compelling player i to play the action $s_i^*(\omega_i)$, but i finds it optimal to do so after seeing the signal ω_i . It might be helpful to think of the **traffic lights** as an analogy for a correlated equilibrium. The light color that a player sees as she arrives at the intersection is her signal and let's imagine a world where there is no traffic police or cameras enforcing traffic rules. Each driver would nevertheless still find it optimal to stop when she sees a red light, because she infers that her seeing the red light signal must mean the driver on the intersecting street received the green light signal, and further the other driver is playing the strategy of going through the intersection if he sees a green light. Even though the red light (ω_i) merely **recommends** an action ($s_i^*(\omega_i)$), i finds it optimal to **obey this recommendation** given how others are acting on their own signals.

Example 25. Consider the usual coordination game, given by the payoff matrix:

¹²This is without loss of generality since any 0 probability signal of player i may be deleted to generate a smaller signal space.

	L	R
L	1,1	0,0
R	0,0	1,1

The following is a correlated equilibrium: $\Omega_1 = \Omega_2 = \{l, r\}$, $p(l, l) = 0.3$, $p(l, r) = 0.1$, $p(r, l) = 0.2$, $p(r, r) = 0.4$, $s_i^*(l) = L$ and $s_i^*(r) = R$ for each $i \in \{1, 2\}$. We can check that no player has a profitable deviation after any signal. For instance, after P1 sees the signal l , he knows that $p(\omega_2 = l | \omega_1 = l) = \frac{3}{4}$, $p(\omega_2 = r | \omega_1 = l) = \frac{1}{4}$. Since $s_2^*(l) = L$, $s_2^*(r) = R$, the expected payoff for P1 to playing L is $\frac{3}{4} \cdot 1 + \frac{1}{4} \cdot 0 = \frac{3}{4}$, whereas the expected payoff to playing R is $\frac{3}{4} \cdot 0 + \frac{1}{4} \cdot 1 = \frac{1}{4}$. As such, P1 does not want to deviate to playing R after seeing signal l . Similar arguments can be made for P1 after signal r , P2 after signal l , and P2 after signal r .

Here's another correlated equilibrium: $\Omega_1 = \Omega_2 = \{l, r\}$, $p(l, l) = 0.8$, $p(r, r) = 0.2$, $p(l, r) = 0$, $p(r, l) = 0$, $s_i^*(l) = L$ and $s_i^*(r) = R$ for each $i \in \{1, 2\}$. Note that the signal structures of different correlated equilibria **need not be the same**. In this example, the signal structure is effectively a “public randomization device” that picks the (L, L) Nash equilibrium 80% of the time, the (R, R) Nash equilibrium 20% of the time. This can be made more general. \blacklozenge

Example 26. (Public randomization device) Fix any normal-form game \mathcal{G} and fix K of its pure Nash equilibria, $\mathcal{E}_1, \dots, \mathcal{E}_K$, where each \mathcal{E}_k abbreviates some pure strategy profile $(\bar{s}_1^{(k)}, \dots, \bar{s}_N^{(k)})$. Then, for any K probabilities p_1, \dots, p_K with $p_k > 0$, $\sum_{k=1}^K p_k = 1$, consider the signal structure with $\Omega_i = \{1, \dots, K\}$, $p(k, \dots, k) = p_k$ for each $1 \leq k \leq K$, and $p(\omega) = 0$ for any ω where not all N dimensions match. Consider the strategies $s_i^*(k) = \bar{s}_i^{(k)}$ for each $i \in \mathcal{N}$, $1 \leq k \leq K$. Then (Ω, p, s^*) is a correlated equilibrium. Indeed, after seeing the signal k , each player i knows that others must be playing their part of the k -th Nash equilibrium, $(\bar{s}_1^{(k)}, \dots, \bar{s}_N^{(k)})$. As such, her best response must be $\bar{s}_i^{(k)}$, so $s_i^*(k) = \bar{s}_i^{(k)}$ is optimal. \blacklozenge

Example 27. (Coordination game with an eavesdropper) Three players Alice (P1), Bob (P2), and Eve (P3, the “eavesdropper”) play a zero-sum coordination game. Alice and Bob win only if they show up at the same location, and furthermore Eve is not there to spy on their conversation. The payoffs are given below. Alice chooses a row, Bob chooses a column, and Eve chooses a matrix.

	L	R
L	-1,-1,2	-1,-1,2
R	-1,-1,2	1,1,-2

matrix L

	L	R
L	1,1,-2	-1,-1,2
R	-1,-1,2	-1,-1,2

matrix R

The following is a correlated equilibrium. $\Omega_1 = \Omega_2 = \Omega_3 = \{l, r\}$, $p(l, l, l) = 0.25$, $p(l, l, r) = 0.25$, $p(r, r, l) = 0.25$, $p(r, r, r) = 0.25$, $s_i^*(l) = L$ and $s_i^*(r) = R$ for all $i \in \{1, 2, 3\}$. The information structure models a situation where Alice and Bob jointly observe some randomization device unseen by Eve¹³ and use it to coordinate on either both playing L or both playing R. Eve's signals are uninformative of Alice and Bob's actions. Indeed, after seeing either $\omega_3 = l$ or $\omega_3 = r$, Eve thinks the chances are 50-50 that Alice and Bob are both playing L or both playing R, so she has no profitable deviation from the prescribed actions $s_3^*(l) = L$, $s_3^*(r) = R$. On the other hand, after seeing $\omega_1 = l$,

¹³Perhaps an encrypted message.

Alice knows for sure that Bob is playing L while Eve has a 50-50 chance of playing L or R. Her payoff is maximized by playing the recommended $s_1^*(l) = L$. (You can check the other deviations similarly.) Eve's expected payoff in this correlated equilibrium is $\frac{1}{2} \cdot 2 + \frac{1}{2} \cdot (-2) = 0$. However, if Alice and Bob were to play **independent** mixed strategies, then Eve's best response leaves her with an **expected payoff of at least 1**. To see this, suppose Alice plays L with probability q_A and Bob plays L with probability q_B . If $q_A \cdot q_B \geq (1 - q_A) \cdot (1 - q_B)$, so that it is more likely that Alice and Bob coordinate on L than on R, Eve may play L to get an expected payoff of:

$$\underbrace{(-2) \cdot (1 - q_A) \cdot (1 - q_B)}_{\text{Alice and Bob meet without Eve}} + \underbrace{(2) \cdot [1 - (1 - q_A) \cdot (1 - q_B)]}_{\text{otherwise}} \geq (-2) \cdot \frac{1}{4} + (2) \cdot \frac{3}{4} = 1$$

where we used the fact that $q_A \cdot q_B \geq (1 - q_A) \cdot (1 - q_B) \Rightarrow q_A + q_B \geq 1 \Rightarrow (1 - q_A) \cdot (1 - q_B) \leq \frac{1}{4}$. On the other hand, if $q_A \cdot q_B \leq (1 - q_A) \cdot (1 - q_B)$, then Eve may play R to get an expected payoff of at least 1. \blacklozenge

4.2 Strong equilibrium. Whereas Nash equilibrium rules out profitable unilateral deviations, strong equilibrium (StrE) rules out **profitable simultaneous deviations** involving a coalition of players.

Definition 28. In a normal-form game \mathcal{G} , a **coalition** is a non-empty subset of players $C \subseteq \mathcal{N}$, $C \neq \emptyset$. Say the coalition C **blocks** the strategy profile (s_1, \dots, s_N) if there is some $(\hat{s}_j)_{j \in C}$ such that for every $j \in C$,

$$u_j(\hat{s}_C, s_{\mathcal{N} \setminus C}) > u_j(s_1, \dots, s_N)$$

That is to say, if there is some group of players who can simultaneously deviate from the profile (s_1, \dots, s_N) in a way that makes every group member strictly better off than under the strategy profile (s_1, \dots, s_N) , then the group is said to block the strategy profile.

Definition 29. A strategy profile (s_1^*, \dots, s_N^*) is a **strong equilibrium** if it is not blocked by any coalition.

In particular, absence of blocking coalitions of size 1 means every strong equilibrium is a Nash equilibrium. Absence of blocking coalitions of size N means every strong equilibrium is Pareto efficient.

Example 30. In the game of assurance,

	L	R
T	1,1	0,0
B	0,0	2,2

the strategy profile (T, L) is a Nash equilibrium, but it is blocked by the coalition $C = \{1, 2\}$. The coalition may play (B, R) instead and every coalition member will be strictly better off. \blacklozenge

(1) Rationalizability; (2) Nash implementation; (3) Bayesian games; (4) Universal type space

TF: Kevin He (he02@fas.harvard.edu)

1 Rationalizability

1.1 Two algorithms. Consider a normal-form game \mathcal{G} . Here we review the two algorithms of iterative strategy elimination studied in lecture.

Algorithm 31. (*Iterated elimination of strictly dominated strategies, “IESDS”*) Put $\hat{S}_i^{(0)} := S_i$ for each i . Then, having defined $\hat{S}_i^{(t)}$ for each i , we define $\hat{S}_i^{(t+1)}$ in the following way:

$$\hat{S}_i^{(t+1)} := \left\{ s_i \in \hat{S}_i^{(t)} : \nexists \sigma_i \in \Delta \left(\hat{S}_i^{(t)} \right) \text{ s.t. } u_i(\sigma_i, s_{-i}) > u_i(s_i, s_{-i}) \forall s_{-i} \in \hat{S}_{-i}^{(t)} \right\}$$

Finally, define $\hat{S}_i^\infty := \bigcap_{t \geq 0} \hat{S}_i^{(t)}$.

The idea behind IESDS is that if some mixed strategy σ_i yields **strictly** more payoff than the action s_i regardless of what other players do, then i should never use the action s_i . The “**iterated**” part comes from requiring that (i) the dominating mixed strategy must be supported on i ’s actions that **survived the previous rounds** of eliminations; (ii) the conjecture of what other players might do must be taken from their strategies that **survived the previous rounds** of eliminations.

Algorithm 32. (*Iterated elimination of never best responses, “IENBR”*) Put $\tilde{S}_i^{(0)} := S_i$ for each i . Then, having defined $\tilde{S}_i^{(t)}$ for each i , we define $\tilde{S}_i^{(t+1)}$ in the following way:

$$\tilde{S}_i^{(t+1)} := \left\{ s_i \in \tilde{S}_i^{(t)} : \exists \sigma_{-i} \in \Delta \left(\tilde{S}_{-i}^{(t)} \right) \text{ s.t. } u_i(s_i, \sigma_{-i}) \geq u_i(s'_i, \sigma_{-i}) \forall s'_i \in \tilde{S}_{-i}^{(t)} \right\}$$

Finally, define $\tilde{S}_i^\infty := \bigcap_{t \geq 0} \tilde{S}_i^{(t)}$.

It is important to note that $\Delta \left(\tilde{S}_{-i}^{(t)} \right) \neq \times_{k \neq i} \Delta \left(\tilde{S}_k^{(t)} \right)$.¹⁴ The left-hand-side is the set of **correlated mixed strategies** of players other than i , i.e. the set of all joint distributions on $\tilde{S}_{-i}^{(t)}$. Such a correlated mixed strategy might be generated, for example, using a signal-space kind of setup as in correlated equilibrium. The elimination of never best responses can be viewed as asking each action of player i to “justify its existence” by naming a correlated¹⁵ mixed strategy of opponents for which it is a best response. The “**iterated**” part comes from requiring that this conjecture of correlated opponents’ strategy have support in their strategies that **survived the previous rounds** of eliminations.

Another view on these two algorithms is that they make progressively **sharper predictions** about the game’s outcome by making more and more levels of **rationality assumptions**. A “rational”

¹⁴When there are only two players, $\Delta \left(\tilde{S}_{-i}^{(t)} \right) = \times_{k \neq i} \Delta \left(\tilde{S}_k^{(t)} \right)$. This is because $-i$ refers to exactly 1 player, not a group of players, so we do not get anything new by allowing $-i$ to “correlate amongst themselves”. As such, we did not have to worry about correlated vs. independent opponent strategies when we computed rationalizable strategy profiles for a two-player game in lecture.

¹⁵This correlation might reflect i ’s belief that opponents are colluding and coordinating their actions, or it could reflect correlation in i ’s subjective uncertainty about what two of her opponents might do.

player i is someone who maximizes the utility function u_i as given in the normal-form game \mathcal{G} . Rational players are contrasted against the so-called “**crazies**” present in some models of reputation, who are irrational in the sense of either maximizing a different utility function than normal players, or in not choosing actions based on utility maximization at all. From the analysts’ perspective, knowing that every player is rational allows us to predict that only actions in $\hat{S}_i^{(1)}$ (equivalently, $\tilde{S}_i^{(1)}$) will be played by i , since playing any other action is incompatible with maximizing u_i . But we cannot make more progress unless we are **also willing to assume** what i knows about j ’s rationality. If i is rational but i thinks that j might be crazy, in particular that j might take an action in $S_j \setminus \hat{S}_j^{(1)}$, then the step for constructing $\hat{S}_i^{(2)}$ for i does not make sense. As it is written in Algorithm 31, we should eliminate any action of i that does strictly worse than a fixed mixed strategy against all action profiles taken from $\hat{S}_{-i}^{(1)}$, which in particular assumes that j must be playing something in $\hat{S}_j^{(1)}$. In general, the t -th step for each of Algorithm 31 and Algorithm 32 rests upon assumptions of the form “ i knows that j knows that ... that k is rational” with length t .

1.2 Equivalence of the two algorithms. In fact, Algorithm 31 and Algorithm 32 are the equivalent, as we now demonstrate.

Proposition 33. $\hat{S}_i^{(t)} = \tilde{S}_i^{(t)}$ for each $i \in \mathcal{N}$ and $t = 0, 1, 2, \dots$. In particular, $\hat{S}_i^\infty = \tilde{S}_i^\infty$.

In view of this result, we call \tilde{S}_i^∞ the “(correlated) rationalizable strategies of player i ”, but note that it can be computed through either IENBR or IESDS.

Proof. Do induction on t . When $t = 0$, $\hat{S}_i^{(0)} = \tilde{S}_i^{(0)} = S_i$ by definition. Suppose for each $i \in \mathcal{N}$, $\hat{S}_i^{(t)} = \tilde{S}_i^{(t)}$.

To establish that $\tilde{S}_i^{(t+1)} \subseteq \hat{S}_i^{(t+1)}$, take some $s_i^* \in \tilde{S}_i^{(t+1)}$. By definition of IENBR, there is some $\sigma_{-i} \in \Delta(\tilde{S}_{-i}^{(t)})$ s.t. $u_i(s_i^*, \sigma_{-i}) \geq u_i(s'_i, \sigma_{-i}) \forall s'_i \in \tilde{S}_i^{(t)}$. The inductive hypothesis lets us replace all tildes with hats, so that there is some $\sigma_{-i} \in \Delta(\hat{S}_{-i}^{(t)})$ s.t. $u_i(s_i^*, \sigma_{-i}) \geq u_i(s'_i, \sigma_{-i}) \forall s'_i \in \hat{S}_i^{(t)}$. If s_i^* were strictly dominated by some $\hat{\sigma}_i \in \Delta(\hat{S}_i^{(t)})$, then $u_i(s_i^*, \sigma_{-i}) < u_i(\hat{\sigma}_i, \sigma_{-i})$, because the same strict inequality holds at every s_{-i} in the support of σ_{-i} . By Fact 17, there exists some $\hat{s}_i \in \hat{S}_i$ with $\hat{\sigma}_i(\hat{s}_i) > 0$ so that $u_i(s_i^*, \sigma_{-i}) < u_i(\hat{s}_i, \sigma_{-i})$, contradicting s_i^* being a best response to σ_{-i} .

Conversely, suppose $s_i^* \in \hat{S}_i^{(t+1)}$. Combining definition of IESDS and the inductive hypothesis shows that for each $\sigma_i \in \Delta(\hat{S}_i^{(t)})$, there corresponds some $s_{-i} \in \tilde{S}_{-i}^{(t)}$ so that $u_i(s_i^*, s_{-i}) \geq u_i(\sigma_i, s_{-i})$. Now enumerate $\tilde{S}_{-i}^{(t)} = \{s_{-i}^{(1)}, \dots, s_{-i}^{(d)}\}$ and hence construct the following subset of \mathbb{R}^d :

$$V := \left\{ v \in \mathbb{R}^d : \exists \sigma_i \in \Delta(\tilde{S}_i^{(t)}) \text{ s.t. } v_k \leq u_i(\sigma_i, s_{-i}^{(k)}) \forall 1 \leq k \leq d \right\}$$

that is, every $\sigma_i \in \Delta(\tilde{S}_i^{(t)})$ gives rise to a point $(u_i(\sigma_i, s_{-i}^{(1)}) \dots u_i(\sigma_i, s_{-i}^{(d)})) \in \mathbb{R}^d$ and V is the region to the “lower-left” of this collection of points. We can verify that V is convex and non-empty. Now consider the singleton set $W = \left\{ (u_i(s_i^*, s_{-i}^{(1)}) \dots u_i(s_i^*, s_{-i}^{(d)})) \right\}$. We must have $W \cap \text{int}(V) = \emptyset$, where $\text{int}(V)$ is the interior of V . As such, separating hyperplane theorem implies there is some $q \in \mathbb{R}^d \setminus \{\mathbf{0}\}$ with $q \cdot (u_i(s_i^*, s_{-i}^{(1)}) \dots u_i(s_i^*, s_{-i}^{(d)})) \geq q \cdot v$ for all $v \in V$. Since V includes points with arbitrarily large negative numbers in each coordinate, we must in fact have $q \in \mathbb{R}_+^d \setminus \{\mathbf{0}\}$. So then, q may be normalized so that its dimensions are weakly positive numbers that add up to 1, i.e. it can be viewed as some correlated mixed strategy $\sigma_{-i}^* \in \Delta(\tilde{S}_{-i}^{(t)})$. This strategy has the property that $u_i(s_i^*, \sigma_{-i}^*) \geq u_i(\sigma_i, \sigma_{-i}^*)$ for all $\sigma_i \in \Delta(\tilde{S}_i^{(t)})$, showing that in particular s_i^* is a best response to σ_{-i}^* amongst $\tilde{S}_i^{(t)}$, hence $s_i^* \in \tilde{S}_i^{(t+1)}$. This establishes the reverse inclusion $\hat{S}_i^{(t+1)} \subseteq \tilde{S}_i^{(t+1)}$ and completes the inductive step. \square

1.3 Rationalizability and equilibrium concepts. In some sense, the collection of rationalizable strategies is a superset of the collection of correlated equilibrium strategies. To be more precise,

Proposition 34. *If (Ω, p, s^*) is a correlated equilibrium, then $s_i^*(\omega_i) \in \hat{S}_i^\infty$ for every i , $\omega_i \in \Omega_i$.*

Proof. We show for any player k and any s_k such that $s_k \in \text{range}(s_k^*)$, $s_k \in \tilde{S}_k^{(t)}$ for every t . This statement is clearly true when $t = 0$. Suppose this statement is true for $t = T$. Then, for each player i and each signal $\omega_i \in \Omega_i$, consider the correlated opponent strategy $\sigma_{-i}^*(\omega_i)$ constructed by

$$\sigma_{-i}^*(\omega_i)(s_{-i}) := p \left[\omega_{-i} : s_{-i}^*(\omega_{-i}) = s_{-i} \mid \omega_i \right]$$

By definition of CE, $s_i^*(\omega_i)$ best responds to $\sigma_{-i}^*(\omega_i)$. Furthermore, $\sigma_{-i}^*(\omega_i) \in \Delta(\tilde{S}_{-i}^{(T)})$ by inductive hypothesis. Therefore, $\hat{s}_i \in \tilde{S}_i^{(T+1)}$, completing the inductive step. \square

Therefore, we see that correlated equilibria (and in particular, Nash equilibria) **embed the assumption of common knowledge of rationality**: not only is Alice rational, but also Alice knows Bob is rational, and Alice knows that Bob knows Alice is rational, etc.

1.4 Nested solution concepts. Here we summarize the inclusion relationships between several solution concepts. For a normal-form game \mathcal{G} ,

$$\text{Rat}(\mathcal{G}) \supseteq \text{CE}(\mathcal{G}) \supseteq \text{NE}(\mathcal{G}) \supseteq \text{StrE}(\mathcal{G})$$

2 Mechanism Design and Nash Implementation

2.1 Mechanism design as a decentralized solution to the information problem.

Definition 35. A **mechanism design problem** (MDP) consists of the following:

- A finite collection of **players** $\mathcal{N} = \{1, \dots, N\}$
- A set of **states of the world** Θ
- A set of **outcomes** A
- A **state-dependent utility** $u_i : A \times \Theta \rightarrow \mathbb{R}$ for each player $i \in \mathcal{N}$
- A **social choice rule** $f : \Theta \rightrightarrows A$

Every MDP presents an information problem. Consider a Central Authority who is **omnipotent** (all-powerful) but **not omniscient** (all-knowing). It can choose any outcome $x \in A$. However, the outcome it wants to pick depends on the state of the world. When the state of the world is θ , Central Authority's favorite outcomes are $f(\theta)$. While every player knows the state of the world, the Central Authority does not. Think of, for example, a town hall (Central Authority) trying to decide how much taxes to levy (outcomes) on a community of neighbors (players), where the optimal taxation depends on the productivities of different neighbors, a state of the world that every neighbor knows but the town hall does not.

Due to Central Authority's ignorance of θ , it does not know which outcome to pick and must proceed more indirectly. The goal of the Central Authority is to come up with an incentive scheme, called a **mechanism**, that induces self-interested players to choose one of the Central Authority's favorite outcomes. The mechanism enlists the help of the players, who know the state of the world, in selecting an outcome optimal from the point of view of the Central Authority.

More precisely,

Definition 36. Given a MDP, a **mechanism** is a set of pure strategies $(S_i)_{i \in \mathcal{N}}$ for each player and a map $g : S \rightarrow A$.

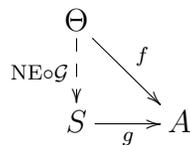
The Central Authority announces a set of pure strategies S_i for each player and a mapping between the profile of pure strategies and the outcome. The Central Authority promises to implement the outcome $g(s_1, \dots, s_N)$ when players choose the strategy profile (s_1, \dots, s_N) .

In state θ , the mechanism $\langle (S_k)_{k \in \mathcal{N}}, g \rangle$ **gives rise to a normal-form game**, $\mathcal{G}(\theta)$, where the set of actions of player i is S_i and the payoff i gets from strategy profile (s_1, \dots, s_N) is $u_i(g(s_1, \dots, s_N), \theta)$. The mechanism solves the Central Authority’s information problem if playing the game $\mathcal{G}(\theta)$ yields the same outcomes as $f(\theta)$. To predict what agents will do when they play the game $\mathcal{G}(\theta)$, the Central Authority must **pick a solution concept**. We will use Nash equilibrium.

Definition 37. The mechanism $\langle (S_k)_{k \in \mathcal{N}}, g \rangle$ **Nash-implements** social choice rule f if $g(\text{NE}(\mathcal{G}(\theta))) = f(\theta)$ for every $\theta \in \Theta$.

If the Central Authority wants to use a solution concept other than Nash equilibrium, then it would simply replace “NE” in the above definition.

We can also represent the definition of implementation diagrammatically. Suppose for the sake of simplicity that f is a (single-valued) function. Then Θ, A , and a function f between them is given by the MDP. The Central Authority chooses $(S_k)_{k \in \mathcal{N}}$ and $g : S \rightarrow A$. The choice of the mechanism induces a function $\text{NE}(\mathcal{G}(\theta))$ giving the Nash equilibrium (assumed unique for simplicity) in each state of the world. Then, mechanism $\langle (S_k)_{k \in \mathcal{N}}, g \rangle$ Nash-implements f if the final, $\text{NE} \circ \mathcal{G}$ arrow makes the following diagram **commute**.



Loosely speaking, mechanism design is “**reverse game theory**”. Whereas a game theorist takes the game as given and analyzes its equilibria, a mechanism designer takes the social choice rule as given and acts as a “Gamemaker”, aiming to engineer a game with suitable equilibria.

2.2 Maskin monotonicity and Nash implementation. It is natural to ask which MDP’s admit Nash implementations. As we saw in lecture the following pair of conditions is important.

Definition 38. A social choice rule f satisfies **Maskin monotonicity** (MM)¹⁶ if for all $\theta, \theta' \in \Theta$, whenever (1) $x \in f(\theta)$, and (2) $\{y : u_i(y, \theta) \leq u_i(x, \theta)\} \subseteq \{y : u_i(y, \theta') \leq u_i(x, \theta')\}$ for every i , then $x \in f(\theta')$ too.

In words, if f chooses x in state θ , then it should also choose x when the set of outcomes weakly worse than x expands for everyone.

Definition 39. A social choice rule f satisfies **no veto power** if for any $i \in \mathcal{N}$ and any $x \in A$, $u_j(x, \theta) \geq u_j(y, \theta)$ for all $j \neq i$ and all $y \in A$ implies $x \in f(\theta)$.

Theorem 40. *If f is Nash-implementable, then it satisfies MM. If $N \geq 3$ and f satisfies MM and no veto power, then f is Nash-implementable.*

¹⁶What Eric Maskin called “monotonicity” in lecture is usually referred to as “Maskin monotonicity” in the literature, cf. Footnote 9.

Proof. See lecture. □

Example 41. (A social choice rule satisfying no veto power but not MM) Suppose $N \geq 3$ and individuals have strict preferences over outcomes A in any state of the world. Consider the social choice rule “top-ranked rule” f^{top} , where $x \in f^{\text{top}}(\theta)$ iff for all $z \in A$,

$$\#\{i : u_i(x, \theta) > u_i(y, \theta) \text{ for all } y \neq x\} \geq \#\{i : u_i(z, \theta) > u_i(y, \theta) \text{ for all } y \neq z\}$$

That is, f^{top} chooses the outcome(s) top-ranked by the largest number of individuals. Then f^{top} satisfies no veto power. Indeed, $u_j(x, \theta) \geq u_j(y, \theta)$ for all $j \neq i$ and all $y \in A$, together with the assumption that all preferences are strict, implies that

$$\#\{i : u_i(x, \theta) > u_i(y, \theta) \text{ for all } y \neq x\} \geq N - 1$$

While for any $z \neq x$,

$$\#\{i : u_i(z, \theta) > u_i(y, \theta) \text{ for all } y \neq z\} \leq 1$$

Since $N \geq 3$, $x \in f(\theta)$.

However, consider the following preferences: In state θ , $u_1(x, \theta) > u_1(y, \theta) > u_1(z, \theta)$, $u_2(y, \theta) > u_2(z, \theta) > u_2(x, \theta)$, and $u_3(z, \theta) > u_3(y, \theta) > u_3(x, \theta)$; in state θ' , the preferences are unchanged except that $u_3(y, \theta) > u_3(z, \theta) > u_3(x, \theta)$. Then outcome x did not drop in ranking relative to any other outcome for any individual from θ to θ' , yet $f^{\text{top}}(\theta) = \{x, y, z\}$ while $f^{\text{top}}(\theta') = \{y\}$. This shows f^{top} does not satisfy MM, hence by Theorem 40 it is not Nash-implementable. ◆

Example 42. (A social choice rule satisfying MM but not no veto power) Suppose individuals have strict preferences over outcomes A in any state of the world. Consider the social choice rule “dictator’s rule” f^{D} , where f^{D} simply chooses the top-ranked outcome of player 1, a dictator. Then f^{D} satisfies MM. To see this, $x \in f^{\text{D}}(\theta)$ implies $u_1(x, \theta) > u_1(y, \theta)$ for any $y \neq x$, but in any state of the world θ' where x does not fall in ranking relative to any other outcome for any individual, it remains true that $u_1(x, \theta') > u_1(y, \theta')$ for any $y \neq x$. As such, $x \in f^{\text{D}}(\theta')$ also. However, f^{D} does not satisfy no veto power. If $A = \{x, y\}$, $N = 3$, then in a state of the world with $u_1(x, \theta) = 1, u_1(y, \theta) = 0, u_2(x, \theta) = 0, u_2(y, \theta) = 1, u_3(x, \theta) = 0, u_3(y, \theta) = 1$, we have y being top-ranked for all individuals except 1, yet $f^{\text{D}}(\theta) = \{x\}$. Theorem 40 does not say whether f^{D} is Nash-implementable or not. However, it is easy to see that a mechanism with $S_i = \{x, y\}$ for $i \in \mathcal{N}$ and $g(s_1, \dots, s_N) = s_1$ Nash-implements f^{D} . The mechanism elicits a message from each player but only implements the action of player 1 while ignoring everyone else. This example shows MM plus no veto power are sufficient for Nash-implementability when $N \geq 3$, but they are **not necessary**. ◆

3 Bayesian Games

3.1 The common prior model of a Bayesian game. In our brief encounter with mechanism design, we considered a setting where the Central Authority is uncertain as to the state of the world $\theta \in \Theta$, but every player knows θ perfectly. Many economic situations involve uncertainty about payoff-relevant state of the world amongst even the **players themselves**. To take some examples:

- Auction participants are uncertain about other bidders’ willingness to pay
- Investors are uncertain about the profitability of a potential joint venture

- Traders are uncertain about the value of a financial asset at a future date

How should a group of **Bayesian** players confront such uncertainty? While there exist some more general approaches (see the optional material on the universal type space, for example), most models of incomplete-information games you will encounter will impose the common prior assumption.

Definition 43. A **Bayesian game with common prior assumption (CPA)** is

$$\mathcal{B} = \langle \mathcal{N}, (S_k)_{k \in \mathcal{N}}, (\Theta_k)_{k \in \mathcal{N}}, \mu, (u_k)_{k \in \mathcal{N}} \rangle$$

consisting of:

- A finite collection of **players** $\mathcal{N} = \{1, 2, \dots, N\}$
- A set of **actions** S_i for each $i \in \mathcal{N}$
- A set of **states of the world** $\Theta = \times_{k=1}^N \Theta_k$.
- A **common prior** $\mu \in \Delta(\Theta)$
- A **utility function** $u_i : \times_{k=1}^N S_k \times \Theta \rightarrow \mathbb{R}$ for each $i \in \mathcal{N}$

Definition 44. A **pure strategy** of player i in a Bayesian game is a function $s_i : \Theta_i \rightarrow S_i$.

For ease of exposition, for now we will focus on the case where Θ is finite. (The CPA Bayesian game model can also accommodate games with infinitely many states of the world, such as auctions with a continuum of possible valuations for each bidder.)

A CPA Bayesian game (or just “Bayesian game” for short) proceeds as follows. A state of the world θ is drawn according to μ . Player i learns the i -th dimension, θ_i , then takes an action from her strategy set S_i . The utility of player i depends on the profile of actions as well as the state of the world θ , so in particular it might depend on the dimensions of θ that i **does not observe**. The subset of Bayesian games where u_i does **not** depend on θ_{-i} are called **private-values** games.

Player i 's strategy is a function of θ_i , not of θ , for i can only condition her action on her partial knowledge of the state of the world. For reasons we make clear later, Θ_i is often called the **type space** of i and one often describes a strategy of i as “type θ'_i does X, while type θ''_i does Y”.

A strategy profile in a Bayesian game might remind you of a **correlated equilibrium**. Indeed, in both setups each player observes some realization (her signal in CE, her type in Bayesian game), then performs an action dependent on her observation. However, unlike (Ω, p) in the definition of a correlated equilibrium, the (Θ, μ) in a Bayesian game is **part of the game**, not part of the solution concept. Furthermore, while the signal profile $\omega \in \Omega$ in a CE is only a coordination device that does not by itself affect players' payoffs (as in an unenforced traffic light), the state of the world in a Bayesian game is **payoff-relevant**.

Example 45. (August 2013 General Exam) Two players play a game. With probability 0.5, the payoffs are given by the left payoff matrix. With probability 0.5, they are given by the right payoff matrix. P1 knows whether the actual game is given by the left or right matrix, while P2 does not. Model this situation as a Bayesian game.

	L	C	R
T	-2,-2	-1,1	0,0
M	1,-1	3,5	3,4
B	0,0	4,2	2,4

	L	C	R
T	0,0	0,0	0,0
M	0,0	0,0	0,0
B	0,0	1,0	4,4

Solution: Let $\Theta_1 = \{l, r\}$, $\Theta_2 = \{0\}$, $\mu \in \Delta(\Theta)$ with $\mu(l, 0) = \mu(r, 0) = 0.5$. In state $(l, 0)$, the payoffs are given by the left matrix. In state $(r, 0)$, they are given by the right matrix. There are thus two types of P1: the type who knows that the payoffs are given by the left matrix, and the type who knows that the payoffs are given by the right one. There is only one type of P2. The utility of each player depends on (s_1, s_2, θ) , for example $u_2(B, C, (l, 0)) = 2$ while $u_2(B, C, (r, 0)) = 0$. In particular, the payoff to P2 depends on θ_1 , so this is **not a private-values game**.

A pure strategy of P1 in this Bayesian game is a function $s_1 : \Theta_1 \rightarrow S_1$, in other words the strategy must specify what the l -type and r -type of P1 will do. A strategy of P2 is a function $s_2 : \Theta_2 \rightarrow S_2$, but since Θ_2 is just a singleton P2 has just one action in each of her pure Bayesian game strategies. \blacklozenge

3.2 Bayesian Nash equilibrium. Here's the most common equilibrium concept for Bayesian games.

Definition 46. A **Bayesian Nash equilibrium (BNE)** is a strategy profile (s_i^*) in a Bayesian game, such that for each player $i \in \mathcal{N}$, each type $\theta_i \in \Theta_i$,

$$s_i^*(\theta_i) \in \arg \max_{\hat{s}_i \in S_i} \left\{ \sum_{\theta_{-i}} u_i(\hat{s}_i, s_{-i}^*(\theta_{-i}), (\theta_i, \theta_{-i})) \cdot \mu(\theta_{-i} | \theta_i) \right\}$$

A BNE might be understood as a “**correlated equilibrium with payoff-relevant signals**”. After observing her type θ_i , player i derives from the common prior a conditional belief $\mu(\cdot | \theta_i) \in \Delta(\Theta_{-i})$ about the types of other players. She knows $s_{-i}^*(\cdot)$, so she knows how these opponent types translate into opponent actions. Unlike in a CE, however, she knows that her payoff also depends on the complete state of the world, $\theta = (\theta_i, \theta_{-i})$. Analogous to CE, a BNE is a strategy profile such that, after player i observes her type θ_i and calculates her expected payoffs to different actions, she finds it optimal to play the prescribed action $s_i^*(\theta_i)$ across all of her choices in S_i .

Example 47. (August 2013 General Exam) Find all the pure-strategy BNEs in Example 45.

Solution: The best way to do this is to systematically check all strategy profiles. Since P2 has only one type, it is easiest to break things down by P2's action in equilibrium.

- Can there exist a BNE where $s_2^*(0) = L$? In any such BNE, we must have $s_1^*(l) = M$ since the type- l P1 knows for sure that P2 is playing L and that payoffs are given by the left matrix, leading to a unique best response of M. Yet this means P2 (of type 0) has a profitable deviation. Playing C yields an expected payoff of $\frac{1}{2} \cdot 5 + \frac{1}{2} \cdot 0 = 2.5$ (regardless of what $s_1^*(r)$ is), which is better than playing L and getting an expected payoff of $\frac{1}{2} \cdot (-1) + \frac{1}{2} \cdot 0 = -0.5$. Therefore, there is no BNE with $s_2^*(0) = L$.
- Can there exist a BNE where $s_2^*(0) = C$? In any such BNE, we must have $s_1^*(l) = s_1^*(r) = B$ for similar reasoning as above. But that means P2 gets an expected payoff of $\frac{1}{2} \cdot 2 + \frac{1}{2} \cdot 0 = 1$ by playing C, yet he can get $\frac{1}{2} \cdot 4 + \frac{1}{2} \cdot 4 = 4$ by playing R. Therefore, there is no BNE with $s_2^*(0) = C$.

- Can there exist a BNE where $s_2^*(0) = R$? In any such BNE, we must have $s_1^*(l) = M$, $s_1^*(r) = B$ for similar reasoning as above. As such, P2 gets an expected payoff of $\frac{1}{2} \cdot 4 + \frac{1}{2} \cdot 4 = 4$ from playing R. By comparison, he would get an expected $\frac{1}{2} \cdot (-1) + \frac{1}{2} \cdot 0 = -0.5$ from playing L and an expected $\frac{1}{2} \cdot 5 + \frac{1}{2} \cdot 0 = 2.5$ from playing C. (It is not feasible for P2 to “play C in the left matrix, play R in the right matrix” since he can only condition his action on his type. P2 has only one type since he knows only the prior probabilities of the two matrices, but not which one is actually being played.) Therefore, we see that $s_1^*(l) = M$, $s_1^*(r) = B$, $s_2^*(0) = R$ is the unique pure-strategy BNE of the game.

◆

4 The Universal Type Space

This is purely for those of you who want to learn more about what happens when we drop the common prior assumption. It won't come up on the exams.

4.1 Higher orders of belief. We have considered a Bayesian game with CPA as a model of how a group of Bayesian players confront uncertainty. But the CPA model makes several assumptions: (1) Θ is assumed to have a **product structure**; (2) it is **common knowledge** that θ is drawn according to μ . That is to say, everyone knows μ , everyone knows that everyone else knows μ , etc. What if we relax the common prior assumption? That is to say, how should a group of Bayesian players **in general** behave when confronting uncertainty Θ ?

If there is only one player, then the answer is simple. The Bayesian player comes up with a prior $\mu \in \Delta(\Theta)$ through introspection, then chooses some $s_1 \in S_1$ as to maximize $\int_{\theta \in \Theta} u_1(s_1, \theta) d\mu(\theta)$. The prior μ is trivially a common prior, since there is only one player.

However, in a game involving two players¹⁷, the answer becomes far more complex. P1 is uncertain not only about state of the world Θ , but **also about P2's belief** over state of the world. P2's belief matters for P1's decision-making, since P1's utility depends on the pair (P1's action, P2's action) while P2's action depends on his belief. As a Bayesian must form a prior distribution over any relevant uncertainty, P1 should entertain not only a belief about state of the world, but also a **belief about P2's belief**, which is also unknown to P1.

To take a more concrete example, suppose there are two players Alice and Bob and the states of the world concern the weather tomorrow, $\Theta = \{\text{sunny, rain}\}$. Alice believes that there is a 60% chance that it is sunny tomorrow, 40% chance that it rains, so we say she has a **first-order belief** $\mu_{\text{Alice}}^{(1)} \in \Delta(\Theta)$ with $\mu_{\text{Alice}}^{(1)}(\text{sunny}) = 0.6$, $\mu_{\text{Alice}}^{(1)}(\text{rain}) = 0.4$. Now Alice needs to form a belief about Bob's belief regarding tomorrow's weather. Alice happens to know that Bob is a meteorologist who has access to more weather information than she does. In particular, Alice believes Bob's belief about weather tomorrow is **correlated with the actual weather** tomorrow. Either it is the case that tomorrow will be sunny and Bob believes today that it will be sunny tomorrow with probability 90%, or it is the case that tomorrow will rain and today Bob believes it will be sunny with probability 20%. Alice assigns 60-40 odds to these two cases. We say Alice has a **second-order belief** $\mu_{\text{Alice}}^{(2)} \in \Delta(\Theta \times \Delta(\Theta))$, where $\mu_{\text{Alice}}^{(2)}$ is supported on two points $(\text{sunny}, \mu_{\text{case 1}}^{(1)})$, $(\text{rain}, \mu_{\text{case 2}}^{(1)})$ with $\mu_{\text{Alice}}^{(2)}[\text{sunny}, \mu_{\text{case 1}}^{(1)}] = 0.6$, $\mu_{\text{Alice}}^{(2)}[\text{rain}, \mu_{\text{case 2}}^{(1)}] = 0.4$. Here $\mu_{\text{case 1}}^{(1)}$ and $\mu_{\text{case 2}}^{(1)}$ are elements of $\Delta(\Theta)$ and $\mu_{\text{case 1}}^{(1)}(\text{sunny}) = 0.9$ while $\mu_{\text{case 2}}^{(1)}(\text{sunny}) = 0.2$. We are not finished. Surely Bob, like Alice, also holds some second-order belief. Alice is uncertain about Bob's second-order belief, so she

¹⁷All of this extends to games with 3 or more players, but with more cumbersome notations.

must form a **third-order belief**

$$\mu_{\text{Alice}}^{(3)} \in \Delta(\Theta \times \Delta(\Theta) \times \Delta(\Theta \times \Delta(\Theta)))$$

that is a joint distribution over (i) the weather tomorrow; (ii) Bob’s first-order belief about the weather; (iii) Bob’s second-order belief about the weather. Alice further needs a fourth-order belief, fifth-order belief, and so on.

We highlight the following features of the above example, which will be relevant to the subsequent theory on the universal type space:

- Alice entertains beliefs of order 1, 2, 3, ... about the state of the world, where k -th order belief is a joint distribution over state of the world, Bob’s first-order belief, Bob’s second-order belief, ..., Bob’s $(k - 1)$ -th-order belief.
- Alice’s second-order belief is **consistent** with her first-order belief, in the sense that whereas $\mu_{\text{Alice}}^{(1)}$ assigns probability of 60% to sunny weather tomorrow, $\mu_{\text{Alice}}^{(2)}$ marginalized to a distribution only over the weather also says there is a 60% chance that it is sunny tomorrow.
- There is **no common prior** over the weather and no signal structure is explicitly given.

Harsanyi first conjectured [2] in 1967 that for each specification of states of the world Θ , there corresponds an object now called the “**universal type space**”¹⁸, say T . Points in the universal type space correspond to all “reasonable” hierarchies of first-order belief, second-order belief, third-order belief, ... that a player could hold about Θ . Furthermore, there exists a “natural” homeomorphism

$$f : T \rightarrow \Delta(\Theta \times T)$$

so that each universal type t **encodes a joint belief** $f(t)$ over the state of the world and opponent’s universal type. The universal type space is thus “universal” in the senses of (i) capturing all possible hierarchies of beliefs that might arise under some signal structure about Θ ; (ii) putting an end to the seemingly infinite regress of having to resort to $(k + 1)$ -th-order beliefs in order to model beliefs about k -th-order beliefs, then having to discuss $(k + 2)$ -th-order beliefs to describe beliefs about the $(k + 1)$ -th-order beliefs just introduced, etc.

4.2 Constructing the universal type space. Mertens and Zamir first constructed the universal type space in 1985 [3]. Brandenburger and Dekel gave an alternative, simpler construction¹⁹ in 1993, which we sketch here [4].

There are two players, i and j . The states of the world Θ is a Polish space (complete, separable metric space). For each Polish space Z , write $\Delta(Z)$ for the set of probability measures on Z ’s Borel σ -algebra. It is known that $\Delta(Z)$ is metrizable by the Prokhorov metric, which makes $\Delta(Z)$ a Polish space of its own right.

Iteratively, define $X_0 := \Theta$, $X_1 := \Theta \times \Delta(X_0)$, $X_2 := \Theta \times \Delta(X_0) \times \Delta(X_1)$, etc. Each player has a first-order belief $\mu_i^{(1)}, \mu_j^{(1)} \in \Delta(X_0)$ that describes her belief about state of the world, a second-order belief $\mu_i^{(2)}, \mu_j^{(2)} \in \Delta(X_1)$ that describes her joint belief about state of the world and opponent’s first-order

¹⁸Harsanyi initially called members of such space “attribute vectors”. The word “type” only appeared in a later draft after Harsanyi discussed his research with Aumann and Maschler, who were also working on problems in information economics.

¹⁹Brandenburger and Dekel’s construction was based on a slightly different set of assumptions than that of Mertens and Zamir. For instance, Mertens and Zamir assumed Θ is compact, but Brandenburger and Dekel required Θ to be a complete, separable metric space. Neither is strictly stronger.

belief, and in general a k -th-order belief $\mu_i^{(k)}, \mu_j^{(k)} \in \Delta(X_{k-1}) = \Delta(\Theta \times \Delta(X_0) \times \dots \times \Delta(X_{k-2}))$ that describes her joint belief about state of the world, opponent's first-order belief, ... and opponent's $(k-1)$ -th-order belief. Since X_0 is Polish, each X_k is Polish.

A **hierarchy of beliefs** is a sequence of beliefs of all orders, $(\mu_i^{(1)}, \mu_i^{(2)}, \dots) \in \times_{k=0}^{\infty} \Delta(X_k) =: T_0$. Note that there is a great deal of **redundancy** within the hierarchy. Indeed, as $\mu_i^{(k)}$ is a distribution over the first k elements of $\Theta, \Delta(X_0), \Delta(X_1), \dots$, each $\mu_i^{(k)}$ can be appropriately marginalized to obtain a distribution over the same domain as $\mu_i^{(k')}$ for any $1 \leq k' < k$. Call a hierarchy of beliefs **consistent** if each $\mu_i^{(k)}$ marginalized on all except the last dimension equals $\mu_i^{(k-1)}$ and write $T_1 \subset T_0$ for the subset of consistent hierarchies. Then, **Kolmogorov extension theorem** implies for each consistent hierarchy $(\mu_i^{(1)}, \mu_i^{(2)}, \dots)$, there exists a measure $f(\mu_i^{(1)}, \mu_i^{(2)}, \dots)$ over the infinite product $\Theta \times (\times_{k=0}^{\infty} \Delta(X_k))$ such that $f(\mu_i^{(1)}, \mu_i^{(2)}, \dots)$ marginalized to $\Theta \times \dots \times \Delta(X_{k-1})$ equals $\mu_i^{(k)}$ for each $k = 0, 1, 2, \dots$. But $\Theta \times (\times_{k=0}^{\infty} \Delta(X_k))$ is in fact $\Theta \times T_0$, so that f associates each consistent hierarchy with a joint belief over state of the world and (possibly inconsistent) hierarchy of the opponent. Further, this association is **natural** in the sense that $f(\mu_i^{(1)}, \mu_i^{(2)}, \dots)$ describes the same beliefs and higher-order beliefs about Θ as the hierarchy $(\mu_i^{(1)}, \mu_i^{(2)}, \dots)$. We may further verify the map $f : T_1 \rightarrow \Delta(\Theta \times T_0)$ is bijective, continuous, and has a continuous inverse, so that it is a homeomorphism.

To close the construction, define a sequence of decreasing subsets of T_1 ,

$$T_k := \{t \in T_1 : f(t)(\Theta \times T_{k-1}) = 1\}$$

That is, T_k is the subset of consistent types who put probability 1 on opponent's type being in the subset T_{k-1} . Let $T := \bigcap_k T_k$, which is the class of types with "common knowledge of consistency": i "knows"²⁰ j 's type is consistent, i "knows" that j "knows" i 's type is consistent, etc. **This is the universal type space** over Θ . The map f can be restricted to the subset T to give a natural homeomorphism from T to $\Delta(\Theta \times T)$.

4.3 CPA Bayesian game as a belief-closed subset of the universal type space. Here we discuss how the CPA Bayesian game model relates to the universal type space.

Take a CPA Bayesian game $\mathcal{B} = \langle \mathcal{N}, (S_k)_{k \in \mathcal{N}}, (\Theta_k)_{k \in \mathcal{N}}, \mu, (u_k)_{k \in \mathcal{N}} \rangle$ and suppose for simplicity there are two players, i and j . Each $\theta_i^* \in \Theta_i$ corresponds to a unique point in the universal type space T over Θ , which we write as $t(\theta_i^*) \in T$. To identify $t(\theta_i^*)$, note that i of the type θ_i^* has a first-order belief $\mu_i^{(1)}[\theta_i^*] \in \Delta(\Theta)$, such that for $E_1 \subseteq \Theta$,

$$\mu_i^{(1)}[\theta_i^*](E_1) := \mu(E_1 | \theta_i^*)$$

where $\mu(\cdot | \theta_i^*) \in \Delta(\Theta)$ is the conditional distribution on Θ derived from the common prior, given that $\theta_i = \theta_i^*$.

Furthermore, θ_i^* also leads to a second-order belief $\mu_i^{(2)}[\theta_i^*] \in \Delta(\Theta \times \Delta(\Theta))$, where for $E_1 \subseteq \Theta$, $E_2 \subseteq \Delta(\Theta)$,

$$\mu_i^{(2)}[\theta_i^*](E_1 \times E_2) := \mu\left(\left\{\hat{\theta} \in \Theta : \hat{\theta} \in E_1 \text{ and } \mu_j^{(1)}[\hat{\theta}_j] \in E_2\right\} | \theta_i^*\right)$$

here $\mu_j^{(1)} : \Theta_j \rightarrow \Delta(\Theta)$ is defined analogously to $\mu_i^{(1)}$. One may similarly construct the entire hierarchy $t(\theta_i^*) = (\mu_i^{(k)}[\theta_i^*])_{k=1}^{\infty}$ and verify that it satisfies common knowledge of consistency. Hence, $t(\theta_i^*) \in T$.

²⁰More precisely, "knows" here means "puts probability 1 on".

This justifies calling elements of Θ_i “types” of P1, for indeed they correspond to universal types over the states of the world.

The set of universal types present in the Bayesian game \mathcal{B} , namely

$$T(\mathcal{B}) := \{t(\theta_k^*) : \theta_k^* \in \Theta_k, k \in \{i, j\}\}$$

is a **belief-closed subset** of T . That is, each $t \in T(\mathcal{B})$ satisfies $f(t)(\Theta \times T(\mathcal{B})) = 1$, putting probability 1 on the event that opponent is drawn from the set of universal types $T(\mathcal{B})$.

(1) Auction model; (2) Solving for auction BNEs; (3) Revenue equivalence

TF: Kevin He (he02@fas.harvard.edu)

1 The Auction Model

1.1 Definition of an auction. In section, we will make a number of simplifying assumptions instead of studying the most general auction model. We will assume that: (1) auctions are **private values**, so the types of $-i$ do not matter for i 's payoff; (2) the type distribution is **symmetric** and **independent** across players; (3) there is one seller who sells **one indivisible item**; (4) players are **risk neutral**, so getting the item with probability H and having to pay P in expectation gives a player i with type θ_i a utility of $\theta_i \cdot H - P$.

Definition 48. An auction $\mathcal{A} = \langle \mathcal{N}, F, [0, \bar{\theta}], (H_k)_{k \in \mathcal{N}}, (P_k)_{k \in \mathcal{N}} \rangle$ consists of:

- A finite set of **bidders** $\mathcal{N} = \{1, \dots, N\}$
- A **type distribution** F over $[0, \bar{\theta}] \subseteq \mathbb{R}$, which admits a continuous density f with $f(\theta_i) > 0$ for all $\theta_i \in [0, \bar{\theta}]$.
- For each $i \in \mathcal{N}$, an **allocation rule** $H_i : \mathbb{R}_+^N \rightarrow [0, 1]$ that specifies the probability that player i gets the item for every profile of N bids
- For each $i \in \mathcal{N}$, a **payment rule** $P_i : \mathbb{R}_+^N \rightarrow \mathbb{R}$ that specifies the expected payment of player i for every profile of N bids

At the start of the auction, each player i learns her own valuation θ_i . The valuations of different players are drawn i.i.d. from F , which is supported on the interval $[0, \bar{\theta}]$. Each player simultaneously submits a non-negative real number as her bid. When the profile of bids (s_1, \dots, s_N) is submitted, player i gets the item with probability $H_i(s_1, \dots, s_N)$ and pays $P_i(s_1, \dots, s_N)$ in expectation.

1.2 Some examples of (H, P) pairs. In lecture, we showed that a number of well-known auction formats – namely, first-price and second-price auction – can be written in terms of some (allocation rule, payment rule) pairs. Now, we turn to a number of unusual auctions to further illustrate the definition²¹.

- **Raffle.** Each player chooses how many raffle tickets to buy. Each ticket costs \$1. A winner is selected by drawing a raffle ticket at random. This corresponds to $H_i(s_1, \dots, s_N) = \frac{s_i}{\sum_k s_k}$, $P_i(s_1, \dots, s_N) = s_i$. Unlike the usual auction formats like first-price and second-price auctions, the allocation rule H_i involves **randomization** for almost all profiles of “bids”.
- **War of attrition.** A strategic territory is contested by two generals. Each general chooses how much resources to use in fighting for this territory. The general who commits more resources

²¹Some of these examples may seem to have nothing to do with auctions at a first glance, yet our definition of an auction in terms of (H, P) is general enough to apply to them. Here, as elsewhere in Economics, theory allows us to unify our modeling and understanding of seemingly disparate phenomena under a single framework.

destroys all of her opponents' forces and wins the territory, but suffers as much losses as the losing general. This corresponds to

$$H_i(s_1, s_2) = \begin{cases} 1 & \text{if } s_i > s_{-i} \\ 0 & \text{if } s_i < s_{-i} \\ 0.5 & \text{if } s_i = s_{-i} \end{cases}$$

and $P_i(s_1, s_2) = \min(s_1, s_2)$, so it is as if two bidders each submits a bid and **everyone** pays the losing bid.

- **All-pay auction.** Each player submits a bid and the highest bidder gets the item. Every player, win or lose, must **pay her own bid**. Here, $H_i(s_1, \dots, s_N)$ is the same as in first-price auction, but the payment rule is $P_i(s_1, \dots, s_N) = s_i$.

1.3 Auctions as private-value Bayesian games. Auctions form an **important class of examples in Bayesian games**. As defined above, an auction is a private value Bayesian game with a continuum of types for each player. Referring back to Definition 43, an auction $\mathcal{A} = \langle \mathcal{N}, F, [0, \bar{\theta}], (H_k)_{k \in \mathcal{N}}, (P_k)_{k \in \mathcal{N}} \rangle$ can be viewed as a Bayesian game $\mathcal{B} = \langle \mathcal{N}, (S_k)_{k \in \mathcal{N}}, (\Theta_k)_{k \in \mathcal{N}}, \mu, (u_k)_{k \in \mathcal{N}} \rangle$, where:

- \mathcal{N} is the set of players
- Player i 's action set is $S_i = \mathbb{R}_+$, interpreted as bids
- States of the world is $\Theta = [0, \bar{\theta}]^N$, where the i -th dimension is the valuation of player i
- The common prior μ on Θ is the product distribution on $[0, \bar{\theta}]^N$ derived from F
- Utility function u_i specifies $u_i(s_1, \dots, s_N, \theta) = \theta_i \cdot H_i(s_1, \dots, s_N) - P_i(s_1, \dots, s_N)$

As such, many terminologies from general Bayesian games carry over to auctions. A strategy of bidder i is a function $s_i : \Theta_i \rightarrow \mathbb{R}_+$, mapping i 's valuation to a nonnegative bid. A BNE in an auction is a strategy profile $(s_k^*)_{k \in \mathcal{N}}$ such that for each player i and valuation $\theta_i \in \Theta_i$,

$$s_i^*(\theta_i) \in \arg \max_{\hat{s}_i \in \mathbb{R}_+} \mathbb{E}_{\theta_{-i}} \left[\theta_i \cdot H_i(\hat{s}_i, s_{-i}^*(\theta_{-i})) - P_i(\hat{s}_i, s_{-i}^*(\theta_{-i})) \right]$$

As usual in a BNE, player i of type θ_i knows the **mapping** from opponent's types θ_{-i} to their actions $s_{-i}^*(\theta_{-i})$, i.e. how each opponent would bid as a function of their valuation, but she does **not** know opponents' **realized valuations**. She does know the **distribution** over opponents' valuations, so she can compute the expected payoff of playing different bids, with expectation²² taken over opponents' types.

2 Solving for Auction BNEs

Given an auction, here are two approaches for identifying some of its BNEs. But be warned: an auction may have multiple BNEs and the following methods may not find all of them.

2.1 Weakly dominant BNEs. The following holds in general for private-value Bayesian games.

²²This is analogous to Definition 46. However, in Definition 46 we spelled out a weighted sum over $\theta_{-i} \in \Theta_{-i}$ instead of writing an expectation. This was possible since we focused on the case of a finite Θ in that section.

Definition 49. In a private-value Bayesian game, a strategy $s_i : \Theta_i \rightarrow S_i$ is **weakly dominant** for i if for all $s_{-i} \in S_{-i}$ and all $\theta_i \in \Theta_i$,

$$s_i(\theta_i) \in \arg \max_{\hat{s}_i \in S_i} \{u_i(\hat{s}_i, s_{-i}, \theta_i)\}$$

Proposition 50. In a private-value Bayesian game, consider a strategy profile $(s_k^*)_{k \in \mathcal{N}}$ where for each $i \in \mathcal{N}$, s_i^* is weakly dominant for i . Then $(s_k^*)_{k \in \mathcal{N}}$ is a BNE.

Proof. For each $i \in \mathcal{N}$ and $\theta_i \in \Theta_i$, definition of weakly dominant strategy says

$$s_i^*(\theta_i) \in \arg \max_{\hat{s}_i \in S_i} \{u_i(\hat{s}_i, s_{-i}, \theta_i)\}$$

for every $s_{-i} \in S_{-i}$, so in particular $s_i^*(\theta_i)$ maximizes $\hat{s}_i \mapsto u_i(\hat{s}_i, s_{-i}^*(\theta_{-i}), \theta_i)$ for each $\theta_{-i} \in \Theta_{-i}$. But this means $s_i^*(\theta_i)$ must also maximize the expectation of $u_i(\hat{s}_i, s_{-i}^*(\theta_{-i}), \theta_i)$ taken over θ_{-i} ,

$$\hat{s}_i \mapsto \mathbb{E}_{\theta_{-i}} [u_i(\hat{s}_i, s_{-i}^*(\theta_{-i}), \theta_i)]$$

since it maximizes the integrand pointwise in θ_{-i} . □

As a result, if we can identify a weakly dominant strategy for each player in an auction, then a profile of such strategies forms a BNE.

Example 51. Consider a second-price auction **with a reserve price**. The seller sets reserve price $r \in \mathbb{R}_+$, then every bidder submits a bid simultaneously.

- If every bid is less than r , then no bidder gets the item and no one pays anything.
- If the highest bid is r or higher, then the highest bidder pays either the bid of the second highest bidder or r , whichever is larger. If several players tie for the highest bidder, then one of these high bidders is chosen uniformly at random, gets the item, and pays the second highest bid (which is equal to her own bid).

We argue that $s_i^*(\theta_i) = \theta_i$ is a weakly dominant strategy.

If $\theta_i < r$, then against any profile of opponent bids s_{-i} , $u_i(\theta_i, s_{-i}, \theta_i) = 0$ since i will never win the item from a bid less than the reserve price. Yet, any other bid can only get expected utility no larger than 0, since any bid that wins must be larger than r , which is larger than i 's valuation.

If $\theta_i \geq r$, then profiles of opponent bids s_{-i} may be classified into 3 cases.

Case 1: highest rival bid is $y \geq \theta_i$. Then $u_i(\theta_i, s_{-i}, \theta_i) = 0$, while any other bid either loses the item or wins the item at a price of y , which can only lead to non-positive payoffs in the event of winning.

Case 2: highest rival bid is $y \in [r, \theta_i)$. Then $u_i(\theta_i, s_{-i}, \theta_i) = u_i(\hat{s}_i, s_{-i}, \theta_i) = \theta_i - y > 0$ for any $\hat{s}_i \in (y, \infty)$, since all bids higher than y lead to winning the item at a price of y . Bidding y leads to an expected payoff no larger than $\frac{1}{2}(\theta_i - y)$ from tie-breaking, which is worse than $\theta_i - y$. Bidding less than y loses the item and gets 0 utility.

Case 3: highest rival bid is $y < r$. Then $u_i(\theta_i, s_{-i}, \theta_i) = u_i(\hat{s}_i, s_{-i}, \theta_i) = \theta_i - r > 0$ for any $\hat{s}_i \in [r, \infty)$. Bidding less than the reserve price r loses the item and gets 0 utility.

Therefore, we have verified that playing $s_i^*(\theta_i) = \theta_i$ is optimal for type θ_i , regardless of opponents' bid profile. This means bidding own valuation is weakly dominant. By Proposition 50, every player bidding own valuation is therefore a BNE. ◆

2.2 The FOC approach. In the BNE of an auction, fixing i and an interior valuation $\theta_i \in (0, \bar{\theta})$ we have:

$$s_i^*(\theta_i) \in \arg \max_{\hat{s}_i \in \mathbb{R}_+} \mathbb{E}_{\theta_{-i}} \left[\theta_i \cdot H_i(\hat{s}_i, s_{-i}^*(\theta_{-i})) - P_i(\hat{s}_i, s_{-i}^*(\theta_{-i})) \right] \quad (1)$$

so in particular,

$$\theta_i \in \arg \max_{\hat{\theta}_i \in \Theta_i} \mathbb{E}_{\theta_{-i}} \left[\theta_i \cdot H_i(s_i^*(\hat{\theta}_i), s_{-i}^*(\theta_{-i})) - P_i(s_i^*(\hat{\theta}_i), s_{-i}^*(\theta_{-i})) \right] \quad (2)$$

because (2) **restricts** the optimization problem in (1) to the domain of $s_i^*(\Theta_i) \subseteq \mathbb{R}_+$. Consider now the objective function of this second optimization problem,

$$\hat{\theta}_i \mapsto \mathbb{E}_{\theta_{-i}} \left[\theta_i \cdot H_i(s_i^*(\hat{\theta}_i), s_{-i}^*(\theta_{-i})) - P_i(s_i^*(\hat{\theta}_i), s_{-i}^*(\theta_{-i})) \right] \quad (3)$$

If it is differentiable (which will hold provided H_i, P_i , and the distribution F are “nice enough”) and $\theta_i \in (0, \bar{\theta})$, then first order condition (FOC) of optimization implies

$$\frac{d}{d\hat{\theta}_i} \left\{ \mathbb{E}_{\theta_{-i}} \left[\theta_i \cdot H_i(s_i^*(\hat{\theta}_i), s_{-i}^*(\theta_{-i})) - P_i(s_i^*(\hat{\theta}_i), s_{-i}^*(\theta_{-i})) \right] \right\} (\theta_i) = 0$$

In auctions without a weakly dominant strategy, sometimes this FOC can help identify a BNE by giving us a closed-form expression of $s_i^*(\theta_i)$ after manipulation.

Example 52. Consider a first-price auction with two bidders. The two bidders’ types are distributed i.i.d. with $\theta_i \sim \text{Uniform}[0, 1]$. Each bidder submits a nonnegative bid and whoever bids higher wins the item and pays her own bid. If there is a tie, then each bidder gets to buy the item at her bid with equal probability. It is known that this auction has a symmetric BNE where (i) $s_i^*(\theta_i)$ is differentiable, strictly increasing in θ_i ; (ii) the associated Equation (3) is differentiable. Find a closed-form expression for $s_i^*(\theta_i)$.

Solution: In the BNE $(s_i^*)_{i=1,2}$, the expected probability of P1 winning the item by playing the BNE strategy of type $\hat{\theta}_1$ is $\hat{\theta}_1$. This is because s_2^* is strictly increasing and symmetric to s_1^* , so that bidding $s_1^*(\hat{\theta}_1)$ wins is exactly when $\theta_2 < \hat{\theta}_1$, which happens with probability $\hat{\theta}_1$ since $\theta_2 \sim \text{Uniform}[0, 1]$. At the same time, the expected payment for submitting the BNE bid of type θ_i is $\hat{\theta}_i \cdot s_i^*(\hat{\theta}_i)$, because bidding $s_i^*(\hat{\theta}_i)$ wins with probability $\hat{\theta}_i$ and pays $s_i^*(\hat{\theta}_i)$ in the event of winning. The relevant optimization problem is therefore

$$\max_{\hat{\theta}_i \in [0, \bar{\theta}]} \theta_i \cdot \hat{\theta}_i - \hat{\theta}_i \cdot s_i^*(\hat{\theta}_i)$$

FOC implies that $\theta_i - s_i^*(\theta_i) - \theta_i \cdot (s_i^*)'(\theta_i) = 0$. This is a first-order differential equation in θ_i that holds for $\theta_i \in (0, \bar{\theta})$. We may rearrange it to get

$$\theta_i = s_i^*(\theta_i) + \theta_i \cdot (s_i^*)'(\theta_i)$$

Integrating both sides,

$$\frac{1}{2} \theta_i^2 + C = \theta_i \cdot s_i^*(\theta_i) \quad (4)$$

Evaluating at $\theta_i = 0$ recovers²³ the constant of integration $C = 0$. Therefore, $s_i^*(\theta_i) = \frac{1}{2}\theta_i$ is the desired symmetric BNE. \blacklozenge

3 Revenue Equivalence: Theorem and Applications

3.1 The revenue-equivalence theorem. While you may be familiar with statements like “first-price auction and second-price auction are revenue equivalent” before taking this course, it is important to gain a more precise understanding of the revenue-equivalence theorem (RET). To see how a cursory reading of the RET might lead you astray, consider the **asymmetric** second-price auction BNE from lecture, where bidder 1 always bids $\bar{\theta}$ and everyone else always bids 0, regardless of their types. The seller’s expected revenue is 0!

Strictly speaking, RET is not a statement comparing two auction formats, but a statement **comparing two equilibria of two auction formats**. “Revenue” is an equilibrium property and an auction game might admit multiple BNEs with different expected revenues.

So let a BNE $(s_k^*)_{k \in \mathcal{N}}$ of auction game \mathcal{A} be given²⁴. Let us define two functions $G_i, R_i : \Theta_i \rightarrow \mathbb{R}$ for each player i , so that $G_i(\hat{\theta}_i)$ and $R_i(\hat{\theta}_i)$ give the **expected probability of winning** and **expected payment** when bidding as though valuation is $\hat{\theta}_i$.

$$G_i(\hat{\theta}_i) := \mathbb{E}_{\theta_{-i}} \left[H_i(s_i^*(\hat{\theta}_i), s_{-i}^*(\theta_{-i})) \right]$$

$$R_i(\hat{\theta}_i) := \mathbb{E}_{\theta_{-i}} \left[P_i(s_i^*(\hat{\theta}_i), s_{-i}^*(\theta_{-i})) \right]$$

The expectations are taken over opponents’ types. Importantly, G_i and R_i are dependent on the BNE $(s_k^*)_{k \in \mathcal{N}}$.²⁵ If we consider a different BNE of the same auction, then we will have a different pair \tilde{G}_i, \tilde{R}_i .

To illustrate, consider the symmetric BNE we derived in the two-player auction in Example 52, where $s_i^*(\theta_i) = \frac{\theta_i}{2}$. It should be intuitively clear that $G_i(\hat{\theta}_i) = \hat{\theta}_i$ and $R_i(\hat{\theta}_i) = \hat{\theta}_i \cdot \frac{\hat{\theta}_i}{2} = \frac{\hat{\theta}_i^2}{2}$. We can also derive these expressions from definition,

$$G_1(\hat{\theta}_1) = \int_{\theta_2=0}^{\theta_2=1} H_1 \left(\frac{\hat{\theta}_1}{2}, \frac{\theta_2}{2} \right) d\theta_2 = \int_0^{\hat{\theta}_1} 1 d\theta_2 + \int_{\hat{\theta}_1}^1 0 d\theta_2 = \hat{\theta}_1$$

$$R_1(\hat{\theta}_1) = \int_{\theta_2=0}^{\theta_2=1} R_1 \left(\frac{\hat{\theta}_1}{2}, \frac{\theta_2}{2} \right) d\theta_2 = \int_0^{\hat{\theta}_1} \frac{\hat{\theta}_1}{2} d\theta_2 + \int_{\hat{\theta}_1}^1 0 d\theta_2 = \frac{\hat{\theta}_1^2}{2}$$

As we have seen in lecture, the celebrated RET is just a corollary of the following nameless result:

Proposition 53. (*Nameless Result*) Fix a BNE $(s_k^*)_{k \in \mathcal{N}}$ of the auction game. Under regularity conditions, $R_i(\theta_i) = \int_0^{\theta_i} x G_i'(x) dx + R_i(0)$ for all i and θ_i .

Proof. See lecture. □

²³Even though the FOC only applies for interior $\theta_i \in (0, \bar{\theta})$, continuity of s_i^* implies Equation (4) holds even at the boundary points. (This is sometimes called “value matching”.)

²⁴We can in fact define G_i and R_i for any arbitrary profile of strategies $(s_k^*)_{k \in \mathcal{N}}$, without imposing that it is a BNE. However, Proposition 53 only holds when $(s_k^*)_{k \in \mathcal{N}}$ is a BNE.

²⁵So they more correctly notated as $G_i^{(s_k^*)_{k \in \mathcal{N}}}, R_i^{(s_k^*)_{k \in \mathcal{N}}}$, but such notation is too cumbersome.

This result expresses the expected payment of an **arbitrary type** of player i in a BNE as a function of: (i) expected payment of the **lowest type** of player i in this BNE; (ii) the **expected probabilities of winning** for various types of player i in this BNE. It then follows that:

Corollary 54. (*Revenue-equivalence theorem*) Under regularity conditions, for two BNEs of two auctions such that $G_i(\theta_i) = G_i^\circ(\theta_i)$ for all i, θ_i and $R_i(0) = R_i^\circ(0)$ for all i , we have $R_i(\theta_i) = R_i^\circ(\theta_i)$ for all i, θ_i .

This follows directly from Proposition 53. Since in BNE the expected payment of an arbitrary type is entirely determined by the winning probabilities of different types and the expected payment of the lowest type, two BNEs where these two objects match must have the same expected payment for all types.

Here are two examples where RET is not applicable due to G_i and G_i° not matching up for two BNEs.

Example 55. In a second-price auction, the asymmetric BNE does not satisfy the conditions of RET when compared to the symmetric BNE of bidding own valuation. In the asymmetric equilibrium, $G_i(\theta_i) = 0$ for all $i \neq 1$, $\theta_i \in [0, \bar{\theta}]$, since bidders other than P1 never win. Therefore, we cannot conclude from RET that these two BNEs yield the same expected revenue. (In fact, they do not.) ♦

Example 56. In Example 51, we showed that bidding own valuation is a BNE in a second-price auction with reserve price. When reserve price is $r > 0$, this BNE does not satisfy the conditions of RET when compared to the BNE of bidding own valuation in a second-price auction **without** reserve price. In the former BNE, $G_i(\theta_i) = 0$ for any $\theta_i \in (0, r)$, whereas in the latter BNE these types have a strictly positive probability of winning the item. Therefore, we cannot conclude from RET that these two BNEs in two auction formats yield the same expected revenue. (In fact, different reserve prices may lead to different expected revenues.) ♦

3.2 Using RET to solve auctions. Sometimes, we can use RET to derive a closed-form expression of the BNE strategy profile (s_i^*).

Example 57. As in Example 52, consider a first-price auction with two bidder whose valuations are i.i.d. with $\theta_i \sim \text{Uniform}[0, 1]$. Assume this auction has a symmetric BNE where $s_i^*(\theta_i)$ strictly increasing in θ_i . Then this BNE is revenue equivalent to the BNE of second price auction where each player bids own valuation. To see this, since both BNEs feature strategies strictly increasing in type, i of type θ_i wins precisely when player $-i$ has a type $\theta_{-i} < \theta_i$. That is to say, $G_i(\theta_i) = \theta_i = G_i^\circ(\theta_i)$. At the same time, the expected payment from type 0 is 0 in both BNEs – in particular, the type 0 bidder in first-price auction never wins since bids are strictly increasing in type, so never pays anything.

But in the bid-own-valuation BNE of the second price auction, $R_i^\circ(\theta_i) = \theta_i \cdot (\theta_i/2)$, where θ_i is probability of being the highest bidder and $\theta_i/2$ is expected rival bid in the event of winning. By RET, $R_i(\theta_i) = \theta_i \cdot (\theta_i/2)$ also. In first-price auction, i pays own bid $s_i^*(\theta_i)$ whenever she wins, which happens with probability θ_i . Hence $s_i^*(\theta_i) = R_i(\theta_i)/\theta_i = \frac{\theta_i}{2}$. This is the same as what we found using FOC in Example 52. ♦

While in the above example we used RET to verify a result we already knew from FOC, RET can also be used in lieu of FOC to find BNEs. This can be particularly helpful when the differential equation from the FOC approach is harder to solve.

Example 58. (December 2012 Final Exam) Suppose there are two risk-neutral potential buyers of an indivisible good. It is common knowledge that each buyer i 's valuation is drawn independently

from the same distribution on $[0, 1]$ with distribution function $F(\theta) = \theta^3$, but the realizations of the θ_i 's are private information. Calculate the expected payment $R_i(\theta_i)$ that a buyer with reservation price θ_i makes in the unique symmetric equilibrium of a **second-price** auction. Then, using the revenue equivalence theorem, find the equilibrium bid function in a **first-price** auction in the same setting.

Solution: In the second-price auction, it is a weakly dominant BNE to bid own valuation. In this BNE,

$$R_i(\theta_i) = \int_0^{\theta_i} \theta_j \cdot f(\theta_j) d\theta_j = \int_0^{\theta_i} \theta_j \cdot (3\theta_j^2) d\theta_j = \left[\frac{3}{4}\theta_j^4 \right]_{\theta_j=0}^{\theta_j=\theta_i} = \frac{3}{4}\theta_i^4$$

This symmetric BNE of second-price auction is revenue equivalent to any BNE in first-price auction where bid increases strictly with own type. This is because in these BNEs, $G_i(\theta_i) = G_i^\circ(\theta_i) = \theta_i^3$ (since i of type θ_i wins exactly when $-i$ is of type lower than θ_i) and $R_i(0) = R_i^\circ(0) = 0$ (since type 0 never wins, so never pays). But in first-price auction, $R_i^\circ(\theta_i) = s_i^*(\theta_i) \cdot G_i^\circ(\theta_i)$, so then $s_i^*(\theta_i) = \left(\frac{3}{4}\theta_i^4\right) / \theta_i^3 = \frac{3}{4}\theta_i$. ◆

(1) Subgame-perfect eqm; (2) Infinite-horizon games and one-shot deviation; (3) Bargaining

TF: Kevin He (he02@fas.harvard.edu)

1 Subgame-Perfect Equilibrium

1.1 Nash equilibrium in finite-horizon games. Recall the definition of a finite-horizon extensive-form game and the definition of a strategy in extensive-form games from Section 1.

Definition 59. A finite-horizon extensive-form game Γ has the following components:

- A finite-depth **tree** with vertices V and terminal vertices $Z \subseteq V$.
- A set of **players** $\mathcal{N} = \{1, 2, \dots, N\}$.
- A **player function** $J : V \setminus Z \rightarrow \mathcal{N} \cup \{c\}$.
- A set of available **moves** $M_{j,v}$ for each $v \in J^{-1}(j)$, $j \in \mathcal{N}$. Each move in $M_{j,v}$ is associated with a unique child of v in the tree.
- A **probability distribution** $f(\cdot|v)$ over v 's children for each $v \in J^{-1}(c)$.
- A (Bernoulli) **utility function** $u_j : Z \rightarrow \mathbb{R}$ for each $j \in \mathcal{N}$.
- An **information partition** \mathcal{I}_j of $J^{-1}(j)$ for each $j \in \mathcal{N}$, whose elements are **information sets** $I_j \in \mathcal{I}_j$. It is required that $v, v' \in I_j \Rightarrow M_{j,v} = M_{j,v'}$.

Definition 60. In an extensive-form game, a **pure strategy** for player j is a function $s_j : \mathcal{I}_j \rightarrow \cup_{v \in J^{-1}(j)} M_{j,v}$, so that $s_j(I_j) \in M_{j,I_j}$ for each $I_j \in \mathcal{I}_j$. Write S_j for the set of all pure strategies of player j .

A strategy profile $(s_k)_{k \in \mathcal{N}}$ induces a distribution over terminal vertices Z , which we write as $p(\cdot|(s_k)) \in \Delta(Z)$. Hence we may define $U_i : S \rightarrow \mathbb{R}$ where

$$U_i(s_i, s_{-i}) := \mathbb{E}_{z \sim p(\cdot|(s_k))} [u_i(z)]$$

That is, the extensive-game payoff to player i is defined as her expected utility from terminal vertices, according to her Bernoulli utility u_i and the distribution over terminal vertices induced by the strategy profile. A Nash equilibrium in extensive-form game is defined in the natural way: a strategy profile where no player has a profitable unilateral deviation, where potential deviations are different extensive-form game strategies.

Definition 61. A **Nash equilibrium** in finite-horizon extensive-form game is a strategy profile $(s_k^*)_{k \in \mathcal{N}}$ where $U_i(s_i^*, s_{-i}^*) \geq U_i(\hat{s}_i, s_{-i}^*)$ for all $\hat{s}_i \in S_i$.

²⁶Figure 9 is adapted from Osborne and Rubinstein (1994): A Course in Game Theory [5].

Example 62. Figure 3 shows the game tree of an **ultimatum game**, Γ . It models an interaction between P1 and P2 who must split **two identical, indivisible items**. P1 proposes an allocation. Then, P2 accepts or rejects the allocation. If the allocation is accepted, it is implemented. If it is rejected, then neither player gets any of the good.

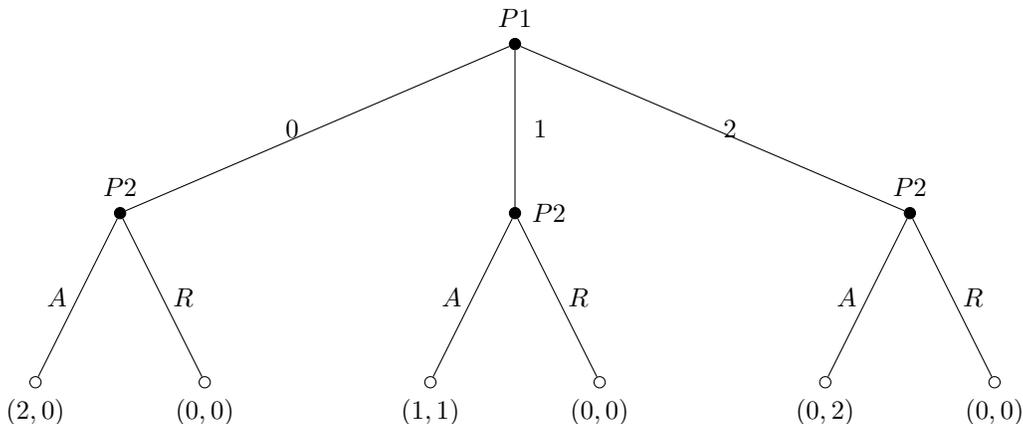


Figure 3: A game tree representation of the ultimatum game.

P1 moves at the root of the game tree. Her move set at the root is $\{0, 1, 2\}$, which correspond to giving 0, 1, 2 units of the good to P2. Regardless of which action P1 chooses, the game moves to a vertex where it is P2’s turn to play. His move set at each of his three decision vertices is $\{A, R\}$, corresponding to accepting and rejecting the proposed allocation.

The strategy profile $s_1^*(\emptyset) = 2, s_2^*(0) = s_2^*(1) = R, s_2^*(2) = A$ is a Nash equilibrium. Certainly P2 has no profitable deviations since $U_2(s_1^*, s_2^*) = 2$, which is the highest he can hope to get in this game. As for P1, she also has no profitable unilateral deviations, since offering 0 or 1 to P2 leads to rejection and no change in her payoff. By the way, this is why we insist that a strategy in an extensive-form game specifies what each player would do at each information set, even those information sets that are not reached when the game is played. What P2 **would have done** if offered 0 or 1 is crucial in sustaining a Nash equilibrium in which P1 offers 2. \blacklozenge

1.2 Subgames and subgame-perfect equilibrium. In some sense, the NE of Example 62 is artificially sustained by a **non-credible threat**. P2 threatens to reject the proposal if P1 offers 1, despite the fact that he has no incentive to carry out the threat if P1 really makes this offer. This threat does not harm P2’s payoff in the game Γ , since P2’s un-optimized decision vertex is never reached when the strategy profile (s_1^*, s_2^*) is played – it is “off the equilibrium path”.

Whether or not strategy profiles like (s_1^*, s_2^*) make sense as predictions of the game’s outcome depends on the availability of commitment devices. If at the start of the game P2 could somehow make it impossible for himself to accept the even-split offer, then this NE is a reasonable prediction. In the absence of such commitment devices, however, we should seek out a refinement of NE in extensive-form games to rule out such non-credible threats.

We begin with the definition of a subgame.

Definition 63. In a finite-horizon extensive-form game Γ , any $x \in V \setminus Z$ such that every information set is either entirely contained in the subtree starting at x or entirely outside of it defines a subgame, $\Gamma(x)$. This subgame is an extensive-form game inherits the payoffs, moves, and information structure of the original game Γ in the natural way.

Example 64. The ultimatum game in Example 62 has 4 subgames: $\Gamma(\emptyset)$ (which is just Γ), as well as $\Gamma(0)$, $\Gamma(1)$, $\Gamma(2)$. We sometimes call $\Gamma(\emptyset)$ the **improper subgame** and $\Gamma(1), \Gamma(2), \Gamma(3)$ the **proper subgames**. \blacklozenge

Definition 65. A strategy profile $(s_k^*)_{k \in \mathcal{N}}$ of Γ is called a **subgame-perfect equilibrium (SPE)** if for every subgame $\Gamma(x)$, $(s_k^*)_{k \in \mathcal{N}}$ restricted to $\Gamma(x)$ forms a Nash equilibrium in $\Gamma(x)$.

We know that $\Gamma(\emptyset)$ is always a subgame of Γ since the root of the game tree is always in a singleton information set. Therefore, every SPE is an NE, but not conversely.

Example 66. The NE (s_1^*, s_2^*) from Example 62 is not an SPE, since (s_1^*, s_2^*) restricted to the subgame $\Gamma(1)$ is not an NE. However, the following is an SPE: $s_1^\circ(\emptyset) = 1$, $s_2^\circ(0) = R$, $s_2^\circ(1) = s_2^\circ(2) = A$. It is easy to see that restricting (s_1°, s_2°) to each of the subgames $\Gamma(0), \Gamma(1), \Gamma(2)$ forms an NE. Furthermore, (s_1°, s_2°) is a NE in $\Gamma(\emptyset) = \Gamma$. P1 gets $U_1(s_1^\circ, s_2^\circ) = 1$ under this strategy profile, while offering 0 leads to rejection and a payoff of 0, offering 2 leads to acceptance but again a payoff of 0. For P2, changing $s_2^\circ(0)$ and $s_2^\circ(2)$ do not change payoff in Γ , since these two vertices are never reached. Changing $s_2^\circ(1)$ from A to R hurts payoff. \blacklozenge

1.3 Backwards induction. **Backwards induction** is an algorithm for finding an SPE in a finite-horizon extensive-form game of perfect information. The idea is to successively replace subgames with terminal vertices corresponding to SPE payoffs of the deleted subgames.

Start with a non-terminal vertex **furthest away** from the root of the game, say v . Since we have picked the deepest non-terminal vertex, all of $J(v)$'s moves at this vertex must lead to terminal vertices. Choose one of $J(v)$'s moves, m^* , that maximizes her payoff in $\Gamma(v)$, then replace the subgame $\Gamma(v)$ with the terminal vertex corresponding to m^* . Repeat this procedure, working backwards from the vertices further away from the root of the game. Eventually, the game tree will be reduced to a single terminal vertex, whose payoff will be an SPE payoff of the extensive-form game, while the moves chosen throughout the deletion process will form a SPE strategy profile.

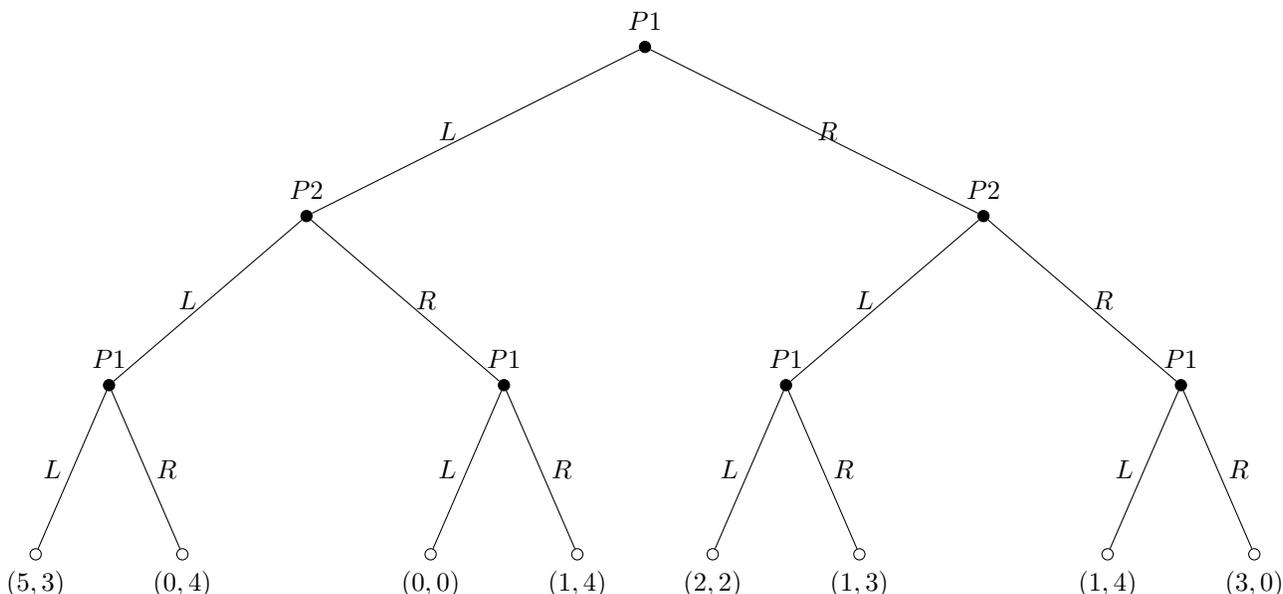


Figure 4: A extensive-form game in the game tree representation.

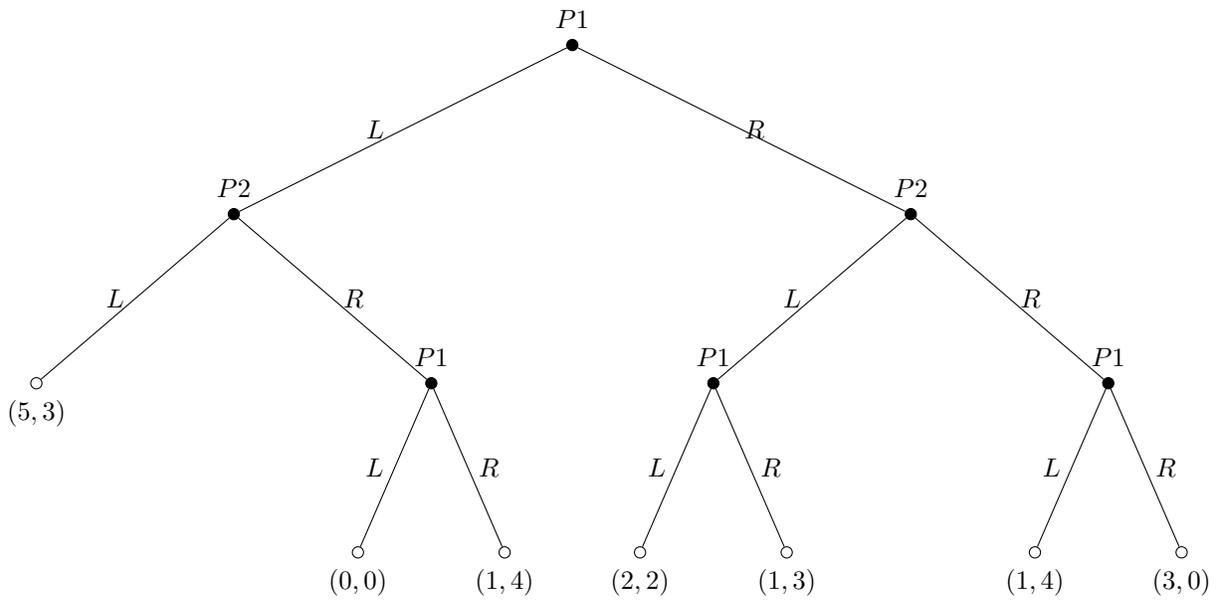


Figure 5: Backwards induction replaces subgames with terminal nodes associated with the SPE payoffs in those subgames. Here is the resulting game tree after one step of backwards induction.

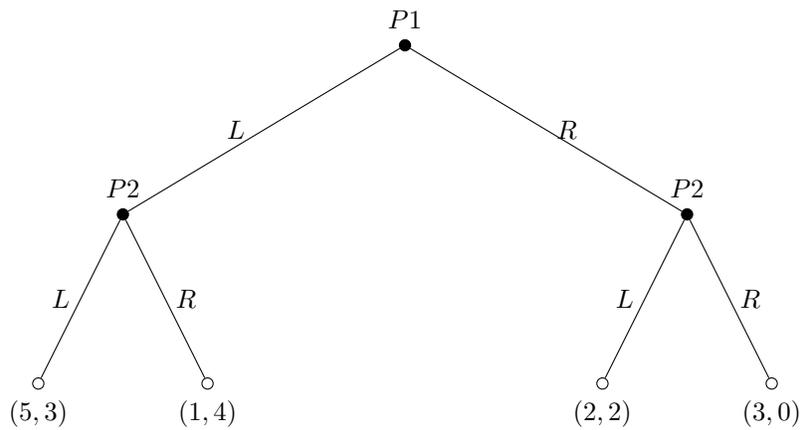


Figure 6: Backwards induction in progress. All nodes at depth 3 in the original tree have been eliminated.

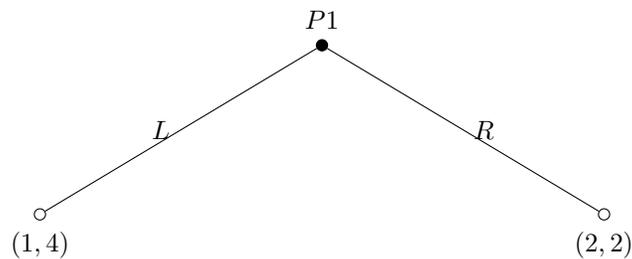


Figure 7: Backwards induction in progress. Only nodes with depth 1 remain.

$$\begin{array}{c} \circ \\ (2,2) \end{array}$$

Figure 8: Backwards induction finds the unique SPE payoff in this game.

If $u_i(z) \neq u_i(z')$ for every i and $z, z' \in Z$ with $z \neq z'$, then backwards induction finds the unique SPE of the extensive-form game. Otherwise, the game may have multiple SPEs and backwards induction may involve choosing between several indifferent moves. Depending on the moves chosen, backwards induction may lead to different SPEs.

————— 2 Infinite-Horizon Games and One-Shot Deviation —————

2.1 Infinite-horizon games. So far, we have only dealt with finite-horizon games. These games are represented by finite-depth game trees and must end within M turns for some $M \in \mathbb{N}$. But games such as the Rubinstein-Stahl bargaining are not finite horizon, for players could reject each other's offers forever. We modify Definition 59 to accommodate such infinite-horizon games. For simplicity, we assume the game has perfect information and no chance moves.

Definition 67. An **extensive-form game with perfect information and no chance moves** has the following components:

- A **possibly infinite-depth tree** with vertices V and terminal vertices $Z \subseteq V$.
- A set of **players** $\mathcal{N} = \{1, 2, \dots, N\}$.
- A **player function** $J : V \setminus Z \rightarrow \mathcal{N}$.
- A set of available **moves** $M_{j,v}$ for each $v \in J^{-1}(j)$, $j \in \mathcal{N}$. Each move in $M_{j,v}$ is associated with a unique child of v in the tree.
- A (Bernoulli) **utility function** $u_j : Z \cup H^\infty \rightarrow \mathbb{R}$ for each $j \in \mathcal{N}$, where H^∞ refers to the set of all infinite length paths.

When an infinite-horizon game is played, it might end at a terminal vertex (such as when one player accepts the other's offer in the bargaining game), or it might **never reach a terminal vertex** (such as when both players use a strategy involving never accepting any offer in the bargaining game). Therefore, each player must have a preference not only over the set of terminal vertices, but also **over the set of infinite histories**. In the bargaining game, for instance, it is specified that $u_j(h) = 0$ for any $h \in H^\infty$, $j = 1, 2$, that is to say every infinite history in the game tree (i.e. never reaching an agreement) gives 0 utility to each player.

Many definitions from finite-horizon extensive-form games directly translate into the infinite-horizon setting. For instance, any nonterminal vertex x in the perfect-information infinite-horizon game defines a subgame $\Gamma(x)$. NE is defined in the obvious way, taking into account distribution over both terminal vertices and infinite histories induced by a strategy profile. SPE is still defined as those strategy profiles that form an NE when restricted to each of Γ 's (possibly infinitely many) subgames.

2.2 One-shot deviation principle. It is often difficult to verify directly from definition whether a given strategy profile forms an SPE in an infinite-horizon game. Indeed, given an SPE candidate $(s_k^*)_{k \in \mathcal{N}}$ of game Γ , we would have to consider each subgame $\Gamma(x)$, which is potentially an infinite-horizon

extensive-form game of its own right, and ask whether player i can improve her payoff in $\Gamma(x)$ by choosing a different extensive-form game strategy \hat{s}_i , modifying some or all of her choices at various vertices in $J^{-1}(i)$ relative to s_i^* . This is not an easy task since i 's set of strategies in $\Gamma(x)$ is a very rich set. The one-shot deviation principle says for extensive-form games satisfying certain regularity conditions, we need only check that i does not have a profitable deviation amongst a **very restricted set** of strategies in each subgame $\Gamma(x)$, namely those that differ from s_i^* only at x .

Theorem 68. (One-shot deviation principle) *If Γ is consistent and continuous at infinity, then a strategy profile $(s_k^*)_{k \in \mathcal{N}}$ is an SPE of Γ iff for every player i , every subgame $\Gamma(x)$ with $J(x) = i$, and every strategy \hat{s}_i of $\Gamma(x)$ such that $s_i^*(v) = \hat{s}_i(v)$ at every $v \in J^{-1}(i)$ except possibly at $v = x$,*

$$U_i^{\Gamma(x)}(s_i^*, s_{-i}^*) \geq U_i^{\Gamma(x)}(\hat{s}_i, s_{-i}^*)$$

where $U_i^{\Gamma(x)}$ denotes the payoff of i in subgame $\Gamma(x)$.

Assuming consistency and continuity at infinity, to verify whether $(s_k^*)_{k \in \mathcal{N}}$ is an SPE we need only examine each subgame $\Gamma(x)$ and consider whether $J(x)$ can improve her payoff in $\Gamma(x)$ by changing her move only at x (a “**one-shot deviation**”). These two regularity conditions are satisfied by all finite-horizon extensive-form games, as well as all infinite-horizon games studied in lecture, including bargaining and repeated games.

3 Rubinstein-Stahl Bargaining

3.1 Bargaining as an extensive-form game. The **Rubinstein-Stahl bargaining game**, or simply “**bargaining game**”²⁷ for short, is an important example of infinite-horizon, perfect-information extensive-form game. It is comparable to the ultimatum game from Example 62, but with two important differences: (i) the game is **infinite-horizon**, so that first rejection does not end the game. Instead, players alternate in making offers; (ii) the good that players bargain over is assumed **infinitely divisible**, so that any allocation of the form $(x, 1 - x)$ for $x \in [0, 1]$ is feasible.

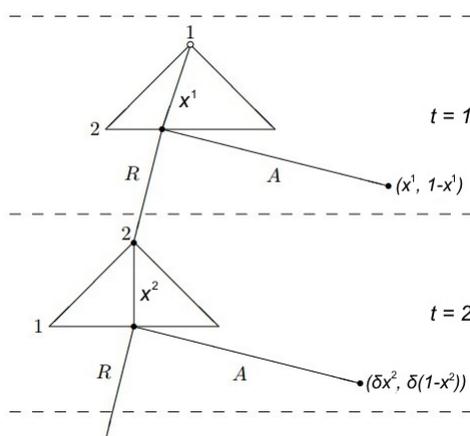


Figure 9: Part of the bargaining game tree, showing only some of the branches in the first two periods. The root \emptyset has (uncountably) infinitely many children of the form $(x^1, 1 - x^1)$ for $x^1 \in [0, 1]$. At each such child, P2 may play R or A. Playing A leads to a terminal node with payoffs $(x^1, 1 - x^1)$, while playing R continues the game with P2 to make the next offer.

²⁷Not to be confused with axiomatic Nash bargaining, which you will study in 2010b.

Let's think about what a strategy in the bargaining game looks like. Figure 9 shows a sketch of the bargaining game tree. P1's strategy specifies $s_1(\emptyset)$, that is to say what P1 will offer at the start of the game. For each $x^1 \in [0, 1]$, P2's strategy specifies $s_2((x^1, 1 - x^1)) \in \{A, R\}$, that is whether he accepts or rejects a period 1 offer of $(x^1, 1 - x^1)$. In addition, P2's strategy must also specify $s_2((x^1, 1 - x^1), R)$ for each $x^1 \in [0, 1]$, that is what he offers in period $t = 2$ if he rejected P1's offer in $t = 1$.²⁸ This offer could in principle depend on what P1 offered in period $t = 1$. Now for every $x^1, x^2 \in [0, 1]$, P1's strategy must specify $s_1((x^1, 1 - x^1), R, (x^2, 1 - x^2)) \in \{A, R\}$, which could in principle depend on what she herself offered in period $t = 1$, as well as what P2 offered in the current period, $(x^2, 1 - x^2)$.

3.2 Asymmetric bargaining power. Here is a modified version of the bargaining game that introduces **asymmetric bargaining power** between the two players.

Example 69. P1 gets to make offers in periods $3k + 1$ and $3k + 2$ for $k \in \mathbb{N}$, while P2 gets to make offers in periods $3k + 3$ for $k \in \mathbb{N}$. As in the usual bargaining game, reaching an agreement of $(x, 1 - x)$ in period t yields the payoff profile $(\delta^{t-1} \cdot x, \delta^{t-1} \cdot (1 - x))$. If the players never reach an agreement, then payoffs are $(0, 0)$.

Consider the following strategy profile: whenever P1 makes an offer in period $3k + 1$, she offers $(\frac{1+\delta}{1+\delta+\delta^2}, \frac{\delta^2}{1+\delta+\delta^2})$. Whenever P1 makes an offer in period $3k + 2$, she offers $(\frac{1+\delta^2}{1+\delta+\delta^2}, \frac{\delta}{1+\delta+\delta^2})$. Whenever P2 makes an offer, he offers $(\frac{\delta+\delta^2}{1+\delta+\delta^2}, \frac{1}{1+\delta+\delta^2})$. Whenever P2 responds to an offer in period $3k + 1$, he accepts iff he gets at least $\frac{\delta^2}{1+\delta+\delta^2}$. Whenever P2 responds to an offer in period $3k + 2$, he accepts iff he gets at least $\frac{\delta}{1+\delta+\delta^2}$. Whenever P1 responds to an offer, she accepts iff she gets at least $\frac{\delta+\delta^2}{1+\delta+\delta^2}$. You may verify that this verbal description indeed defines a strategy profile that plays a valid move at every non-terminal node of the bargaining game tree.

We use the **one-shot deviation principle** to verify that this strategy profile is SPE. By the principle, we need only ensure that in each subgame, the player to move at the root of the subgame cannot gain by changing her move only at the root. Subgames of this bargaining game may be classified into **six families**:

- Subgame starting with P1 making an offer in period $3k + 1$.
- Subgame starting with P1 making an offer in period $3k + 2$.
- Subgame starting with P2 making an offer in period $3k + 3$.
- Subgame starting with P2 responding to an offer $(x, 1 - x)$ in period $3k + 1$.
- Subgame starting with P2 responding to an offer $(x, 1 - x)$ in period $3k + 2$.
- Subgame starting with P1 responding to an offer $(x, 1 - x)$ in period $3k + 3$.

We consider these six families one by one, showing in no subgame is there a profitable one-shot deviation. By the one-shot deviation principle, this shows the strategy profile is an SPE.

(1) Subgame starting with P1 making an offer in period $3k + 1$. Not deviating gives P1 $\delta^{(3k)} \cdot \frac{1+\delta}{1+\delta+\delta^2}$. Offering P2 more than $\frac{\delta^2}{1+\delta+\delta^2}$ leads to acceptance but yields strictly less utility to P1.

²⁸Remember, an extensive-form game strategy for j is a complete contingency plan that specifies a valid move at any vertex in the game tree where it is j 's turn to play, even those vertices that would never be reached due to how j plays in previous rounds. Even if P2's strategy specifies accepting every offer from P1 in $t = 1$, P2 still needs to specify what he would do after a history of the form $((x^1, 1 - x^1), R)$ for each $x^1 \in [0, 1]$.

Offering P2 less than $\frac{\delta^2}{1+\delta+\delta^2}$ leads to rejection. In the next period, P1 will offer herself $\frac{1+\delta}{1+\delta+\delta^2}$, which P2 will accept. Therefore, this deviation gives P1 utility $\delta^{(3k+1)} \cdot \frac{1+\delta}{1+\delta+\delta^2} < \delta^{(3k)} \cdot \frac{1+\delta}{1+\delta+\delta^2}$. So we see P1 has no profitable one-shot deviation at the start of this subgame.

(2) Subgame starting with P1 making an offer in period $3k + 2$. Not deviating gives P1 $\delta^{(3k+1)} \cdot \frac{1+\delta}{1+\delta+\delta^2}$. Offering P2 more than $\frac{\delta}{1+\delta+\delta^2}$ leads to acceptance but yields strictly less utility to P1. Offering P2 less than $\frac{\delta}{1+\delta+\delta^2}$ leads to rejection. In the next period, P2 will offer P1 $\frac{\delta+\delta^2}{1+\delta+\delta^2}$, which P1 will accept. Therefore, this deviation gives P1 utility $\delta^{(3k+2)} \cdot \frac{\delta+\delta^2}{1+\delta+\delta^2} < \delta^{(3k+1)} \cdot \frac{1+\delta}{1+\delta+\delta^2}$. So we see P1 has no profitable one-shot deviation at the start of this subgame.

(3) Subgame starting with P2 making an offer in period $3k + 3$. Not deviating gives P2 $\delta^{(3k+2)} \cdot \frac{1}{1+\delta+\delta^2}$. Offering P1 more than $\frac{\delta+\delta^2}{1+\delta+\delta^2}$ leads to acceptance but yields strictly less utility to P2. Offering P1 less than $\frac{\delta+\delta^2}{1+\delta+\delta^2}$ leads to rejection. In the next period, P1 will offer P2 $\frac{\delta^2}{1+\delta+\delta^2}$, which P2 will accept. Therefore, this deviation gives P2 utility $\delta^{(3k+3)} \cdot \frac{\delta^2}{1+\delta+\delta^2} < \delta^{(3k+2)} \cdot \frac{1}{1+\delta+\delta^2}$. So we see P2 has no profitable one-shot deviation at the start of this subgame.

(4) Subgame starting with P2 responding to an offer $(x, 1 - x)$ in period $3k + 1$.

If $1 - x < \frac{\delta^2}{1+\delta+\delta^2}$, the strategy for P2 prescribes rejection. In the next period, P1 will offer P2 $\frac{\delta}{1+\delta+\delta^2}$ which P2 will accept, giving P2 a utility of $\delta^{(3k+1)} \cdot \frac{\delta}{1+\delta+\delta^2}$. On the other hand, the deviation of accepting $1 - x$ in the current period gives utility $\delta^{(3k)} \cdot (1 - x) < \delta^{(3k)} \cdot \frac{\delta^2}{1+\delta+\delta^2} = \delta^{(3k+1)} \cdot \frac{\delta}{1+\delta+\delta^2}$. So P2 has no profitable one-shot deviation.

If $1 - x \geq \frac{\delta^2}{1+\delta+\delta^2}$, the strategy for P2 prescribes acceptance, giving P2 a utility of $\delta^{(3k)} \cdot (1 - x) \geq \delta^{(3k)} \cdot \frac{\delta^2}{1+\delta+\delta^2}$. If P2 rejects instead, then in the next period P1 will offer P2 $\frac{\delta}{1+\delta+\delta^2}$ which P2 will accept, giving P2 a utility of $\delta^{(3k+1)} \cdot \frac{\delta}{1+\delta+\delta^2} \leq \delta^{(3k)} \cdot \frac{\delta^2}{1+\delta+\delta^2}$. So P2 has no profitable one-shot deviation.

(5) Subgame starting with P2 responding to an offer $(x, 1 - x)$ in period $3k + 2$.

If $1 - x < \frac{\delta}{1+\delta+\delta^2}$, the strategy for P2 prescribes rejection. In the next period, P2 will offer himself $\frac{1}{1+\delta+\delta^2}$, which P1 will accept, giving P2 a utility of $\delta^{(3k+2)} \cdot \frac{1}{1+\delta+\delta^2}$. On the other hand, the deviation of accepting $1 - x$ in the current period gives P2 utility $\delta^{(3k+1)} \cdot (1 - x) < \delta^{(3k+1)} \cdot \frac{\delta}{1+\delta+\delta^2} = \delta^{(3k+2)} \cdot \frac{1}{1+\delta+\delta^2}$. So P2 has no profitable one-shot deviation.

If $1 - x \geq \frac{\delta}{1+\delta+\delta^2}$, the strategy for P2 prescribes acceptance, giving P2 utility of $\delta^{(3k+1)} \cdot (1 - x) \geq \delta^{(3k+1)} \cdot \frac{\delta}{1+\delta+\delta^2}$. If P2 rejects instead, then in the next period P2 will offer himself $\frac{1}{1+\delta+\delta^2}$, which P1 will accept, giving P2 a utility of $\delta^{(3k+2)} \cdot \frac{1}{1+\delta+\delta^2} = \delta^{(3k+1)} \cdot \frac{\delta}{1+\delta+\delta^2} \leq \delta^{(3k+1)} \cdot (1 - x)$. So P2 has no profitable one-shot deviation.

(6) Subgame starting with P1 responding to an offer $(x, 1 - x)$ in period $3k + 3$.

If $x < \frac{\delta+\delta^2}{1+\delta+\delta^2}$, the strategy for P1 prescribes rejection. In the next period, P1 will offer herself $\frac{1+\delta}{1+\delta+\delta^2}$, which P2 will accept, giving P1 a utility of $\delta^{(3k+3)} \cdot \frac{1+\delta}{1+\delta+\delta^2}$. On the other hand, the deviation of accepting x in the current period gives P1 utility $\delta^{(3k+2)} \cdot (x) < \delta^{(3k+2)} \cdot \frac{\delta+\delta^2}{1+\delta+\delta^2} = \delta^{(3k+3)} \cdot \frac{1+\delta}{1+\delta+\delta^2}$. So P1 has no profitable one-shot deviation.

If $x \geq \frac{\delta+\delta^2}{1+\delta+\delta^2}$, the strategy for P1 prescribes acceptance, giving P1 utility of $\delta^{(3k+2)} \cdot (x) \geq \delta^{(3k+3)} \cdot \frac{1+\delta}{1+\delta+\delta^2}$. If P1 rejects instead, then in the next period P1 will offer herself $\frac{1+\delta}{1+\delta+\delta^2}$, which P2 will accept, giving P1 a utility of $\delta^{(3k+3)} \cdot \frac{1+\delta}{1+\delta+\delta^2} \leq \delta^{(3k+2)} \cdot (x)$. So P1 has no profitable one-shot deviation.

Along the equilibrium path, P1 offers $\left(\frac{1+\delta}{1+\delta+\delta^2}, \frac{\delta^2}{1+\delta+\delta^2}\right)$ in $t = 1$ and P2 accepts. This is a better outcome for P1 than in the symmetric bargaining game where P1 gets $\frac{1}{1+\delta}$. P1's SPE payoff improves when she has more bargaining power. \blacklozenge

(1) Intro. to repeated games; (2) Folk theorem, infinitely repeated; (3) Folk theorem, finitely repeated

TF: Kevin He (he02@fas.harvard.edu)

1 Introduction to Repeated Games

1.1 What is a repeated game? Many of the normal-form and extensive-form games studied so far can be viewed as **models of one-time encounters**. After players finish playing Rubinstein-Stahl bargaining or high-bid auction, they part ways and never interact again. In many economic situations, however, a group of players may **play the same game again and again** over a long period of time. For instance, a customer might approach a printing shop every month with a major printing job. While the printing shop has an incentive to shirk and produce low-quality output in a one-shot version of this interaction, in a long-run relationship the shop might never shirk as to avoid losing the customer in the future. In general, repeated games study what outcomes can arise in such repeated interactions.

Formally speaking, repeated games (with perfect monitoring²⁹) form an important class of examples within extensive-form games with finite- or infinite-horizon, depending on the length of repetition.

Definition 70. For a normal-form game³⁰ $\mathcal{G} = \langle \mathcal{N}, (A_k)_{k \in \mathcal{N}}, (u_k)_{k \in \mathcal{N}} \rangle$ and a positive integer T , denote by $\mathcal{G}(T)$ the extensive-form game where \mathcal{G} is played in every period for T periods and players observe the action profiles from all previous periods. \mathcal{G} is called the **stage game** and $\mathcal{G}(T)$ the **(T-times) repeated game**. Terminal vertices of $\mathcal{G}(T)$ are of the form $(a^1, a^2, \dots, a^T) \in \left(\times_{k=1}^T A_k\right)^T =: A^T$ and payoff to player i at such a terminal vertex is $\sum_{t=1}^T u_i(a^t)$. A pure strategy for player i maps each non-terminal history of action profiles to a stage game action, $s_i : \cup_{k=0}^{T-1} A^k \rightarrow A_i$.

Definition 71. For a normal-form game $\mathcal{G} = \langle \mathcal{N}, (A_k)_{k \in \mathcal{N}}, (u_k)_{k \in \mathcal{N}} \rangle$ and a $\delta \in [0, 1)$, denote by $\mathcal{G}^\delta(\infty)$ the extensive-form game where \mathcal{G} is played in every period for infinitely many periods and players act like exponential discounters with discount rate δ . This is called the **(infinitely) repeated game with discount rate δ** . An infinite history of the form $(a^1, a^2, \dots) \in A^\infty$ gives player i the payoff $\sum_{t=1}^{\infty} \delta^{t-1} u_i(a^t)$. A pure strategy for player i maps each finite history of action profiles to a stage game action, $s_i : \cup_{k=0}^{\infty} A^k \rightarrow A_i$.

A strategy for player i in $\mathcal{G}(T)$ or $\mathcal{G}^\delta(\infty)$ must specify a valid action of the stage game \mathcal{G} after **any** non-terminal history $(a^1, \dots, a^k) \in A^k$, including those histories that would never be reached under i 's strategy. For example, even if P1's strategy in repeated prisoner's dilemma is to always play defect, she still needs to specify $s_1((C, C), (C, C))$, that is what she will play in period 3 if both players cooperated in the first two periods.

As defined above, our treatment of repeated games focuses on the the simplest case where payoffs in period t are **independent of actions taken in all previous periods**. This rules out, for instance, investment games where players choose a level of contribution every period and the utility in period t depends on the sum of all accumulated capital up to period t .

²⁹That is to say, actions taken in previous periods are common knowledge. There exists a rich literature on repeated games with coarser monitoring structures – for instance, all players observe an imperfect public signal of each period's action profile, or each player privately observes such a signal – and folk theorems in these generalized settings [6, 7, 8].

³⁰When discussing repeated games, we use A_i (for “actions”) to denote i 's pure strategies in the stage game \mathcal{G} as to avoid confusion with i 's pure strategies in the repeated game $\mathcal{G}(T)$, which we write as S_i .

When discussing repeated games, we are often interested in the “average” stage game payoff under a repeated game strategy profile. The following definitions are just normalizations: they ensure that if i gets payoff c in every period, then she is said to have an average payoff of c .

Definition 72. In $\mathcal{G}(T)$, the **average payoff** to i at a terminal vertex $(a^1, a^2, \dots, a^T) \in A^T$ is $\frac{1}{T} \sum_{t=1}^T u_i(a^t)$. In $\mathcal{G}^\delta(\infty)$, the **(discounted) average payoff** to i at the infinite history $(a^1, a^2, \dots) \in A^\infty$ is $(1 - \delta) \cdot \sum_{t=1}^\infty \delta^{t-1} u_i(a^t)$.

1.2 Some immediate results. We first require some additional definitions.

Definition 73. Given a normal-form game \mathcal{G} , the set of **feasible payoffs** is defined as $\text{co}(\{u(a) : a \in A\}) \subseteq \mathbb{R}^N$, where co is the convex hull operator.

These are the payoffs that can be obtained if players use a **public randomization device** to correlate their actions. Specifically, as every $v \in \text{co}(\{u(a) : a \in A\})$ can be written as a weighted average $v = \sum_{k=1}^r p_k \cdot u(a^k)$ where $p_k \geq 0$, $\sum_{k=1}^r p_k = 1$ and $a^k \in A$ for each k , one can construct a correlated strategy profile where all players observe a public random variable that realizes to k with probability p_k , then player i plays a_i^k upon observing k . The expected payoff profile under this correlated strategy profile is v .

Definition 74. In a normal-form game \mathcal{G} , player i 's **minimax payoff** is defined as

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

Call a payoff profile $v \in \mathbb{R}^N$ **individually rational** if $v_i \geq \underline{v}_i$ for every $i \in \mathcal{N}$. Call v **strictly individually rational** if $v_i > \underline{v}_i$ for every $i \in \mathcal{N}$.

Two technical remarks which you can skip:

(1) The outer minimization in minimax payoff is across the set of **correlated** strategy profiles of $-i$. As demonstrated in the coordination game with an eavesdropper (Example 27), the correlated minimax payoff of a player could be **strictly lower** than her independent minimax payoff (when opponents in $-i$ play independently mixed actions). This distinction is not very important for this course, as we will almost always consider two-player stage games when studying repeated games, so that the set of “correlated” strategy profiles of $-i$ is just the set of mixed strategies of $-i$.

(2) We claimed to have described repeated games with **perfect monitoring** in Definitions 71 and 72, but the monitoring structure as written was less than perfect. Players only observe past actions and **cannot always detect deviations from mixed strategies or correlated strategies**, so in particular they do not know for sure if everyone is faithfully playing a correlated minimax strategy profile against i . (Even when there are only 2 players, the minimax strategy against P1 might be a purely mixed strategy of P2. By observing only past actions, P1 does not know if P2 is really randomizing with the correct probabilities.) To remedy this problem, we can assume that every coalition (including singleton coalitions) observes a correlating signal at the start of every period, which they use to implement correlated strategies and mixed strategies. Furthermore, the realizations of such correlating signals become publicly known at the end of each period, so that even correlated strategies and mixed strategies are “observable”. This remark is again not very important for this course, for the minimax action profile turns out to be pure in most stage games we examine. In addition, Fudenberg and Maskin showed that their 1986 folk theorem continues to hold, albeit with a modified proof, even when players only observe past actions and not the realizations of past correlating devices.

The following two results are now immediate.

Proposition 75. *Suppose $(s_k^*)_{k \in \mathcal{N}}$ is an NE for $\mathcal{G}(T)$ or $\mathcal{G}^\delta(\infty)$. Then the average payoff profile associated with $(s_k^*)_{k \in \mathcal{N}}$ is feasible and individually rational for the stage game \mathcal{G} .*

Proof. Evidently, the payoff profile in every period of the repeated game must be in $\text{co}(\{u(a) : a \in A\})$. In $\mathcal{G}(T)$, the average payoff profile under $(s_k^*)_{k \in \mathcal{N}}$ is the simple average of T such points, while in $\mathcal{G}^\delta(\infty)$ it is a weighted average of countably many such points, so in both cases the average payoff profile must still be in $\text{co}(\{u(a) : a \in A\})$ by the convexity of this set. Suppose now player i 's average payoff is strictly less than \underline{v}_i . Then consider a new repeated game strategy \hat{s}_i for i , where $\hat{s}_i(h)$ best responds to the (possibly correlated) action profile $s_{-i}^*(h)$ after every non-terminal history h . Then playing \hat{s}_i guarantees i at least \underline{v}_i in every period, so that his average payoff will be at least \underline{v}_i . This would contradict the optimality of s_i^* in the NE $(s_k)_{k \in \mathcal{N}}$. \square

Proposition 76. *If \mathcal{G} has a unique NE, then for any finite T the repeated game $\mathcal{G}(T)$ has a unique SPE. In this SPE, players play the unique stage game NE after every non-terminal history.*

Proof. Let $(s_k^*)_{k \in \mathcal{N}}$ be an SPE of $\mathcal{G}(T)$. For any history h_{T-1} of length $T-1$, $(s_k^*(h_{T-1}))$ must be the unique NE of \mathcal{G} . Else, some player must have a strictly profitable deviation in the subgame starting at h_{T-1} . So we deduce $(s_k^*)_{k \in \mathcal{N}}$ plays the unique NE of \mathcal{G} in period T regardless of what happened in previous periods. But this means $(s_k^*(h_{T-2}))$ must also be the unique NE of \mathcal{G} for any history h_{T-2} of length $T-2$. Otherwise, consider the subgame starting at h_{T-2} . If $(s_k^*(h_{T-2}))$ does not form an NE, some player i can improve her payoff in the current period by changing $s_i^*(h_{T-2})$, and furthermore this change does not affect her payoff in future periods, since we have argued the unique NE of \mathcal{G} will be played in period T regardless of what happened earlier in the repeated game. So we have found a strictly profitable deviation for i in the subgame, contradicting the fact that $(s_k^*)_{k \in \mathcal{N}}$ is an SPE. Hence, we have shown $(s_k^*)_{k \in \mathcal{N}}$ plays the unique NE of \mathcal{G} in the last two periods of $\mathcal{G}(T)$, regardless of what happened earlier. Continuing this argument shows the unique NE of \mathcal{G} is played after any non-terminal history. \square

2 Folk Theorem for Infinitely Repeated Games

2.1 The folk theorem for infinitely repeated games under perfect monitoring. It is natural to ask what payoff profiles can arise in $\mathcal{G}^\delta(\infty)$. Write $\mathcal{E}(\mathcal{G}^\delta(\infty))$ for the set of average payoff profiles attainable in SPEs of $\mathcal{G}^\delta(\infty)$. Since every SPE is an NE, in view of Proposition 75, the most we could hope for are results of the following form: “ $\lim_{\delta \rightarrow 1} \mathcal{E}(\mathcal{G}^\delta(\infty))$ equals the set of feasible, individually rational payoffs of \mathcal{G} .” Theorems along this line are usually called “**folk theorems**”, for such results were widely believed and formed part of the economic folklore long before anyone obtained a formal proof. It is important to remember that folk theorems are not merely efficiency results. They are more correctly characterized as “**anything-goes results**”. Not only do they say that there exist SPEs with payoff profiles close to the Pareto frontier, but they also say there exist other SPEs with payoff profiles close to players’ minimax payoffs.

The following is a folk theorem for infinitely repeated games with perfect monitoring.

Theorem 77. *(Fudenberg and Maskin 1986 [9]) Write V^* for the set of feasible, strictly individually rational payoff profiles of \mathcal{G} . Assume \mathcal{G} has full dimensionality. For any $v^* \in V^*$, there corresponds a $\bar{\delta} \in (0, 1)$ so that $v^* \in \mathcal{E}(\mathcal{G}^\delta(\infty))$ for all $\delta \in (\bar{\delta}, 1)$.*

Proof. See lecture. \square

2.2 Rewarding minimaxers. The proof of Theorem 77 is constructive and explicitly defines an SPE with average payoff v^* . To ensure subgame-perfection, the construction must ensure that $-i$ **have an incentive to minimax** i in the event that i deviates. It is possible that the minimax action against i hurts some other player $j \neq i$ so much that j would prefer to be minimized instead of minimizing i . The solution, as we saw in lecture, is to **promise a reward** of $\epsilon > 0$ in **all future periods** to players who successfully carry out their roles as minimaxers³¹. This way, at a history that calls for players to minimax i , deviating from the minimax action loses an infinite stream of ϵ payoffs. As players become more patient, this infinite stream of strictly positive payoffs matters far more than the utility cost from finitely many periods of minimizing i .

Sometimes, no such rewards are necessary. This is the case when minimizing i is not particularly costly for her opponents, as the following example demonstrates.

Example 78. (December 2012 Final Exam) Consider an infinitely repeated game with the following symmetric stage game.

	L	C	R
T	-4,-4	12,-8	3,1
M	-8,12	8,8	5,0
B	1,3	0,5	4,4

Construct a **pure** strategy profile of the repeated game with the following properties: (i) the strategy profile is an SPE of the repeated game for all δ close enough to 1; (ii) the average payoffs are $(8, 8)$; (iii) in every subgame, both players' payoffs are nonnegative in each period.

Solution: We quickly verify that each player's pure minimax payoff (i.e. when minimaxers are restricted to using pure strategies) is 1. P1 minimaxes P2 with T , who best responds with R , leading to the payoff profile $(3, 1)$. Symmetrically, P2 minimaxes P1 with L , who best responds with B , giving us the payoff profile $(1, 3)$. So, $(8, 8)$ is feasible and strictly individually rational, even when restricting attention to pure strategies.

However, we cannot directly recite the Fudenberg-Maskin SPE, for their construction uses a public randomization device in several places – for instance to give the $\epsilon > 0$ reward to minimaxers – but the question asks for a pure strategy profile. Even if we are allowed to use public randomizations, we still face the additional restriction that we cannot let any player get a negative payoff in any period, even off-equilibrium. If we publicly randomize over some action profiles, then we are restricted to those action profiles in the lower right corner of the payoff matrix in all subgames.

Perhaps the easiest solution is to **build a simpler SPE** and forget about giving the $\epsilon > 0$ reward to minimaxers altogether. This is possible because for this particular stage game, the minimaxer gets utility 3 while the minimizee gets utility 1, so it is better to minimax than to get minimized. Consider an SPE given by three phases: in **normal phase**, play (M, C) ; in **minimax P1 phase**, play (B, L) ; in **minimax P2 phase**, play (T, R) . If player i deviates during normal phase, go to minimax P i phase. If player j deviates during minimax P i phase, go to minimax P j phase, where possibly $j = i$. If minimax P i phase completes without deviations, go to normal phase.

³¹The strategy profile used in Fudenberg and Maskin (1986) is often called the “stick and carrot strategy”. If a player deviates during the normal phase, the deviator is hit with a “stick” for finitely many periods. Then, all the other players are given a “carrot” for having carried out this sanction.

We verify this strategy profile is an SPE for δ near enough 1 using one-shot deviation principle. Due to symmetry, it suffices to verify P1 has no profitable one-shot deviation in any subgame. For any subgame in normal phase, deviating gives at most³²

$$12 + \delta \cdot 1 + \frac{\delta^2}{1 - \delta} \cdot 8 \quad (5)$$

while not deviating gives

$$8 + \delta \cdot 8 + \frac{\delta^2}{1 - \delta} \cdot 8 \quad (6)$$

(6) minus (5) gives

$$-4 + \delta \cdot 7$$

which is positive for $\delta \geq \frac{4}{7}$.

For a subgame in the minimax P1 phase, deviating not only hurts P1's current period payoff, but also leads to another period of P1 being minmaxed. So P1 has no profitable one-shot deviation in such subgames for any δ .

For a subgame in the minimax P2 phase, deviating to playing M instead of D gives:

$$5 + \delta \cdot 1 + \frac{\delta^2}{1 - \delta} \cdot 8 \quad (7)$$

while not deviating gives:

$$3 + \delta \cdot 8 + \frac{\delta^2}{1 - \delta} \cdot 8 \quad (8)$$

(8) minus (7) gives

$$-2 + \delta 7$$

which is positive for $\delta \geq \frac{2}{7}$.

Therefore this strategy profile is an SPE whenever $\delta \geq \frac{4}{7}$. ◆

It turns out minimaxer rewards are generally unnecessary when there are only two players³³, as the following theorem shows. In particular, this says we can drop the full-dimensionality assumption from the Fudenberg-Maskin theorem when $N = 2$.

Theorem 79. (*Fudenberg and Maskin 1986 [9]*) Write V^* for the set of feasible, strictly individually rational payoff profiles of \mathcal{G} where $N = 2$. For any $v^* \in V^*$, there corresponds a $\bar{\delta} \in (0, 1)$ so that $v^* \in \mathcal{E}(\mathcal{G}^\delta(\infty))$ for all $\delta \in (\bar{\delta}, 1)$.

³²Recall that for $\delta \in [0, 1)$,

- $1 + \delta + \delta^2 + \dots = \frac{1}{1 - \delta}$
- $1 + \delta + \delta^2 + \dots + \delta^T = \frac{1 - \delta^{T+1}}{1 - \delta}$

³³However, the SPE from the proof of Theorem 79 is not allowed in Example 78, as it involves players getting payoffs $(-4, -4)$ in some periods of some off-equilibrium subgames.

Proof. We may without loss assume each player's minimax payoff is 0.

Consider a strategy profile with two phases. In **normal phase**, players publicly randomize over vertices $\{u(a) : a \in A\}$ to get v^* as an expected payoff profile. In the **mutual minimax phase**, P1 plays the minimax strategy against P2 while P2 also plays the minimax strategy against P1, for M periods. If any player deviates in the normal phase, go to the mutual minimax phase. If any player deviates in the mutual minimax phase, restart the mutual minimax phase. If the mutual minimax phase completes without deviations, go to normal phase.

We show that for suitable choice of M , this strategy profile forms an SPE for all δ near 1. Write

$$\bar{v}_i := \max_{a_i, a_{-i} \in A_{-i}} u_i(a_i, a_{-i})$$

and $\bar{v} := \max\{\bar{v}_1, \bar{v}_2\}$. Choose $M \in \mathbb{N}$ so that $M \cdot v_i^* \geq 2\bar{v}$ for each $i \in \{1, 2\}$. Write \underline{u}_i as the payoff to i when i and $-i$ both play the minimax actions against each other and note $\underline{u}_i \leq 0$.

Consider a subgame in normal phase. If player i makes a one-shot deviation, she gets at most:

$$\bar{v} + \delta \cdot \underline{u}_i + \delta^2 \cdot \underline{u}_i + \dots + \delta^M \cdot \underline{u}_i + \frac{\delta^{M+1}}{1 - \delta} \cdot v_i^* \quad (9)$$

while conforming gives:

$$v_i^* + \delta \cdot v_i^* + \delta^2 \cdot v_i^* + \dots + \delta^M \cdot v_i^* + \frac{\delta^{M+1}}{1 - \delta} \cdot v_i^* \quad (10)$$

(10) minus (9) gives:

$$v_i^* - \bar{v} + (\delta + \dots + \delta^M) \cdot (v_i^* - \underline{u}_i) \geq -\bar{v} + (\delta + \dots + \delta^M) \cdot (v_i^*)$$

where we used $v_i^* > 0$, $\underline{u}_i \leq 0$. But for δ close to 1, $\delta + \dots + \delta^M \geq M/2$, hence implying $-\bar{v} + (\delta + \dots + \delta^M) \cdot (v_i^*) \geq 0$ by choice of M . So for $\delta \geq \bar{\delta}$, there are no profitable one-shot deviations in normal phase.

Consider a subgame in the first period of the mutual minimax phase. If player i deviates, she gets at most:

$$0 + \delta \cdot \underline{u}_i + \delta^2 \cdot \underline{u}_i + \dots + \delta^M \cdot \underline{u}_i + \frac{\delta^{M+1}}{1 - \delta} \cdot v_i^* \quad (11)$$

where $-i$ playing the minimax strategy against i implies her payoff in the period of deviation is bounded by 0. On the other hand, conforming gives:

$$\underline{u}_i + \delta \cdot \underline{u}_i + \delta^2 \cdot \underline{u}_i + \dots + \delta^M \cdot \underline{u}_i + \frac{\delta^{M+1}}{1 - \delta} \cdot v_i^* \quad (12)$$

(12) minus (11) gives:

$$(1 - \delta^M) \cdot \underline{u}_i + \delta^M \cdot v_i^*$$

But $\underline{u}_i \leq 0$ while $v_i^* > 0$, so for δ sufficiently close to 1 this expression is positive. This shows i does not have a profitable one-shot deviation in the first period of the mutual minimax phase. *A fortiori*, she cannot have a profitable one-shot deviation in later periods of the mutual minimax phase either. \square

3 Folk Theorem for Finitely Repeated Games

In view of Proposition 76, the stage game \mathcal{G} must have multiple NEs for $\mathcal{G}(T)$ to admit more than one SPE. The following result is not the most general one, but it shows how one can use the multiplicity of NEs in the stage game to incentivize cooperative behavior for most of the T periods.

Proposition 80. *Suppose \mathcal{G} has a pair of NEs $\bar{a}^{(i)}$ and $\underline{a}^{(i)}$ for each player $i \in \mathcal{N}$ such that $u_i(\bar{a}^{(i)}) > u_i(\underline{a}^{(i)})$. Write $d_i := \max_{\hat{a}_i \in A_i, a_{-i} \in A_{-i}} u_i(\hat{a}_i, a_{-i}) - u_i(a_i, a_{-i})$ for an upperbound on deviation utility to i in the stage game and set $M_i \in \mathbb{N}$ with $M_i \cdot (u_i(\bar{a}^{(i)}) - u_i(\underline{a}^{(i)})) \geq d_i$. For any feasible payoff profile v^* where $v_i^* \geq u_i(\underline{a}^{(i)})$ for each i and $T \geq \sum_{k \in \mathcal{N}} M_k$, there exists an SPE of $\mathcal{G}(T)$ where the average payoff for all except the last $\sum_{k \in \mathcal{N}} M_k$ periods is v^* .*

Unlike an infinitely repeated game, a finitely repeated game “unravels” because some NE must be played in the last period. However, if \mathcal{G} has multiple NEs, then conditioning which NEs get played in the last few periods of $\mathcal{G}(T)$ on players’ behavior in the early periods of the repeated game provides incentives for cooperation.

Proof. Consider the following strategy profile. In the first $T - \sum_{k \in \mathcal{N}} M_k$ periods, if no one has deviated so far, publicly randomize so that expected payoff profile is v^* . If some players have deviated and player i was the first to deviate, then play $\underline{a}^{(i)}$. In the last $\sum_{k \in \mathcal{N}} M_k$ periods, if no one deviated in the first $T - \sum_{k \in \mathcal{N}} M_k$ periods, then play $\bar{a}^{(1)}$ for M_1 periods, followed by $\bar{a}^{(2)}$ for M_2 periods, ..., finally $\bar{a}^{(N)}$ for M_N periods. If someone deviated in the first $T - \sum_{k \in \mathcal{N}} M_k$ periods and i was the first to deviate, then do the same as before except play $\underline{a}^{(i)}$ for M_i periods instead of $\bar{a}^{(i)}$.

We use the one-shot deviation principle to argue this strategy profile forms an SPE. At a subgame starting in the first $T - \sum_{k \in \mathcal{N}} M_k$ periods without prior deviations, suppose i deviates. Compared with conforming to the SPE strategy, i gains at most d_i in the current period, but gets weakly worse payoffs for the remainder of these first $T - \sum_{k \in \mathcal{N}} M_k$ periods as $v_i^* \geq u_i(\underline{a}^{(i)})$. In addition, i loses at least d_i across M_i periods in the last $\sum_{k \in \mathcal{N}} M_k$ periods of the game by choice of M_i . Therefore, i does not have a profitable one-shot deviation.

At a subgame starting in the first $T - \sum_{k \in \mathcal{N}} M_k$ periods with prior deviation, the SPE specifies playing some NE action profile of the stage game. Deviation can only hurt current period payoff with no effect on the payoffs of any future periods. Similarly for subgames starting in the last $\sum_{k \in \mathcal{N}} M_k$ periods. \square

Example 81. (December 2013 Final Exam) Suppose the following game is repeated T times and each player maximizes the sum of her payoffs in these T plays. Show that, for every $\epsilon > 0$, we can choose T big enough so that there exists an SPE of the repeated game in which each player’s average payoff is within ϵ of 2.

	A	B	C
A	2,2	-1,3	0,0
B	3,-1	1,1	0,0
C	0,0	0,0	0,0

Solution: For given $\epsilon > 0$, choose large enough T such that $\frac{2(T-1)+1}{T} > 2 - \epsilon$. Consider the following strategy profile: in period $t < T$, play A if (A,A) has been played in all previous periods, else play

C. In period T , play B if (A,A) has been played in all previous periods, else play C. At a history in period $t < T$ where (A,A) has been played in all previous periods, a one-shot deviation at most gains 1 in the current period but loses 2 in each of periods $t + 1, t + 2, \dots, T - 1$, and finally loses 1 in period T . At a history in period $t < T$ with prior deviation, one-shot deviation hurts current period payoff and does not change future payoffs. At a history in period T , clearly there is no profitable one-shot deviation as this is the last period of the repeated game and the strategy profile prescribes playing a Nash equilibrium of the stage game. \blacklozenge

(1) Refinements of NE in extensive- and normal-form games; (2) Signaling games

TF: Kevin He (he02@fas.harvard.edu)

1 Refinements of NE in Extensive- and Normal-Form Games

1.1 Four refinements. In lecture we studied four refinements of NE for extensive-form games: perfect Bayesian equilibrium (PBE), sequential equilibrium (SE), trembling-hand perfect equilibrium (THPE), and strategically stable equilibrium (SSE). Whereas specifying an NE or SPE just requires writing down a profile of strategies, PBE and SE are defined in terms of not only a strategy profile, but also a **belief system** – that is, a distribution $\pi_j(\cdot|I_j)$ over the vertices in information set I_j for each information set of each player j . They differ in terms of some **consistency conditions** they impose on the belief system.

Definition 82. A **perfect Bayesian equilibrium (PBE)** is a behavioral strategy profile $(p_j)_{j \in \mathcal{N}}$ together with a belief system $(\pi_j(\cdot|I_j))_{j \in \mathcal{N}, I_j \in \mathcal{I}_j}$ so that:

- $p_j(I_j)$ maximizes expected payoffs starting from information set I_j according to belief $\pi_j(\cdot|I_j)$, for each $j \in \mathcal{N}, I_j \in \mathcal{I}_j$
- beliefs are derived from Bayes' rule at all on-path information sets (an information set is called **on-path** if it is reached with strictly positive probability under $(p_j)_{j \in \mathcal{N}}$. Else, it is called **off-path**.)

If an information set I_j is reached with strictly positive probability under $(p_j)_{j \in \mathcal{N}}$, then the conditional probability of having reached each vertex $v \in I_j$ given that I_j is reached is well-defined. This is $\pi_j(v|I_j)$. On the other hand, we cannot use Bayes' rule to compute the conditional probability of reaching various vertices in an off-path information set, as we would be dividing by 0. As such, PBE places no restrictions on these off-equilibrium beliefs.

Definition 83. A **sequential equilibrium (SE)** is a behavioral strategy profile $(p_j)_{j \in \mathcal{N}}$ together with a belief system $(\pi_j(\cdot|I_j))_{j \in \mathcal{N}, I_j \in \mathcal{I}_j}$ so that:

- $p_j(I_j)$ maximizes expected payoffs starting from information set I_j according to belief $\pi_j(\cdot|I_j)$, for each $j \in \mathcal{N}, I_j \in \mathcal{I}_j$.
- there exists a sequence of strictly mixed behavioral strategies $(p_j^{(m)})$ so that $\lim_{m \rightarrow \infty} (p_j^{(m)}) = (p_j)$, and furthermore $\lim_{m \rightarrow \infty} (\pi_j^{(m)}) = (\pi_j)$, where $(\pi_j^{(m)})$ is the unique belief system consistent with $(p_j^{(m)})$.

Though it is not part of the definition, it is easy to show that in an SE, all on-path beliefs are given by Bayes' rule, just as in PBE.

Compared to PBE, SE places some **additional restrictions on off-equilibrium beliefs**. Instead of allowing them to be completely arbitrary, SE insists that these off-equilibrium beliefs must be attainable as the limiting beliefs of a sequence of strictly mixed strategy profiles that converge to $(p_j)_{j \in \mathcal{N}}$ – hence the name **sequential** equilibrium. Given a strictly mixed $(p_j^{(m)})_{j \in \mathcal{N}}$, every information set is

reached with strictly positive probability. Therefore, the belief system $(\pi_j^{(m)})_{j \in \mathcal{N}}$ is well-defined, as there exists exactly one such system consistent with the behavioral strategies $(p_j^{(m)})_{j \in \mathcal{N}}$ under Bayes' rule.

Importantly, there are **no assumptions of rationality** on the sequence of strategies $(p_j^{(m)})$. It is merely a device used to justify how the belief system (π_j) might arise. In particular, there is no requirement that $(p_j^{(m)})$ forms any kind of equilibrium under beliefs $(\pi_j^{(m)})$.

The next two equilibrium concepts, THPE and SSE, are defined in terms of trembles. A **tremble** ϵ in an extensive-form game associates a small, positive probability to each move in each information set, interpreted as the minimum weight that any strategy must assign to the move. The strategy profile (p_j, p_{-j}) is said to be an ϵ -**equilibrium** if at each information set I_j and with respect to belief $\pi_j(\cdot|I_j)$, p_j maximizes j 's expected utility subject to the constraint of minimum weights from the tremble ϵ .

Definition 84. A **trembling-hand perfect equilibrium (THPE)** is a behavioral strategy profile $(p_j)_{j \in \mathcal{N}}$ so that **there exists** a sequence of trembles $\epsilon^{(m)}$ converging to 0 and a sequence of strictly mixed behavioral strategies $(p_j^{(m)})$, such that $(p_j^{(m)})$ is an $\epsilon^{(m)}$ -equilibrium and $\lim_{m \rightarrow \infty} (p_j^{(m)}) = (p_j)$.

Definition 85. A **strategically stable equilibrium (SSE)** is a behavioral strategy profile $(p_j)_{j \in \mathcal{N}}$ so that **for every** sequence of trembles $\epsilon^{(m)}$ converging to 0, there exists a sequence of strictly mixed behavioral strategies $(p_j^{(m)})$ such that $(p_j^{(m)})$ is an $\epsilon^{(m)}$ -equilibrium and $\lim_{m \rightarrow \infty} (p_j^{(m)}) = (p_j)$.

THPE and SSE are also defined for normal-form games, where the tremble ϵ specifies minimum weights for the different actions of various players.

The following table summarizes some of the key comparisons between these four equilibrium concepts.

	belief at on-path info. set	belief at off-path info. set	robustness to trembles
PBE	determined by Bayes' rule	no restriction	none
SE	determined by Bayes' rule	limit of beliefs associated with strictly mixed profiles	none
THPE	N/A	N/A	robust to one sequence of trembles
SSE	N/A	N/A	robust to any sequence of trembles

Finally here are some useful facts:

Fact 86. For extensive-form game Γ , the following inclusions³⁴ hold:

$$NE(\Gamma) \supseteq \begin{matrix} PBE(\Gamma) \\ SPE(\Gamma) \end{matrix} \supseteq SE(\Gamma) \supseteq THPE(\Gamma) \supseteq SSE(\Gamma)$$

There is no inclusion relationship between $PBE(\Gamma)$ and $SPE(\Gamma)$.

³⁴Inclusion in terms of strategy profiles. Technically, NE does not belong to the same universe as PBE, SE, etc., as these later objects require a belief system in addition to a strategy profile.

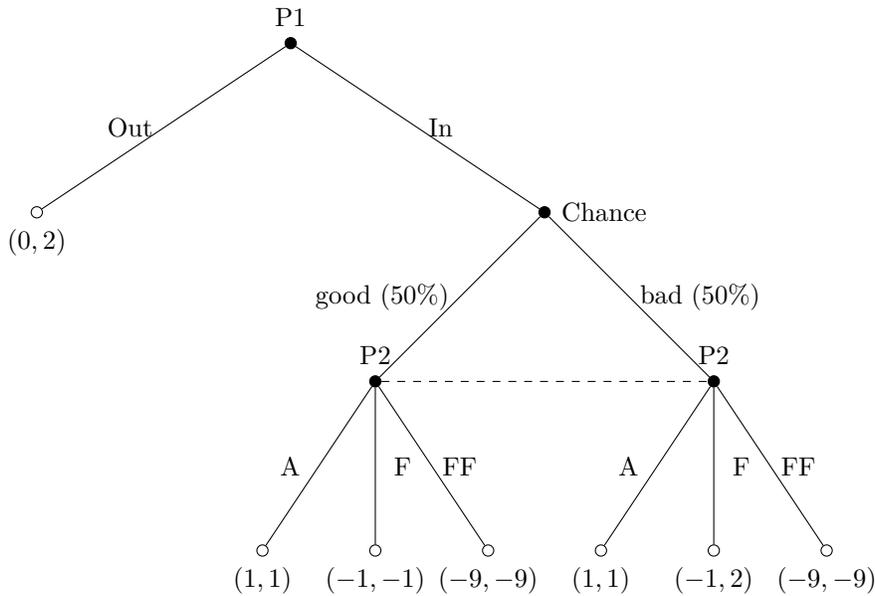


Figure 10: A modified market entry game.

Fact 87. For a finite extensive-form game Γ , $THPE(\Gamma) \neq \emptyset$, though it is possible that $SSE(\Gamma) = \emptyset$.

That is, there is always at least one THPE in a finite extensive-form game (but it may be in behavioral strategies, not pure strategies). The immediate implication is that SE, SPE, PBE, and NE are also non-empty equilibrium concepts.

1.2 Some examples. We illustrate these refinement concepts through two examples. The first example shows an extensive-form game where we have strict inclusions $NE(\Gamma) \supsetneq PBE(\Gamma) \supsetneq SE(\Gamma)$.

Example 88. Consider the following modification to the entry game. The entrant (P1) chooses whether to stay out or enter the market. If she enters, nature then determines whether her product is good or bad, each with 50% probability. Incumbent (P2) observes entry decision, but not whether the product is good or bad. If entrant enters, the incumbent can choose (A) to allow entry, (F) to fight, or (FF) to fight fiercely. This extensive game is depicted in Figure 10. Let's write I for P2's information set and abbreviate strategies in the obvious way (eg. (Out, F) is the strategy profile where P1 plays Out, P2 plays F) rather than writing them out in formal notations. Restrict attention to pure equilibria.

What are the **pure NEs** of this game?

- It is easy to see that (Out, F), (Out, FF) are NEs. P2's action has no effect on his payoff, since P1 never enters. P1 does not have a profitable deviation either, as playing In gets an expected utility of -1 if P2 is playing F, -9 if P2 is playing FF.
- In addition, (In, A) is also an NE. P2 does not have a profitable deviation to F, since doing so yields an expected payoff of $0.5 \cdot (-1) + 0.5 \cdot (2) = 0.75$. P2 does not have a profitable deviation to FF, since doing so yields an expected payoff of -9.

What are the **pure PBEs** of this game?

- These must form a subset of pure NEs.

- There cannot be a PBE with strategy profile (Out, FF). No matter what off-equilibrium belief P2 holds, he will find it strictly better to play A rather than FF.
- However, for any belief $\pi_2(\text{bad}|I) \geq \frac{2}{3}$, $((\text{Out}, \text{F}), \pi_2(\cdot|I))$ forms a PBE³⁵. Recall that P2 is allowed to hold arbitrary beliefs at I since this information set is off-path under strategy profile (Out, F).
- In addition, $((\text{In}, \text{A}), \pi_2(\text{bad}|I) = 0.5)$ is another PBE. In fact, this is the only PBE featuring strategy profile (In, A), since the information set I is on-path for this strategy profile and so $\pi_2(\cdot|I)$ must be derived from Bayes' rule.

What are the **pure SEs** of this game?

- These must form a subset of pure PBEs.
- There is no SE of the form $((\text{Out}, \text{F}), \pi_2(\cdot|I))$. This is because, in every strictly mixed behavioral strategy profile $(p_i^{(m)})$, the information set I is reached with strictly positive probability, meaning Bayes' rule requires $\pi_2^{(m)}(\text{bad}|I) = 0.5$. But SE requires that $\lim_{m \rightarrow \infty} \pi_2^{(m)}(\text{bad}|I) = \pi_2(\text{bad}|I)$, so we see in any SE of this form we must have $\pi_2(\text{bad}|I) = \frac{1}{2}$. SE requires that P2's action at the information set maximizes expected utility subject to belief, yet under the belief $\pi_2(\text{bad}|I) = \frac{1}{2}$, P2 finds it strictly profitable to deviate to A.
- We finally check that $((\text{In}, \text{A}), \pi_2(\text{bad}|I) = 0.5)$ is an SE³⁶. It is straightforward to verify that actions maximize expected payoff at each information set given belief in $((\text{In}, \text{A}), \pi_2(\text{bad}|I) = 0.5)$. Now, consider a sequence of strictly mixed behavioral strategy profiles $(p_i^{(m)})$ where in the m -th profile, P1 plays Out with probability $\frac{1}{m}$, plays In with probability $1 - \frac{1}{m}$, P2 plays each of F and FF with probability $\frac{1}{2m}$ and plays A with probability $1 - \frac{1}{m}$. It is easy to see that $\lim_{m \rightarrow \infty} (p_i^{(m)}) = (\text{In}, \text{A})$. Furthermore, for each such profile, $\pi_2^{(m)}(\text{bad}|I) = 0.5$, so we get $\lim_{m \rightarrow \infty} \pi_2^{(m)}(\cdot|I) = \pi_2(\cdot|I)$.

◆

The second example illustrates THPE and SSE in a normal-form game.

Example 89. (December 2013 Final Exam) In Example 22, we considered the normal-form game

	L	R	Y
T	2, 2	-1, 2	0, 0
B	-1, -1	0, 1	1, -2
X	0, 0	-2, 1	0, 2

³⁵This PBE is not an SPE since the strategy profile does not form an NE when restricted to the subgame starting with the chance move. Recall in lecture we saw an example of an SPE that is not a PBE. This shows neither the set of SPEs nor the set of PBEs nests the other one.

³⁶This does not follow from the non-emptiness of SE as an equilibrium concept, since we have restricted attention to pure equilibria. The game could have a unique mixed SE but no pure SE.

and found that it has two pure NEs, (T, L) and (B, R) , as well as infinitely many mixed NEs, $(T, pL \oplus (1-p)R)$ for $p \in [\frac{1}{4}, 1)$. Now find all the THPEs and SSEs of this game.

Solution: First, we show $(T, pL \oplus (1-p)R)$ is not a THPE for any $p \in [\frac{1}{4}, 1]$ (so we rule out also the pure (T, L)). Suppose there is a sequence of strictly mixed strategy profiles $(p_i^{(m)})$ which are ϵ -equilibria for the sequence of trembles, $\epsilon^{(m)}$, that converge to 0. Since $p_1^{(m)}(B) > 0$ and $p_1^{(m)}(X) > 0$ for each m , R is a **strictly better response** than L against $p_1^{(m)}$ for each m . This means $p_2^{(m)}(L) = \epsilon^{(m)}(L)$ for each m , so that we cannot have $\lim_{m \rightarrow \infty} p_2^{(m)} = pL \oplus (1-p)R$ for $p > 0$. But $\text{THPE}(\mathcal{G}) \subseteq \text{NE}(\mathcal{G})$ and $\text{THPE}(\mathcal{G}) \neq \emptyset$, so (B, R) must be the unique THPE.

Now we check that (B, R) is an SSE³⁷. Suppose P2 plays $\epsilon_1 L \oplus (1 - \epsilon_1 - \epsilon_2)R \oplus \epsilon_2 Y$. Then P1 gets $2\epsilon_1 - (1 - \epsilon_1 - \epsilon_2)$ from T, $-\epsilon_1 + \epsilon_2$ from B, and $-2(1 - \epsilon_1 - \epsilon_2)$ from X. Whenever $\epsilon_1, \epsilon_2 < 0.1$, it is clear that B is the strict best response. Similarly, suppose P1 plays $\epsilon_1 T \oplus (1 - \epsilon_1 - \epsilon_2)B \oplus \epsilon_2 X$. Then P2 gets $2\epsilon_1 - (1 - \epsilon_1 - \epsilon_2)$ from playing L, $2\epsilon_1 + (1 - \epsilon_1 - \epsilon_2) + \epsilon_2$ from playing R, and $-2(1 - \epsilon_1 - \epsilon_2) + 2\epsilon_2$ from playing Y. Whenever $\epsilon_1, \epsilon_2 < 0.1$, it is clear that R is the strict best response. This means that, given **any** sequence of trembles $\epsilon^{(m)}$ converging to 0, eventually the purely mixed strategy profile $(p_i^{(m)})$ where P1 puts **as much weight as possible** on B while P2 puts as much weight as possible on R will be an $\epsilon^{(m)}$ -equilibrium – in fact, this happens as soon as the minimum tremble in $\epsilon^{(m)}$ falls below 0.1. Then, $\lim_{m \rightarrow \infty} (p_i^{(m)}) = (B, R)$ shows (B, R) is an SSE. \blacklozenge

2 Signaling Games

2.1 Strategies, beliefs, and PBEs in signaling games. Signaling games form an important class of examples within extensive-form games with incomplete information. For a schematic representation, see Figure 11.

Nature determines state of the world, $\theta \in \{\theta_1, \theta_2\}$, according to a common prior. P1 is informed of this state. P1 then selects a message from a possibly infinite message set \mathcal{M} and sends it to P2. In the buyer-seller example from class, for instance, the state of the world is the quality of the product while the message is a (price, quantity) pair that the seller offers to buyer.

P2 does not observe the state of the world, but observes the message that P1 sends. This means P2 has **one information set for every message** in A_1 .

A PBE in the signaling game must then have the following components:

- A strategy for P1, that is what message to send in state θ_1 , what message to send in state θ_2
- A strategy for P2, that is how to respond to every $a_1 \in A_1$ that P1 could send (even the off-path messages not sent by P1's strategy)
- A belief system $\{\pi_2(\cdot|a_1)\}_{a_1 \in A_1}$ over states of the world. There is one such $\pi_2(\cdot|a_1) \in \Delta(\{\theta_1, \theta_2\})$ for each message a_1
- P2's belief system is derived from Bayes' rule whenever possible
- P2's strategy after every message a_1 is optimal given $\pi_2(\cdot|a_1)$
- P1's strategy is optimal in every state of the world

³⁷Since SSE is not always non-empty, we cannot immediately conclude that (B, R) must be an SSE.

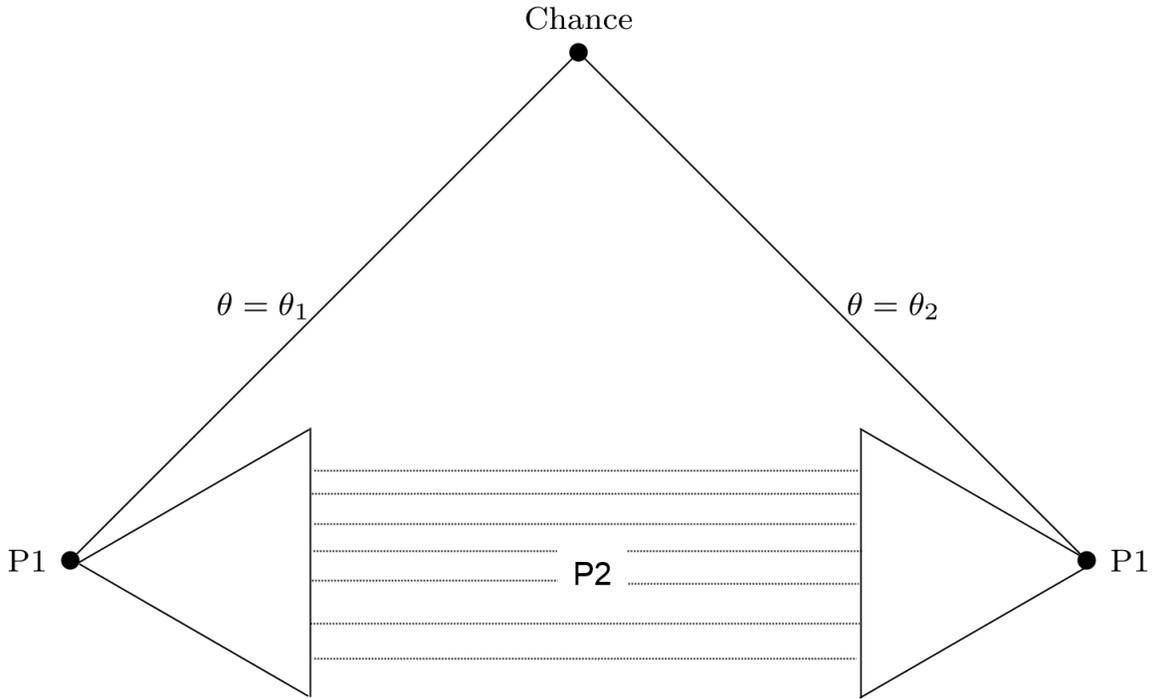


Figure 11: Schematic representation of a signaling game.

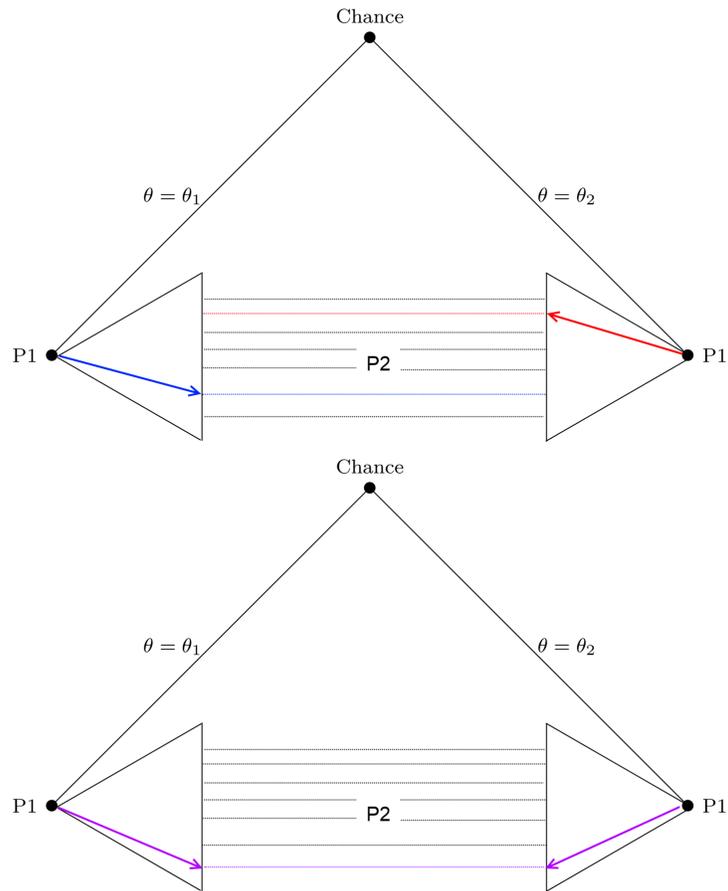


Figure 12: Schematic representations of separating and pooling PBEs.

When there are two states of the world (i.e. two “types” of P1 to borrow terminology from Bayesian games), pure PBEs can be classified into two families, as illustrated in Figure 12. In a **separating PBE**, the two types of P1 send different messages, say $a'_1 \neq a''_1$. By Bayes’ rule, each of these two messages perfectly reveals the state of the world in the PBE. In a **pooling PBE**, the two types of P1 send the same message, say a''' . By Bayes’ rule, P2 should keep his prior about the state of the world after seeing a''' in such a PBE. In a PBE from either family, most of P2’s information sets (i.e. messages he could receive from P1) are off-path. PBE allows P2 to hold arbitrary beliefs on these off-path information sets. In fact, we (the analysts) will often want to pick “**pessimistic**” **off-path beliefs** to help support some strategy profile as a PBE. The following example will illustrate the role of these off-path beliefs in sustaining equilibrium.

2.2 An example. We illustrate separating and pooling PBEs in a civil lawsuit example.

Example 90. Consider a plaintiff (P1) and a defendant (P2) in a civil lawsuit. Plaintiff knows whether she has a strong case (θ_H) or weak case (θ_L), but the defendant does not. Defendant has prior belief that $\pi(\theta_H) = \frac{1}{3}$, $\pi(\theta_L) = \frac{2}{3}$. The plaintiff can ask for a low settlement or a high settlement ($A_1 = \{1, 2\}$). The defendant accepts or refuses, $A_2 = \{y, n\}$. If the defendant accepts a settlement offer of x , the two players **settle out-of-court** with payoffs $(x, -x)$. If defendant refuses, the case goes to trial. If the case is strong ($\theta = \theta_H$), plaintiff wins for sure and the payoffs are $(3, -4)$. If the case is weak ($\theta = \theta_L$), the plaintiff loses for sure and the payoffs are $(-1, 0)$.

Separating equilibrium: Typically, there are **multiple** potential separating equilibria, depending on what action each type of P1 plays. Be sure to check all of them.

- **Separating equilibrium, version 1.** Can there be a PBE where $a_1(\theta_H) = 2$, $a_1(\theta_L) = 1$? If so, in any such PBE we must have $\pi(\theta_H|2) = 1$, $\pi(\theta_L|1) = 1$, $a_2(2) = y$, $a_2(1) = n$. But this means type θ_L gets -1 in PBE and has a profitable unilateral deviation by playing $\hat{a}_1(\theta_L) = 2$ instead. Asking for the high settlement makes P2 think P1 has a strong case, so that P2 will settle and P1 will get 2 instead of -1 . Therefore no such PBE exists.
- **Separating equilibrium, version 2.** Can there be a PBE where $a_1(\theta_H) = 1$, $a_1(\theta_L) = 2$? (It seems very counterintuitive that the plaintiff with a strong case asks for a lower settlement than the plaintiff with a weak case, but this is still a candidate for a separating PBE so we cannot ignore it.) If so, in any such PBE we must have $\pi(\theta_L|2) = 1$, $\pi(\theta_H|1) = 1$, $a_2(2) = n$, $a_2(1) = y$. But this means type θ_H gets 1 in PBE and has a profitable unilateral deviation by playing $\hat{a}_1(\theta_H) = 2$ instead. Asking for the high settlement makes P2 think P1 has a weak case, so that P2 will let the trial go to court. But this is great when P1 has a strong case, giving her a payoff of 3 instead of 1. Therefore no such PBE exists.

Pooling equilibrium: In a pooling equilibrium all types of P1 play the same action. When this “pooled” action a_1^* is observed, P2’s posterior belief is the **same as the prior**, $\pi(\theta|a_1^*) = \pi(\theta)$, since the action carries no additional information about P1’s type. When any other action is observed (i.e. an off-equilibrium action is observed), PBE allows P2’s belief to be **arbitrary**. Every member of A_1 could serve as a pooled action, so we need to **check for all of them systematically**.

- **Pooling on low settlement.** Can there be a PBE where $a_1(\theta_H) = a_1(\theta_L) = 1$? If so, in any such PBE we must have $\pi(\theta_H|1) = \frac{1}{3}$. Under this posterior belief, P2’s expected payoff to $a_2 = n$ is $\frac{1}{3}(-4) + \frac{2}{3}(0) = -\frac{4}{3}$, while playing $a_2 = y$ always yields -1 . Therefore in any such PBE we must have $a_2(1) = y$. But then the θ_H type of P1 has a profitable unilateral deviation of $\hat{a}_1(\theta_H) = 2$, regardless of what $a_2(2)$ is! If $a_2(2) = y$, that is P2 accepts the high settlement,

then type θ_H P1's deviation gives her a payoff of 2 rather than 1. If $a_2(2) = n$, that is P2 refuses the high settlement, then this is even better for the type θ_H P1 as she will get a payoff of 3 when the case goes to court. Therefore no such PBE exists.

- **Pooling on high settlement.** Can there be a PBE where $a_1(\theta_H) = a_1(\theta_L) = 2$? If so, in any such PBE we must have $\pi(\theta_H|2) = \frac{1}{3}$. Under this posterior belief, P2's expected payoff to $a_2 = n$ is $\frac{1}{3}(-4) + \frac{2}{3}(0) = -\frac{4}{3}$, while playing $a_2 = y$ always yields -2 . Therefore in any such PBE we must have $a_2(2) = n$. In order to prevent a deviation by type θ_L , we must ensure $a_2(1) = n$ as well. Else, if P2 accepts the low settlement offer, θ_L would have a profitable deviation where offering the low settlement instead of following the pooling action of high settlement yields her a payoff of 1 instead of -1 . But whether $a_2(1) = n$ is optimal for P2 **depends on the belief**, $\pi(\theta_H|1)$. Fortunately, this is an **off-equilibrium belief** and PBE allows such beliefs to be arbitrary. Suppose $\pi(\theta_H|1) = \lambda \in [0, 1]$. Then P2's expected payoff to playing $a_2(1) = n$ is $\lambda(-4) + (1 - \lambda)(0) = -4\lambda$, while deviating to $a_2(1) = y$ yields -1 for sure. Therefore to ensure P2 does not have a profitable deviation at the information set of seeing a low settlement offer, we need $\lambda \leq \frac{1}{4}$. If P2's off-equilibrium belief is that P1 has a strong case with probability less than $\frac{1}{4}$ upon seeing a low-settlement offer, then it is optimal for P2 to reject such low-settlement offers and θ_L will not have a profitable deviation. At the same time, since P2 rejects all offers, the strong type θ_H of P1 does not have a profitable deviation either, since whichever offer she makes the case will go to court. In summary, there is a family of pooling equilibria where $a_1(\theta_H) = a_1(\theta_L) = 2$, $a_2(1) = a_2(2) = n$, $\pi(\theta_H|2) = \frac{1}{3}$, $\pi(\theta_H|1) = \lambda$ where $\lambda \in [0, \frac{1}{4}]$. Crucially, it is the judicious choice of **off-equilibrium belief** $\pi(\cdot|1)$ that sustains an action of $a_2(1) = n$, which in turn sustains the pooling equilibrium.

All of these pooling PBEs also satisfy the intuitive criterion. Recall that:

Definition 91. Say a PBE $(a_1(\theta), a_2(a_1)), \pi_2(\cdot|a_1)$ satisfies Intuitive Criterion if there do **not** exist $(\hat{a}_1, \hat{a}_2, \hat{\theta})$ such that:

$$(C1) \hat{a}_1 \notin (a_1(\theta_1), a_1(\theta_2))$$

$$(C2) u_1(\hat{a}_1, \hat{a}_2, \hat{\theta}) > u_1(a_1^*(\hat{\theta}), a_2^*(a_1(\hat{\theta})), \hat{\theta})$$

$$(C3) u_1(\hat{a}_1, a_2, \theta) < u_1(a_1^*(\theta), a_2^*(a_1(\theta)), \theta) \text{ for all } a_2 \in A_2, \theta \neq \hat{\theta}$$

$$(C4) \hat{a}_2 \in \arg \max_{a_2 \in A_2} u_2(\hat{a}_1, a_2, \hat{\theta})$$

The only off-equilibrium message \hat{a}_1 in the pooling PBE is the low settlement offer.

Suppose $(\hat{a}_1, \hat{a}_2, \hat{\theta}) = (1, \hat{a}_2, \theta_L)$. Then by (C4), $\hat{a}_2 = n$. But this shows (C2) must not hold, since θ_L gets the same payoff in the PBE as under $(\hat{a}_1, \hat{a}_2, \hat{\theta})$ – the defendant rejects the settlement in both cases.

Suppose $(\hat{a}_1, \hat{a}_2, \hat{\theta}) = (1, \hat{a}_2, \theta_H)$. Then by (C4), $\hat{a}_2 = y$. But this shows (C2) must not hold, since θ_H was actually getting higher payoff in PBE when defendant rejects settlement than under $(\hat{a}_1, \hat{a}_2, \hat{\theta})$, where defendant accepts settlement.

So there are no $(\hat{a}_1, \hat{a}_2, \hat{\theta})$ satisfying (C1) through (C4), meaning the high settlement pooling PBEs satisfy the Intuitive Criterion. \blacklozenge

References

- [1] M. Maschler, E. Solan, and S. Zamir, *Game Theory*, 1st ed. Cambridge: Cambridge University Press, May 2013.
- [2] J. C. Harsanyi, “Games with Incomplete Information Played by Bayesian Players, I-III Part I. The Basic Model,” *Management Science*, vol. 14, no. 3, pp. 159–182, Nov. 1967.
- [3] P. J.-F. Mertens and P. S. Zamir, “Formulation of Bayesian Analysis for Games with Incomplete Information,” *Int J Game Theory*, vol. 14, no. 1, pp. 1–29, Mar. 1985.
- [4] A. Brandenburger and E. Dekel, “Hierarchies of Beliefs and Common Knowledge,” *Journal of Economic Theory*, vol. 59, no. 1, pp. 189–198, Feb. 1993.
- [5] M. J. Osborne and A. Rubinstein, *A Course in Game Theory*. Cambridge: The MIT Press, Jul. 1994.
- [6] D. Fudenberg, D. Levine, and E. Maskin, “The Folk Theorem with Imperfect Public Information,” *Econometrica*, vol. 62, no. 5, pp. 997–1039, Sep. 1994. [Online]. Available: <http://www.jstor.org/stable/2951505>
- [7] M. Kandori and H. Matsushima, “Private Observation, Communication and Collusion,” *Econometrica*, vol. 66, no. 3, pp. 627–652, May 1998. [Online]. Available: <http://www.jstor.org/stable/2998577>
- [8] T. Sugaya, “Folk Theorem in Repeated Games with Private Monitoring,” Social Science Research Network, Rochester, NY, SSRN Scholarly Paper ID 1789775, Feb. 2011. [Online]. Available: <http://papers.ssrn.com/abstract=1789775>
- [9] D. Fudenberg and E. Maskin, “The Folk Theorem in Repeated Games with Discounting or with Incomplete Information,” *Econometrica*, vol. 54, no. 3, pp. 533–554, May 1986. [Online]. Available: <http://www.jstor.org/stable/1911307>