

Part III: Uncertainty and Strategy

7. Uncertainty

8. Game Theory

Chapter 8

Game Theory

Part II

Ming-Ching Luoh

2020.12.11.

Repeated Games

Incomplete Information

Simultaneous Bayesian Games

Signaling Games

Experimental Games

Evolutionary Games and Learning

Extensions: Existence of Nash Equilibrium

Repeated Games

- The simple constituent game (e.g. the Prisoners' Dilemma) that is played repeatedly is called the *stage game*.
- **Repeated play** of the stage game opens up the possibility of **cooperation** in equilibrium.
- Players can adopt *trigger strategies*, whereby they continue to cooperate as long as everyone else does, but revert to playing the Nash equilibrium if anyone deviates from cooperation.
- We will investigate the **conditions** under which trigger strategies work to increase players' payoffs.

Finitely repeated games

- Repeating a stage game for a **known, finite** number of times does **not** increase the possibility for cooperation. This can be shown by **backward induction**.
- Reinhard Selten, winner of the Nobel Prize for his contributions to game theory, showed that for any stage game with a **unique** Nash equilibrium, the unique subgame-perfect equilibrium of the **finitely** repeated game involves playing the Nash equilibrium **every period**.
- If the stage game has **multiple** Nash equilibria, it **may be** possible to achieve cooperation in a finitely repeated game.
- Players can use trigger strategies to maintain cooperation by threatening to play the Nash equilibrium that yields a worse outcome for the player who deviates from cooperation.

Infinitely repeated games

- With infinitely repeated games, there is no definite ending period T from which to **start** backward induction.
- The crucial issue with an infinitely repeated game is not that it goes on forever but that its end is **indeterminate**.
- Players can sustain cooperation in **infinitely** repeated games using trigger strategies. The trigger strategy must be **severe enough** to deter the deviation.
- Suppose both players use the following specific trigger strategy in the Prisoners' Dilemma:

Continue being silent if no one has deviated; fink forever afterward if anyone has deviated to fink in the past.

- To show that this trigger strategy forms a subgame-perfect equilibrium, we need to check that a player **can not gain** from deviating.

		Suspect 2	
		Fink	Silent
Suspect 1	Fink	$U_1 = 1, U_2 = 1$	$U_1 = 3, U_2 = 0$
	Silent	$U_1 = 0, U_2 = 3$	$U_1 = 2, U_2 = 2$

- If both players are silent every period, the present discount value of payoffs, with discount factor δ , over time would be

$$\begin{aligned} V^{eq} &= 2 + 2\delta + 2\delta^2 + 2\delta^3 + \dots \\ &= 2(1 + \delta + \delta^2 + \delta^3 + \dots) = \frac{2}{1 - \delta} \end{aligned}$$

- If a player deviates and then the other finks every period, that player's payoff is

$$\begin{aligned} V^{dev} &= 3 + 1 \cdot \delta + 1 \cdot \delta^2 + \dots \\ &= 3 + \delta(1 + \delta + \delta^2 + \dots) = 3 + \frac{\delta}{1 - \delta} \end{aligned}$$

- The trigger strategies form a subgame-perfect equilibrium if $V^{eq} \geq V^{dev}$, implying that

$$\frac{2}{1-\delta} \geq 3 + \frac{\delta}{1-\delta}, \text{ or } \delta \geq \frac{1}{2}$$

Continued cooperative play is desirable provided players do not **discount** future gains from cooperation **too highly**.

- If $\delta < 1/2$, no cooperation is possible in the infinitely repeated Prisoner's Dilemma; the only subgame-perfect equilibrium involves **finking every period**.
- The trigger strategy in which players revert to stage-game Nash equilibrium **forever** is the **harshest** punishment possible is called the **grim strategy**.
- Less harsh punishment include the so-called **tit-for-tat strategy**, which involves only **one round** of punishment for cheating.

- The grim strategy elicits cooperation for the **largest range** of cases (the lowest value of δ) of any strategy.
- As δ approaches 1, grim-strategy punishments become **infinitely harsh** because they involve an unending stream of undiscounted losses.
- Infinite punishments can be used to sustain a wide range of possible outcomes. This is the logic behind the *folk theorem for infinitely repeated games*.
- Take any stage-game payoff for a player between Nash equilibrium one and the **highest one**, with payoff V , that appears anywhere in the payoff matrix, the **folk theorem** says that the player can earn V in some subgame-perfect equilibrium for δ close enough to **1**.

Incomplete Information

- In the games studies thus far, players knew everything there was to know about the setup of the game, including each others' strategy sets and payoffs.
- Matters become more complicated, and potentially more interesting, if some players have information about the game that others **do not**.
- Games of **incomplete** information can quickly become complicated. Players that lack full information will use what they do know to **make inferences** about what they do not.
- **Probability theory** provides a formula, called *Bayes' rule*, for making inferences about **hidden information**.
- The relevance of Bayes' rule in games of incomplete information has led them to be called *Bayesian games*.

- To limit the complexity of the analysis, we will focus on **two-player** games in which one of the players (player 1) has **private information** and the other (player 2) does not.
- The next section begins with the simple case in which the players move **simultaneously**.
- The subsequent section then analyzes games in which the informed player 1 **move first**.
- Such games, called **signaling games**, are more complicated than simultaneous games because player 1's action may signal something about his or her private information to **uninformed** player 2.
- We will introduce **Bayes' rule** at that point to help analyze player 2's inference about player 1's hidden information based on **observations** of player 1's action.

Simultaneous Bayesian Games

- A two-player, simultaneous-move game in which player 1 has **private information**, but player 2 does not.
- We use “**he**” for player 1 and “**she**” for player 2 to facilitate exposition.

Player types and beliefs

- Player 1 can be of a number of possible *types*, denoted t . Player 1 knows his own **type**.
- Player 2 is uncertain about t and must decide on her strategy based on **beliefs** about t .
- The game begins at a **chance node**, at which a value of t_k is drawn for player 1 from a set of possible types,

$$T = \{t_1, \dots, t_k, \dots, t_K\}.$$

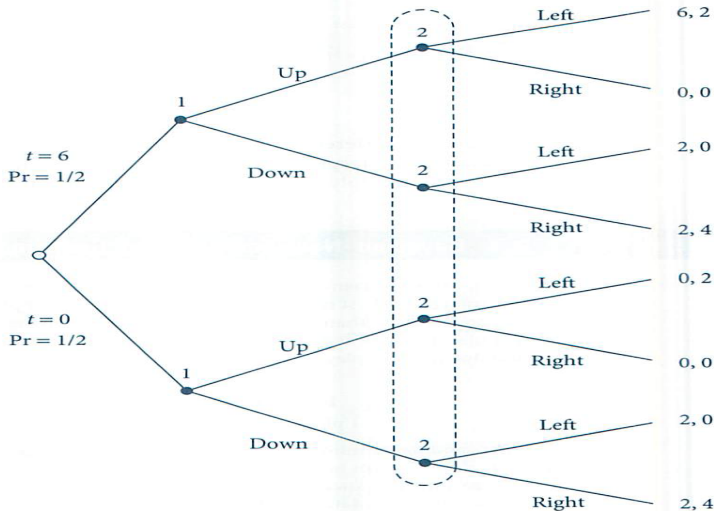
- Let $Pr(t_k)$ be the probability of drawing type t_k .
- Player 1 sees which type is drawn. Player 2 does not see and only **knows** the probabilities, $Pr(t_k)$.
- Since player 1 observes t **before moving**, his strategy can be conditioned on t . Let $s_1(t)$ be player 1's strategy contingent on his type.
- Player 2's strategy is an **unconditional one**, s_2 .
- We write player 1's payoff as $U_1(s_1(t), s_2, t)$ and player 2's as $U_2(s_2, s_1(t), t)$.
- Note that Player 1's type may affect player 2's payoff in two ways: directly and indirectly through player 1's strategy.

Figure 8.13 Simple Game of Incomplete Information

		Player 2	
		Left	Right
Player 1	Up	$t, 2$	$0, 0$
	Down	$2, 0$	$2, 4$

- All payoffs are known to both players except for t in the upper left.
- Player 2 only knows the distribution: an equal chance ($1/2$, $1/2$) that $t = 0$ or $t = 6$.
- Player 1 knows the realized value of t , equivalent to knowing his **type**.

Figure 8.14 Extensive Form for Simple Game of Incomplete Information



Bayesian-Nash equilibrium

- Equilibrium requires that player 1's strategy be the best response for **each and every one** of his types.
- Equilibrium also requires player 2's strategy maximize an **expected payoff**, where the expectation is taken with respect to her **beliefs** about player 1's type.
- The calculations involved in computing the best response to the pure strategies of different types of rivals in a game of **incomplete information** are similar to the calculations involved in computing the best response to a rival's mixed strategy in a game of **complete information**.

- **Bayesian-Nash equilibrium.** In a two-player simultaneous move game in which player 1 has private information, a Bayesian-Nash equilibrium is a strategy profile $(s_1^*(t), s_2^*)$ such that $s_1^*(t)$ is a best response to s_2^* for each type $t \in T$ of player 1,

$$U_1(s_1^*(t), s_2^*, t) \geq U_1(s_1', s_2^*, t) \text{ for all } s_1' \in S_1,$$

and such that s_2^* is a best response to $s_1^*(t)$ given player 2's beliefs $Pr(t_k)$ about player 1's types:

$$\begin{aligned} & \sum_{t_k \in T} Pr(t_k) U_2(s_2^*, s_1^*(t_k), t_k) \\ \geq & \sum_{t_k \in T} Pr(t_k) U_2(s_2', s_1^*(t_k), t_k) \text{ for all } s_2' \in S_2 \end{aligned}$$

Example 8.5 Bayesian-Nash Equilibrium of Game in Figure 8.14

First solve for the informed player's (player 1's) **best response** for each of his types. There are two possible candidates for an equilibrium in **pure strategies**.

- 1 plays (Up|t=6, **Down**|t=0) and 2 plays Left

This is not an equilibrium because player 2 would gain by deviating to right, earning an expected payoff of **2**, instead of earning an expected payoff of **1**.

- 1 plays (Down|t=6, **Down**|t=0) and 2 plays Right

This is a **Bayesian-Nash equilibrium** because both would **not** gain by deviating to the other strategy.

Example 8.6 Tragedy of the Commons as a Bayesian Game

- For an example with **continuous actions**, suppose that Herder 1 has **private information** regarding his value of grazing per sheep,

$$v_1(q_1, q_2, t) = t - (q_1 + q_2)$$

where herder 1's type is $t = 130$ ("high" type) with probability $2/3$ and $t = 100$ ("low" type) with probability $1/3$.

- Herder 2's value remains the same as

$$v_2(q_1, q_2) = 120 - q_1 - q_2$$

- Herder 1's value-maximization problem is

$$\max_{q_1} \{q_1 v_1(q_1, q_2, t)\} = \max_{q_1} \{q_1 (t - q_1 - q_2)\}.$$

The first-order condition for a maximum is

$$t - 2q_1 - q_2 = 0$$

- Therefore,

$$q_{1H} = 65 - \frac{q_2}{2}, q_{1L} = 50 - \frac{q_2}{2}$$

where q_{1H} and q_{1L} are the “high type” and “low type” of herder 1 respectively.

- For herder 2's best response. Herder 2's expected payoff is

$$\frac{2}{3}q_2(120 - q_{1H} - q_2) + \frac{1}{3}q_2(120 - q_{1L} - q_2) = q_2(120 - \bar{q}_1 - q_2),$$

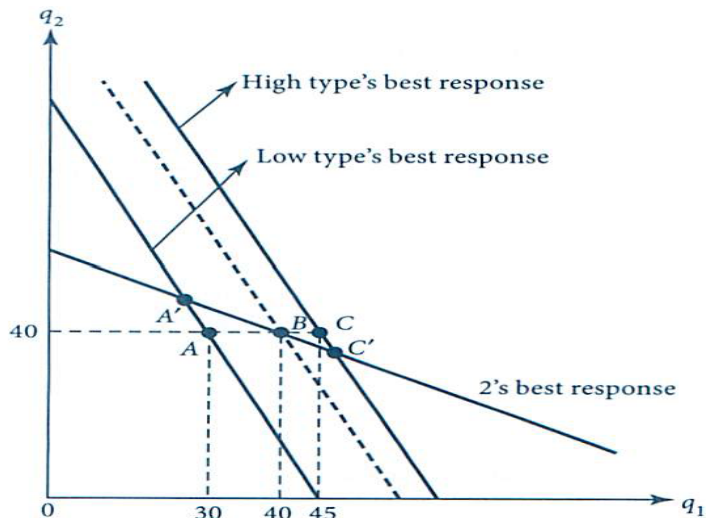
where $\bar{q}_1 = \frac{2}{3}q_{1H} + \frac{1}{3}q_{1L}$.

- The first-order condition for maximizing herder 2's expected payoff is

$$q_2 = 60 - \frac{\bar{q}_1}{2} = 60 - \frac{60 - \frac{q_2}{2}}{2} = 30 + \frac{q_2}{4},$$

implying that $q_2^* = 40$ and $q_{1H} = 45, q_{1L} = 30$.

Figure 8.15 Equilibrium of the Bayesian Tragedy of the Commons




Signaling Games

- In this section we move from **simultaneous-move games** of private information to sequential game in which the informed player, player 1, takes an **action** that is **observable** to player 2 before player 2 moves.
- Player 1's action can serve as a **signal** to player 2 that can be used to **update** her beliefs about player 1's type.
- For example, a prestigious college degree may signal that a job applicant is highly skilled.
- Signaling games are more complicated because we need to model how player 2 **processes** the information in player 1's **signal** and then updates her beliefs about player 1's type.

Job-market signaling

- A version of **Michael Spence's** model of job-market signalling, for which he won the Nobel Prize in economics.
- Player 1 is a worker who can be of two types, high skilled ($t = H$) and low skilled ($t = L$).
- Player 2 is a firm considering hiring player 1.
- Low-skilled worker generates **no revenue** for the firm while high-skilled worker generates revenue of π .
- If a worker is hired, the firm must pay the worker w , with $\pi > w > 0$.
- The firm cannot observe the worker's skill, only his education level.
- Let c_L be the cost of obtaining an education for the low type. Let c_H be the cost of obtaining an education for the high type, and $c_H < c_L$.

Figure 8.16 Job-Market Signaling



- Player 1 observes his type at the start; player 2 observes only player 1's **action** (educational **signal**) before moving.
- Let $Pr(H)$ and $Pr(L)$ be player 2's **beliefs** before observing player 1's educational signal that player 1 is high- or low-skilled. These are called player 2's **prior beliefs**.
- Observing player 1's action will lead player 2 to **revise** his or her beliefs to form what are called **posterior beliefs**.
- Suppose player 2 sees player 1 choose E , then player 2's **expected** pay-off from playing J is

$$\begin{aligned} & Pr(H|E)(\pi - w) + Pr(L|E)(-w) \\ = & Pr(H|E)\pi - (Pr(H|E) + Pr(L|E)) \cdot w = Pr(H|E)\pi - w \end{aligned}$$

Player 2's payoff from playing NJ is 0 . Therefore, player 2's best response is J if and only if $Pr(H|E) > \frac{w}{\pi}$.

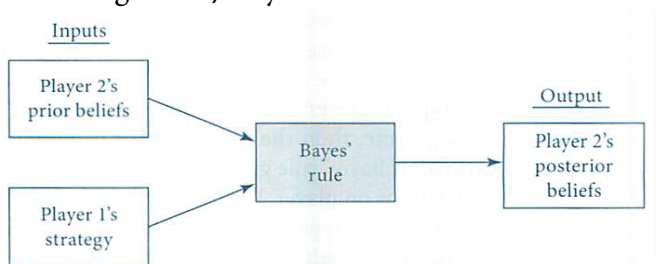
Bayes' rule

- Bayes' rule gives the following formula for computing player 2's **posterior belief** $Pr(H|E)$ and $Pr(H|NE)$;

$$Pr(H|E) = \frac{Pr(E|H)Pr(H)}{Pr(E|H)Pr(H) + Pr(E|L)Pr(L)}$$

$$Pr(H|NE) = \frac{Pr(NE|H)Pr(H)}{Pr(NE|H)Pr(H) + Pr(NE|L)Pr(L)}$$

Figure 8.17 Bayes' Rule as a Black Box



- When player 1 plays a **pure strategy**, Bayes' rule often gives a simple result. Suppose that $Pr(E|H) = 1$ and $Pr(E|L) = 0$, that is, player 1 obtains an education **if and only if** he or she is high-skilled.

$$Pr(H|E) = \frac{1 \cdot Pr(H)}{1 \cdot Pr(H) + 0 \cdot Pr(L)} = 1$$

- Suppose that $Pr(E|H) = Pr(E|L) = 1$, that is, player 1 obtains an education **regardless** of his or her type.

$$Pr(H|E) = \frac{1 \cdot Pr(H)}{1 \cdot Pr(H) + 1 \cdot Pr(L)} = Pr(H)$$

- More generally, when player 1 plays the mixed strategy $Pr(E|H) = p$ and $Pr(E|L) = q$, then

$$Pr(H|E) = \frac{p \cdot Pr(H)}{p \cdot Pr(H) + q \cdot Pr(L)}$$

Perfect Bayesian equilibrium

A **perfect Bayesian equilibrium** consists of a **strategy profile** and a **set of beliefs** such that

- at each information set, the **strategy** of a player moving there maximizes his or her expected payoff, where the expectation is taken with respect to his or her **beliefs**, and
- at each information set, the **beliefs** of the player moving there are formed using **Bayes' rule** (based on prior beliefs and other players' strategies).

- Signaling games may have multiple equilibria, which can be divided into three classes— **separating**, **pooling** and **hybrid**.
- In a *separating equilibrium*, each type of player 1 chooses a **different** action. Player 2 learns player 1's type with certainty after observing player 1's action.
- In a *pooling equilibrium*, different types of player 1 choose the **same action**. Observing player 1's action provides player 2 with **no additional** information about player 1's type.
- In a *hybrid equilibrium*, one type of player 1 plays a strictly mixed strategy. Player 2 learns a little about player 1's type but **doesn't** learn it with **certainty**. Player 2 may respond to the uncertainty by **also playing a mixed strategy**.

Example 8.7 **Separating** Equilibrium in Job-Market Signaling Game

- A good **guess** for a **separating** equilibrium is that the **high-skilled** worker signals his or her type by **getting** an education and the **low-skilled** worker does **not**.
- Given these strategies, player 2's **beliefs** must be $Pr(H|E) = Pr(L|NE) = 1$ and $Pr(H|NE) = Pr(L|E) = 0$.
- Player 2's **best response** if he observes that player 1 obtains an education is to offer a job (J), given the payoff of $\pi - w > 0$, and not offer a job (NJ) if player 1 does not obtain an education because $0 > -w$.
- Lastly, we need to **check** that player 1 would **not want to deviate** from the separating strategy ($E|H, NE|L$) given that player 2 plays ($J|E, NJ|NE$).

- Type H player 1 would not deviate if $w - c_H > 0$, where $w - c_H$ and 0 are payoffs of obtaining and not obtaining an education.
- Type L player 1 would not deviate if $0 > w - c_L$, where 0 and $w - c_L$ are payoffs of obtaining and not obtaining an education.
- There is separating equilibrium in which the worker obtains an education if and only if he or she is high-skilled and in which the firm offers a job only to applicants with an education if and only if $c_H < w < c_L$.

- Another possible separating equilibrium if for player 1 to obtain an education if and only if he or she is **low-skilled**.
- Player 2's best response would be to offer a job if and only if player 1 did **not** obtain an education.
- Under this circumstance, Type L player 1 would earn $0 - c_L$ from playing E because he is not offered an job, and $w - 0$ from playing NE , so he would **deviate** to NE since $-c_L < w$.
- Therefore, this will be ruled out as a perfect Bayesian equilibrium.

Example 8.8 **Pooling** Equilibrium in Job-Market Signalling Game

- A possible pooling equilibrium in which **both** types of player 1 choose E .
- For player 1 not to deviate from choosing E , player 2's strategy **must** be to offer a job if and only if the worker is educated— that is, $(J|E, NJ|NE)$. (**Why?**)
- Player 2's expected payoff from choosing J when player 1 obtains an education is

$$\begin{aligned} & Pr(H|E)(\pi - w) + Pr(L|E)(-w) \\ = & Pr(H)(\pi - w) + Pr(L)(-w) = Pr(H)\pi - w \end{aligned}$$

For J to be a best response for player 2, we need $Pr(H) > \frac{w}{\pi}$.

- For NJ to be a best response to NE for player 2, we need

$$0 > Pr(H|NE)(\pi - w) + Pr(L|NE)(-w) = Pr(H|NE)\pi - w$$

or

$$Pr(H|NE) \leq \frac{w}{\pi}.$$

- In sum, we need $Pr(H|NE) \leq w/\pi \leq Pr(H)$ to have a pooling equilibrium in which both types of player 1 obtain an education. That is, $Pr(H)$ must be sufficiently high and $Pr(H|NE)$ must be sufficiently low.
- In this equilibrium, type L pools with type H to prevent player 2 from learning anything about the worker's skill from the education signal.
- The other possibility for a pooling equilibrium is for both types of player 1 to choose NE . (See **Problem 8.10**)

Example 8.9 Hybrid Equilibrium in Job-Market Signalling Game

- **One possible** hybrid equilibrium is for type H always to obtain an education and type L to randomize between playing E and NE with probabilities e and $1 - e$.
- Player 2's strategy is to offer a job to an educated applicant with probability j and not to offer a job to an uneducated applicant.
- We need to solve for the equilibrium values of e^* and j^* and the posterior beliefs $Pr(H|E)$ and $Pr(H|NE)$ that are **consistent** with perfect Bayesian equilibrium.
- The posterior beliefs are

$$Pr(H|E) = \frac{Pr(H)}{Pr(H) + ePr(L)} = \frac{Pr(H)}{Pr(H) + e[1 - Pr(H)]}$$

$$Pr(H|NE) = 0$$

- For type L of player 1 to be willing to play a strictly **mixed** strategy, he or she must get the **same expected payoff** from playing E and NE . That is

$$jw - c_L = 0, \text{ or } j^* = \frac{c_L}{w}$$

- Player 2 will play a strictly mixed strategy (conditional on **observing E**) only if he or she gets the same expected payoff from playing J and **playing NJ** . That is

$$Pr(H|E)(\pi - w) + Pr(L|E)(-w) = Pr(H|E)\pi - w = 0,$$

or

$$\begin{aligned} Pr(H|E) &= \frac{w}{\pi} = \frac{Pr(H)}{Pr(H) + e[1 - Pr(H)]} \\ e^* &= \frac{(\pi - w)Pr(H)}{w[1 - Pr(H)]} \end{aligned}$$

Experimental Games

- Experimental economics is a recent branch of research that explores how well economic theory **matches** the behavior or experimental subjects in a **laboratory** setting.
- The importance of experimental economics was highlighted in **2002** when Vernon Smith received the Nobel Prize in economics for his pioneering work in the field.
- An important area in this field is the use of experimental methods to **test** game theory.

Experiments with the Prisoners' Dilemma

- In one experiment, subjects played Prisoners' Dilemma game **20** times with each other being matched with a **different**, anonymous opponent to avoid repeated-game effects.
- Players **converged** to the Nash equilibrium as subjects gained **experience** with the game.
- Players played the cooperative action **43%** of the time in the first 5 rounds, falling to only **20%** of the time in the **last five** rounds.
- Although 80% of the decisions were consistent with Nash equilibrium play by the end of the experiment, 20% of them still were **anomalous**.

Experiments with the Ultimatum Game

- In the subgame-perfect equilibrium of the Ultimatum Game, the Proposer offers a **minimal** share of the pot, and this is accepted by the Responder.
- In experiments, the division tends to be **much more** than in the subgame-perfect equilibrium.
- The most common offer is a 50-50 split. Responders tend to receive offers giving them **less than 30%** of the pot.
- Players may care about other factors such as **fairness** and thus obtain a benefit from a more equal division of the pot.

Experiments with the Dictator Game

- In the Dictator Game, the Proposer chooses a split of the pot, and this split is implemented without input from the Responder.
- Proposer tend to offer a **less-even** split than in the Ultimatum Game but still offer the Responder **some** of the pot, suggesting that Proposers have some residual concern for **fairness**.
- In one experiment designed so that the experimenter would **never learn** which Proposer made the offers, Proposers almost never gave an equal split to Responders and indeed took **the whole pot** for themselves **2/3** of the time.

Evolutionary Games and Learning

Evolutionary model

- In the evolutionary model, players do not make rational choices. They play the way they are **genetically** programmed.
- The more successful a player's strategy in the population, the more fit is the player and the more likely the player **survive** to pass his or her genes on the future generations and thus the more likely the strategy spreads in the population.

Learning model

- In a learning model, players are matched at random with others from a large population. Players usually are assumed to have a degree of rationality.
- Players are not fully rational in that they do not distort their strategies to affect others' learning and thus future play.


Extensions: Existence of Nash Equilibrium

- This section will **sketch** John Nash's original proof (1950) that **all finite games have at least one Nash Equilibrium** (in mixed if not in pure strategies).
- Nash's proof is similar to the proof of the existence of a general competitive equilibrium in **Chapter 13**. Both reply on a **fixed point theorem**.
- The proof of the existence of Nash equilibrium requires a slightly more powerful theorem. Instead of Brouwer's fixed point theorem, which applies to **functions**. Nash's proof replies on **Kakutani's fixed point theorem**, which applies to **correspondence**.

E8.1 Correspondences versus functions

- A **function** maps each point in a first set to a **single point** in a second set.
- A **correspondence** maps a single point in the first set to possibly **many points** in the second set. The best response, $BR_i(s_{-i})$, is an example.


Figure E8.1 Comparison of Functions and Correspondences




E8.2 Kakutani's fixed point theorem

- Any **convex**, **upper-semicontinuous** correspondence $[f(x)]$ from a closed, bounded, convex set into itself has at least one fixed point (x^*) such that $x^* \in f(x^*)$.

Figure E8.2 Kakutani's Conditions on Correspondence



(a) Correspondence that is not convex



(b) Correspondence that is not upper semicontinuous

E8.3 Nash's proof


- We use $R(s)$ to denote the **correspondence** that underlies Nash's existence proof.
- This correspondence takes any profile of players' strategies $s = (s_1, s_2, \dots, s_n)$ and maps it into another mixed strategy profile, the profile of best responses:

$$R(s) = (BR_1(s_{-1}), BR_2(s_{-2}), \dots, BR_n(s_{-n}))$$


- A fixed point of the correspondence is a strategy for which $s^* \in R(s^*)$; this is a Nash equilibrium because each player's strategy is a best response to others' strategies.
- The proof checks that all the conditions involved in Kakutani's fixed point theorem are **satisfied** by the best-response correspondence $R(s)$.

1. Show that the set of mixed-strategy profiles is *closed*, *bounded*, and *convex*.
 - A strategy profile is just a list of individual strategies, the set of strategy profiles will be closed, bounded, and convex if each player's strategy set S_i has these properties *individually*.
 - Figure E8.3 shows for the case of two and three actions, the set of mixed strategies over actions has a simple shape.
 - The set is *closed* (contains its boundary), *bounded* (does not go off to infinity in any direction), and *convex* (the segment between any two point in the set is also in the set).

Figure E8.3 Set of Mixed Strategies for an Individual



(a) Two actions



(b) Three actions

2. *Check that the best-response correspondence $R(s)$ is convex.*
- Individual best responses **cannot** look like Figure E8.2a because if any two mixed strategies such as A and B are best responses to others' strategies, then mixed strategies between them must also be best responses.
 - For example, in the Battle of Sexes, if $(1/3, 2/3)$ and $(2/3, 1/3)$ are best responses for husband against his wife playing $(2/3, 1/3)$, then mixed strategies between the two such as $(1/2, 1/2)$ must also be best responses for him. In fact, all possible mixed strategies for the husband are best response to the wife's playing $(2/3, 1/3)$.

3. Check that $R(s)$ is upper semicontinuous

- Individual best responses cannot look like in Figure E8.2b, they can not have holes like point D punched out of them because payoff functions $U_i(s_i, s_{-i})$ are continuous.
- Recall that payoffs, when written as functions of mixed strategies, are actually expected values with probabilities given by the strategies s_i and s_{-i} .
- Expected values are linear functions of the underlying probabilities. Linear functions are, of course, **continuous**.

E8.4 Games with continuous actions

- Nash's existence theorem applies to **finite** games. Nash's theorem does **not** apply to game that feature continuous actions, such as the Tragedy of the Commons in Example 8.5.
- Is a Nash equilibrium guaranteed to exist for these games? Glicksberg(1952) proved that the answer is "yes" as long as payoff functions are **continuous**.

References

- Glickberg, I. L. "A Further Generalization of the Kakutani Fixed Point Theorem with Application to Nash Equilibrium Points." *Proceedings of the National Academy of Sciences* 38 (1952): 170-84.
- Nash, John. "Equilibrium Points in n -Person Games." *Proceedings of the National Academy of Sciences* 36 (1950): 48-49.