

Evolution & Learning in Games


Econ 243B

Jean-Paul Carvalho


Lecture 2.

The Interactive Knowledge Problem

Market Interaction



Social Interaction



Strategic-Form Representation

A strategic-form representation has three ingredients:

- ▶ **Players,**
- ▶ **Strategies,**
- ▶ **Payoffs.**
 - ▶ *Note:* payoffs depend not only on your own strategy, but the strategies of other players.

More formally:

Strategic-Form Games

Definition. A strategic-form (or normal-form) game consists of:

1. *Players.* A set of agents who play the game N , with typical member $i \in N$.
2. *Strategies.* A strategy is a complete plan of action specifying what a player will do at every point at which she may be called upon to play. For each $i \in N$ there is a nonempty set of strategies S_i with typical element $s_i \in S_i$.
3. *Payoffs.* A payoff function $u_i : S \mapsto \mathcal{R}$ assigned to each player i , where $S = \times_{i \in N} S_i$.

Anything with these three features can be written as a strategic-form game:

$$\mathcal{G} = \langle N, \{S_i\}_{i \in N}, \{u_i\}_{i \in N} \rangle.$$

Strategic-Form Games

Aggregate Behavior:

- ▶ A collection of strategies $s \in S = \times_{i \in N} S_i$ is called a **strategy profile**.

Preferences over Strategy Profiles & Utility:

- ▶ There is a preference relation \succeq_i over strategy profiles for each player i ; in fact, we will often require vNM utility functions u_i .

We focus heavily on games with two players and a finite number of strategies:

- ▶ These games can be represented by a **payoff matrix**.

The Prisoners' Dilemma

“Two suspects are arrested for a crime, and interviewed separately. If they both keep quiet (they *cooperate* with each other) they go to prison for a year. If one suspect supplies incriminating evidence (*defects*) then that one is freed, and the other one is imprisoned for nine years. If *both* defect then they are imprisoned for six years. Their preferences are solely contingent on any jail term they individually serve.”

Players. The players are the two suspects $N = \{1, 2\}$.

Strategies. The strategy set for player 1 is $S_1 = \{C, D\}$, and for player 2 is $S_2 = \{C, D\}$.

Payoffs. Represent the payoffs in the strategic-form *payoff matrix*:

	C	D
C	-1, -1	0, -9
D	0, -9	-6, -6

Dominant-Strategy Equilibrium

There are no real strategic issues in the one-shot Prisoners' Dilemma.

As long as players are rational (they choose the best available action given their preferences), each will play D :

		C		D								
C		<table border="1"><tr><td></td><td>-1</td></tr><tr><td>-1</td><td></td></tr></table>		-1	-1		\Rightarrow	<table border="1"><tr><td></td><td>0</td></tr><tr><td>-9</td><td></td></tr></table>		0	-9	
	-1											
-1												
	0											
-9												
		\Downarrow		\Downarrow								
D		<table border="1"><tr><td></td><td>-9</td></tr><tr><td>0</td><td></td></tr></table>		-9	0		\Rightarrow	<table border="1"><tr><td></td><td>-6</td></tr><tr><td>-6</td><td></td></tr></table>		-6	-6	
	-9											
0												
	-6											
-6												

$\{D, D\}$ is a *dominant-strategy equilibrium*.

It is also **inefficient**: both players are better off if (C, C) is played.

Nash Equilibrium

Definition. A *Nash equilibrium* is a strategy profile $s^* \in S$ such that for each $i \in N$,

$$u_i(s_i^*, s_{-i}^*) \geq u_i(s_i, s_{-i}^*) \quad \forall s_i \in S_i.$$

At s^* , no i regrets playing s_i^* . Given all the other players' actions, i couldn't have done better.

Hence a Nash equilibrium is a strategy profile from which **no player has a profitable unilateral deviation**.

Nash Equilibrium

Definition. The *best-reply correspondence* for player $i \in N$ is a set-valued function B_i such that:

$$B_i(s_{-i}) = \{s_i \in S_i \mid u_i(s_i, s_{-i}) \geq u_i(s'_i, s_{-i}), \forall s'_i \in S_i\}.$$

So that $B_i(s_{-i}) \subseteq S_i$ “tells” player i what to do when the other players play s_{-i} . Hence:

Definition. $s^* \in S$ is a Nash equilibrium if and only if $s_i^* \in B_i(s_{-i}^*)$ for all $i \in N$.

In words: a Nash equilibrium is a strategy profile of mutual best replies. Each player picks a best reply to the combination of strategies chosen by the other players.

Mixed Extension

Definition. The *mixed extension* of a game $\mathcal{G} = \langle N, \{S_i\}_{i \in N}, \{u_i\}_{i \in N} \rangle$ is the game Γ , where:

1. $\Gamma = \langle N, \{\Delta(S_i)\}_{i \in N}, \{U_i\}_{i \in N} \rangle$.
2. $\Delta(S_i)$ is the set of probability distributions over S_i , and $\Delta(S) = \times_{i \in N} \Delta(S_i)$.
3. $U_i : \Delta(S) \mapsto \mathcal{R}$ is a vNM expected utility function that assigns to each $\sigma \in \Delta(S)$ the expected value under u_i of the lottery over S induced by σ .

Consider part (3) for finite games. Suppose player i plays mixed strategy $\sigma_i \in \Delta(S_i)$. Denote the probability that this places on pure strategy $s_i \in S_i$ as $\sigma_i(s_i)$. Then,

$$U_i(\sigma) = \sum_{s \in S} u_i(s) \prod_{j \in N} \sigma_j(s_j).$$

Notation. Define $\sigma_{-i} \in \Delta(S_{-i}) = \times_{j \neq i} \Delta(S_j)$ analogously to the pure strategy case.

A Coordination Game

Stag Hunt

	<i>H</i>	<i>S</i>
<i>H</i>	<u>2</u> , <u>2</u>	0, 0
<i>S</i>	0, 0	<u>3</u> , <u>3</u>

NE:

- ▶ $(H, H), (S, S)$
- ▶ $(\frac{1}{3}, \frac{2}{3})$ for both players.

An Anti-Coordination Game

Hawk Dove

	<i>Hawk</i>	<i>Dove</i>
<i>Hawk</i>	-2	<u>0</u>
<i>Dove</i>	<u>4</u>	0

NE:

- ▶ $(Hawk, Dove), (Dove, Hawk)$
- ▶ $(\frac{2}{3}, \frac{1}{3})$ for both players.

Existence of Nash Equilibrium*

Theorem. (Nash 1951) *Every finite strategic-form game has at least one Nash equilibrium.*

Recall that $\sigma^* \in \Delta(S)$ is a Nash equilibrium if and only if $\sigma_i^* \in B_i(\sigma_{-i}^*)$ for all $i \in N$.

This condition can be expressed as $\sigma^* \in B(\sigma^*)$, where the best-reply correspondence $B(\sigma) = \times_{i \in N} B_i(\sigma)$. That is, σ^* must be a fixed point of the best-reply correspondence.

Hence proving the existence of a Nash equilibrium reduces to proving the existence of a fixed point.

Brouwer's fixed-point theorem guarantees the existence of a fixed point for certain *functions*, but does not apply to correspondences. The relevant theorem is due to Kakutani.

Kakutani's Fixed Point Theorem*

Theorem. A correspondence $g : X \rightarrow X$ has a fixed point $x \in X$ if:

1. X is a compact, convex and non-empty set of \mathcal{R}^n .
2. $g(x)$ is nonempty for all x .
3. $g(x)$ is convex for all x .
4. g has a closed graph (no holes).

To prove the existence of a Nash equilibrium, we need to show that every finite strategic-form game satisfies the conditions of Kakutani's theorem.

– Nash equilibria have also been shown to exist when strategy sets are continuous under certain conditions (e.g. continuity of payoffs).

Requirements for Nash Equilibrium Play

The traditional justification for NE as a “rational” outcome of a game requires three things.

1. Players are **rational**:

- ▶ They can formulate *strategies* that maximize their payoff given what everyone else is doing.

2. Players have **knowledge** of the game they are playing:

- ▶ They know the strategy set,
- ▶ They know the payoffs generated by each strategy profile.

3. Players have **equilibrium knowledge**:

- ▶ They can correctly anticipate what other players will do.

Requirements for Nash Equilibrium Play

All of these requirements pose problems:

1. Players are **rational**:

- ▶ They can formulate *strategies* that maximize their payoff given what everyone else is doing. (**consider chess**)

2. Players have **knowledge** of the game they are playing:

- ▶ They know the strategy set, (**consider sports**)
- ▶ They know the payoffs generated by each strategy profile. (**consider traffic congestion**)

3. Players have **equilibrium knowledge**:

- ▶ They can correctly anticipate what other players will do. (**consider games with multiple equilibria such as Stag Hunt and Hawk Dove**)

The Problem of Equilibrium Knowledge

In economics, typically prediction = equilibrium.

Are the following **disequilibria** not reasonable?

- ▶ Player 1 chooses *stag* expecting player 2 to choose *stag*. But, player 2 chooses *hare*, expecting player 1 to choose *hare*.
- ▶ Player 1 chooses *hawk* expecting player 2 to choose *dove*. But, player 2 chooses *hawk*, expecting player 1 to choose *dove*.

These are both collections of **rationalizable strategies**:

- ▶ Each player's strategy is a best response to some *rationalizable belief* about the other player's strategy.

The Problem of Equilibrium Selection

Another problem is that of **multiple equilibria**.

- ▶ Nash theorem guarantees the existence of at least one NE in a large class of games.
- ▶ However, it does NOT in general make a *unique* prediction of play.
- ▶ Which NE is the most “plausible” prediction?
 - ▶ The equilibrium refinements program emerged to address this question.
 - ▶ Placed unreasonable burdens on the rationality of players.

Coordination & Focal Points

- ▶ Schelling's experiments:
 - Two people are to meet in New York city, but have forgotten the time and place. Where would they try to meet and when?
 - The strategy space would seem to be impossibly large.
 - Yet under the clock at Grand Central Station at 12 noon was chosen by the majority of students surveyed.
- ▶ Schelling's theory of **focal points**: some equilibria have intrinsic properties that make them focal or have become conventional through recurrent play.

Classical Game Theory

We repeat most emphatically that our theory is thoroughly static. A dynamic theory would unquestionably be more complete and preferable.

von Neumann and Morgenstern (1944, p. 44-45)

Learning to Play Nash

Fictitious Play (Brown 1951):

- ▶ Start with a random initial history of choices by row and column players.
- ▶ In round 1 of play, row player is able to revise his strategy best responding to column's history of choices.
- ▶ In round 2 of play, column player is able to revise his strategy best responding to row's history of choices.
- ▶ Repeat ...

If the process converges — for some finite period T , row chooses some strategy s and column chooses some strategy s' (possibly s) for all $t \geq T$ — then (s, s') is a Nash equilibrium.

But the process does not always converge. In some games, players may never learn to play a Nash equilibrium via fictitious play.

Evolution and Learning in Games

We shall now take up the “mass-action” interpretation of equilibrium points... It is unnecessary to assume that the participants have full knowledge of the total structure of the game, or the ability and inclination to go through any complex reasoning processes. But the participants are supposed to accumulate empirical information on the relative advantages of the various pure strategies at their disposal.

John Nash (PhD thesis, p. 21)

An Evolutionary Approach

Partly in response to the shortcomings of the equilibrium refinements program, a new field emerged known as **evolutionary game theory**.

The evolutionary approach to the social sciences is based on:

- ▶ **boundedly rational**
- ▶ **populations** of agents,
- ▶ who may (or may not) **learn** or **evolve** their way into equilibrium,
- ▶ by **gradually revising**
- ▶ **simple, myopic rules** of behavior.

Bounded Rationality

This is a model of **bounded rationality** in which knowledge and computation are distributed.

Learning about the environment and the behavior of other players takes place at the population level through dynamic processes including:

- ▶ natural selection
- ▶ imitation
- ▶ reinforcement learning
- ▶ Bayesian social learning
- ▶ myopic best responses
- ▶ cultural transmission.

Out-of-Equilibrium Dynamics

By specifying an explicit dynamic process of adjustment to equilibrium, analyses of evolution and learning in games enable us to address the following questions:

1. Are equilibria stable in the face of random shocks? (**local stability**)
2. Does evolution/learning lead to an equilibrium from any initial state? (**global convergence and nonconvergence**)
3. To which equilibria, if any, does an evolutionary dynamic lead? (**equilibrium selection**)

Population Games

Players

- ▶ The population is a set of agents (possibly a continuum).

Strategies

- ▶ The set of (pure) strategies is $S = \{1, \dots, n\}$, with typical members i, j and s .
- ▶ The mass of agents choosing strategy i is m_i , where $\sum_{i=1}^n m_i = m$.
- ▶ Let $x_i = \frac{m_i}{m}$ denote the proportion of players choosing strategy $i \in S$.

Population States

- ▶ The set of population states (or strategy distributions) is $X = \{x \in \mathbb{R}_+^n : \sum_{i \in S} x_i = 1\}$.
- ▶ X is the unit simplex in \mathbb{R}^n .
- ▶ The set of vertices of X are the pure population states—those in which all agents choose the same strategy.
- ▶ These are the standard basis vectors in \mathbb{R}^n :

$$e_1 = (1, 0, 0, \dots), e_2 = (0, 1, 0, \dots), e_3 = (0, 0, 1, \dots), \dots$$

Payoffs

- ▶ A *continuous* payoff function $F : X \rightarrow \mathbb{R}^n$ assigns to each population state a vector of payoffs, consisting of a real number for each strategy.
- ▶ $F_i : X \rightarrow \mathbb{R}$ denotes the payoff function for strategy i .

Equivalence to Mixed Strategies

Consider the expected payoff to strategy i if i is matched with another strategy drawn uniformly at random from the population to play the following *two-player* game:

	1	2	...	n
i	$u(i, 1)$	$u(i, 2)$...	$u(i, n)$

The expected payoff to strategy i in state x is:

$$\begin{aligned} F_i(x) &= x_1 u(i, 1) + x_2 u(i, 2) \dots + x_n u(i, n) \\ &= \sum_{j=1}^n x_j u(i, j) \\ &= \sum_{j=1}^n x_j F_i(e_j). \end{aligned}$$

Nash Equilibria of Population Games

x^* is a Nash equilibrium of the population game if

$$(x^* - x)'F(x^*) \geq 0 \text{ for all } x \in X.$$

- ▶ **Monomorphic equilibria:** $x^* = e_i$.
- ▶ **Polymorphic equilibria:** $x^* \neq e_i$ for some $i \in S$; requires $F_i(x^*) = F_j(x^*) \geq F_k(x^*)$ for all i, j in support of x^* and k not in the support of x^* .