Optimization Theory and Methods

2025 Autumn

同济经管 TONGJI SEM

魏可佶 kejiwei@tongji.edu.cn
https://kejiwei.github.io/

CAMEA 中国高质量MBA教育认证 AACSB ACCREDITED EQUIS ACCREDITED

- Central Idea
- Formal Definition
- Examples with a unique PSNE
- Examples without a PSNE
- Examples with multiple PSNE
- Nash Equilibrium as a Result of Learning or Evolution
- Existence Properties

- If there exists a combination of strategies of all players such that no player can improve its payoff by deviating from its respective strategy in that combination, then no one would want to deviate.
- Such strategy profile is considered to be *stable*, and hence it is one potential way of predicting the outcome of the game.
- E.g. $\{stag, stag\}$ profile from the stag hunt game:

Player 2 Decision

Player 1 Decision		**Stag**	**Hare**
	Stag	2,2	0,1
	Hare	1,0	1,1

  - For a player deviating to $hare$ strategy payoff reduces from 2 to 1. So there is no incentive to deviate.

- A mixed strategy profile σ^* is a Nash equilibrium, if for all players i, $u_i(\sigma_i{}^*, \sigma_{-i}{}^*) \geq u_i(s_i, \sigma_{-i}{}^*)$, for all $s_i \in S_i$.
- This is same as saying that for a given set of strategies of all other players (denoted by $\sigma_{-i}{}^*$), the best that player i can do is to play $\sigma_i{}^*$ strategy.
- To prove that a strategy profile is a Nash equilibrium, it is sufficient to compare its payoff to only the pure strategies.
- In other words, it is not required to replace the s_i on right hand side of the expression above with σ_i. (Why?)
- In the next few slides, we will examine various ways of finding Nash Equilibria (NE) for various types of games. We will look at both Pure Strategy Nash Equilibria (PSNE) and Mixed Strategy Nash Equilibria (MSNE).

- For finite games (i.e., games where each player has a finite number of pure strategies to choose from), a PSNE is typically found by examining one strategy profile at a time.
- Consider the stag hunt example again. We already saw that $\{stag, stag\}$ is an equilibrium.

Player 2 Decision

Player 1 Decision		Stag	Hare
	Stag	2,2	0,1
	Hare	1,0	1,1

- It can be seen that $\{hare, hare\}$ is also a PSNE: If a player switches to $stag$ strategy then payoff drops from 1 to 0.
- Also, we can see that $\{stag, hare\}$ or $\{hare, stag\}$ is not a PSNE because the player who chooses stag can increase its payoff by switching to $hare$.

- Prisoner's Dilemma: Two people are arrested for a crime.
- Police lack sufficient evidence to convict either suspect and so need them to testify against each other.
- Each suspect is put in a different cell so they can't communicate with each other.
- The police tell each suspect that if he testifies against the other, he will receive a reward.
- Additionally, if he testifies and the other does not testify, he will be set free and the other will go to prison.
- If both testify, both go to prison but will also collect rewards.
- If neither testifies, then they will be released due to insufficient evidence.

↳ Prisoner's Dilemma Contd...

- There are 2 strategies for each prisoner. One is to testify (i.e., defect: D), other is not to testify (i.e., cooperate: C).
- The payoff matrix for prisoner's dilemma is as follows:

	C	D
C	1,1	-1,2
D	2,-1	0,0

- (C, C) is not a PSNE, because if a player defects then his payoff will increase from 1 to 2.
- (C, D) and (D, C) are also not PSNEs because in each case, the player who cooperates can switch to defecting and increase payoff from -1 to 0.
- The only PSNE is (D, D). If any player switches from defect to cooperate then the payoff will decrease from 0 to -1.

- In games with infinite number of pure strategies (called Infinite Games), typical way of solving for a PSNE is to solve simultaneous equations obtained from first order conditions.
- Consider an example with 2 players where each player has infinite pure strategies.
- Cournot Competition model: Two firms simultaneously decide on the output $q_i \in [0, \infty)$ to be produced.
- The price p in a competitive market is a decreasing function of total quantity produced $q = q_1 + q_2$.
- Consider linear production costs $c_i(q_i) = cq_i$ (with $0 \leq c \leq 1$) and assume linear price function $p(q) = \max(0{,}1 - q)$.
- So, profit of firm i is $u_i(q_1, q_2) = q_i * (\max(0{,}1 - q_1 - q_2)) - cq_i$.

- Assuming $q_1 + q_2 < 1$, we get the first order conditions as $0 = \frac{\partial}{\partial q_i}\left(u_i(q_1, q_2)\right) = \frac{\partial}{\partial q_i}(q_i * (1 - q_1 - q_2) - cq_i) = 1 - q_{-i} - 2q_i - c$. So $q_i = -\frac{q_{-i}}{2} + \frac{1-c}{2}$.
- Thus, $q_1 = -\frac{q_2}{2} + \frac{1-c}{2}$ and $q_2 = -\frac{q_1}{2} + \frac{1-c}{2}$.
- Solving the 2 equations simultaneously, we get $q_1 = q_2 = \frac{1-c}{3}$.
- We still need to verify that $q_1 + q_2 < 1$:
  - $q = q_1 + q_2 = \frac{2}{3}(1 - c) \leq \frac{2}{3} < 1$ (since $0 \leq c \leq 1$).
- Also, note that if $q_1 + q_2 > 1$ then PSNE cannot exist because a player can switch to $q_i = 0$ and increase profits from negative to 0.
- Thus, the unique PSNE is found at $({q_1}^*, {q_2}^*) = \left(\frac{1-c}{3}, \frac{1-c}{3}\right)$.

↳ Multiple PSNE in Finite Game: Voting

- Each of the 3 players, $\{1,2,3\}$, has to choose exactly 1 alternative out of $\{A,B,C\}$. Alternative with most votes wins. If each alternative receives 1 vote, alternative A wins. The payoff functions are as follows: $u_1(A) = u_2(B) = u_3(C) = 2$, $u_1(B) = u_2(C) = u_3(A) = 1$, $u_1(C) = u_2(A) = u_3(B) = 0$.

Profile	Win	PSNE	Profile	Win	PSNE	Profile	Win	PSNE
(A,A,A)	A	Y	(B,A,A)	A	N (2)	(C,A,A)	A	N (2,3)
(A,A,B)	A	N (2)	(B,A,B)	B	N (1,3)	(C,A,B)	A	N (2,3)
(A,A,C)	A	N (2)	(B,A,C)	A	N (2)	(C,A,C)	C	N (1)
(A,B,A)	A	Y	(B,B,A)	B	N (1)	(C,B,A)	A	N (2,3)
(A,B,B)	B	N (3)	(B,B,B)	B	Y	(C,B,B)	B	N (3)
(A,B,C)	A	N (2)	(B,B,C)	B	N (1)	(C,B,C)	C	N (1)
(A,C,A)	A	N (3)	(B,C,A)	A	N (2,3)	(C,C,A)	C	N (1)
(A,C,B)	A	N (2,3)	(B,C,B)	B	N (1,3)	(C,C,B)	C	N (1,2)
(A,C,C)	C	Y	(B,C,C)	C	N (2)	(C,C,C)	C	Y

↳ No PSNE in Finite Game: Matching Pennies

- Two players simultaneously announce heads (H) or tails (T).
- If the announcements match then player 1 gets a payoff of 1 unit and player 2 gets a payoff of -1 units.
- If the announcements differ then player 1 gets a payoff of -1 units and player 2 gets a payoff of 1 unit.

	H	T
H	1,-1	-1,1
T	-1,1	1,-1

- It is easy to see that (H,H) and (T,T) is not a PSNE since player 2 will benefit from switching. Similarly, (H,T) and (T,H) is not a PSNE since player 1 will benefit from switching.
- This game has no PSNE.

↳ MSNE for Matching Pennies Game

- In any MSNE of this game, both players must randomize.
- Specifically, if player 1 has pure strategy then player 2 will prefer the opposite pure strategy; If player 2 has a pure strategy then player 1 will prefer the same pure strategy.
- Consider an MSNE with p as the probability of player 1 picking H and q as the probability of player 2 picking H.
- Since both H and T are in the support of the chosen mixed strategy for player 1, **the payoff of player 1 under each pure strategy must be the same** (otherwise player 1 would increase its payoff by shifting to a pure strategy).
- So, $u_1(H,\sigma_2) = u_1(T,\sigma_2) \Rightarrow q * 1 + (1-q) * (-1) = q * (-1) + (1-q) * 1 \Rightarrow q = \frac{1}{2}$.

- Similarly, since both H and T are in the support of the chosen mixed strategy for player 2, the payoff of player 2 under each pure strategy must be the same (otherwise player 2 would increase its payoff by shifting to a pure strategy).
- So, $u_2(\sigma_1, H) = u_2(\sigma_1, T) \Rightarrow p * (-1) + (1-p) * 1 = p * 1 + (1-p) * (-1) \Rightarrow p = \frac{1}{2}$.
- So the unique MSNE is at $\sigma_1 = \left(\frac{1}{2}, \frac{1}{2}\right)$ and $\sigma_2 = \left(\frac{1}{2}, \frac{1}{2}\right)$.
- This is a general method used to find any MSNE in a finite 2-player game.
- Note that when finding the mixed strategy for player 1 (i.e., when finding p) we equated the payoffs for player 2, and vice versa.

↳ Justification of a Nash Equilibrium

- Typically, Nash equilibrium is justified as an output of a process involving introspection and deduction by players.
- All players are assumed to be completely knowledgeable and “infinitely rational”. This means players have full knowledge of opponents’ payoffs, and knowledge that the opponent is rational.
- Furthermore, *each player knows that every other player knows* the same things, and that *each player knows that every other player knows that every other player knows* the same things, and so on.
- This is what we call as “common knowledge”.
- While these assumptions are somewhat extreme, many (not all) real world outcomes are very well described by Nash equilibrium.
- However, this “infinite rationality” is not the only way of justifying a Nash equilibrium.

- Players can predict opponent behavior by closely observing
  - the same opponent's behavior in *similar games* from past, or
  - *similar opponents' behaviors* from the past.
- Then they respond to that predicted behavior of the opponents in a rational way. This is what we call *learning*.
- Such learning process sometimes leads to an equilibrium.
- This is an alternative justification of a Nash equilibrium.
- E.g. consider a Cournot game where the first order conditions give the following: $q_1 = -\frac{q_2}{2} + \frac{1-c}{2}$ and $q_2 = -\frac{q_1}{2} + \frac{1-c}{2}$.
- These equations represent the *best response* of player 1 to player 2's strategy and vice versa.

- For $c = 0$, let us plot the two best response equations: $q_1 = r_1(q_2) = -\frac{q_2}{2} + \frac{1}{2}$ and $q_2 = r_2(q_1) = -\frac{q_1}{2} + \frac{1}{2}$.
- Assume that the process starts from an arbitrary point (q_1, q_2).
- In each iteration, one player adjusts strategy by responding to current strategy of other player.
- For example: start from $\left(\frac{1}{6}, 0\right)$.
- $q_2 = -\frac{1}{12} + \frac{1}{2} = \frac{5}{12}$. => $\left(\frac{1}{6}, \frac{5}{12}\right)$.
- $q_1 = -\frac{5}{24} + \frac{1}{2} = \frac{7}{24}$. => $\left(\frac{7}{24}, \frac{5}{12}\right)$.
- $q_2 = -\frac{7}{48} + \frac{1}{2} = \frac{17}{48}$. => $\left(\frac{7}{24}, \frac{17}{48}\right)$, etc.
- This converges to PSNE $(q_1{}^*, q_2{}^*) = \left(\frac{1}{3}, \frac{1}{3}\right)$.

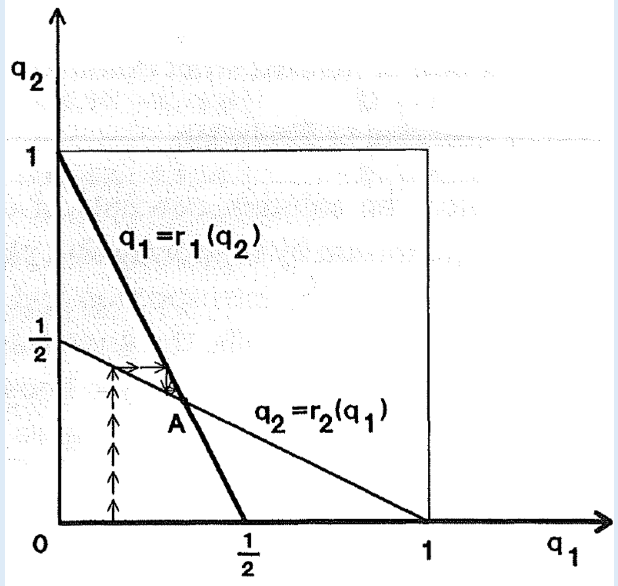

- Fact 1: Every finite strategic-form game has at least one NE.
  - Note that this can be either a PSNE or a MSNE.
  - Recall: A finite game is one where each player has finite number of possible strategies.
- Fact 2: If the players' strategy sets are compact and the payoffs are continuous, then at least one NE exists.
  - Note that this can be either a PSNE or a MSNE.
  - Recall: A compact set is one which is closed and bounded.
- Fact 3: If the players' strategy sets are compact and convex, and the payoff functions are continuous and concave, then at least one PSNE exists.

↳ Stag Hunt Motivation

- Note that situations similar to stag hunt game happen often.
- One common example is where two competing firms with identical costs and identical products decide their prices. Two alternatives are stag (high price) and hare (low price).
- If the prices for both are high then both get half the market share and both have high profits.
- If the prices for both are low then both get half the market share and both have moderate profits.
- If one has high price and the other has low price, then the high price firm gets a low market share and hence very little profit. The low price firm gets a high market share, but has to invest in more capacity to satisfy the extra demand and so earns moderate profit

Objective :

Key Concepts：