Algorithmic Game Theory

Algorithmische Spieltheorie

Nash Equilibria in Mixed Strategies

Wintersemester 2022/2023

Dozent: Prof. Dr. J. Rothe

Pure and Mixed Strategies

In all games so far, all players had to choose **exactly** one strategy:

- Smith and Wesson had to either *confess* or *remain silent* in the prisoners' dilemma;
- George and Helena had to go either to the *soccer match* or the *concert* in the battle of the sexes;
- David and Edgar could only either *swerve* or *go on driving* in the chicken game;
- in the penalty game, the kicker and the goalkeeper had each to choose one side of the goal, *left* or *right*;
- in the paper-rock-scissors game, David and Edgar would form upon *pon* either *paper*, *rock*, or *scissors* with their hands; and
- each player had to choose *exactly one number* in the guessing numbers game.

Pure and Mixed Strategies

- All players play *pure strategies* in these games.
- However, if one such game is played several times in a row, the players might change their minds and choose different strategies.
- It would be pretty dull in certain games to always decide for the same strategy. For example, a goalkeeper who always jumps to the left side will be very predictable; instead he should choose randomly where to jump, sometimes to the left, sometimes to the right.
- If the players make their decisions on which strategy to choose randomly under some probability distribution, we say they use a *mixed strategy*.
- In many games, especially so in those with mixed strategies, one does not win by intelligence only, one has also to be lucky.

Nash Equilibrium in Mixed Strategies

Definition (Nash equilibrium in mixed strategies)

Let $\mathscr{S} = S_1 \times S_2 \times \cdots \times S_n$ be the set of strategy profiles of the n players in a noncooperative game in normal form and let g_i be the gain function of player i, $1 \leq i \leq n$. For simplicity, let us assume that all sets S_i are finite.

1. A *mixed strategy for player i* is a probability distribution π_i on S_i, where $\pi_i(s_j)$ is the probability of the event that i chooses the strategy $s_j \in S_i$. Let Π_i be the set of all probability distributions on S_i (so $\pi_i \in \Pi_i$). Let $\Pi = \Pi_1 \times \Pi_2 \times \cdots \times \Pi_n$.
2. The *expected utility of a mixed-strategy profile $\vec{\pi} = (\pi_1, \pi_2, \ldots, \pi_n)$ for player i* is

$$G_i(\vec{\pi}) = \sum_{\vec{s}=(s_1,\ldots,s_n)\in\mathscr{S}} g_i(\vec{s}) \prod_{j=1}^{n} \pi_j(s_j).$$

Nash Equilibrium in Mixed Strategies

Remark

- *Intuitively, to compute the expected utility of $G_i(\vec{\pi})$ for player i,*
  - *we first calculate the probability of reaching each outcome given $\vec{\pi}$, and*
  - *we then calculate the average of the gains of the outcomes weighted by the probabilities of each outcome.*
- *We assume players to be risk-neutral, i.e., they seek to maximize their expected utility.*
- *The support of a mixed strategy π_i for player i is the set of pure strategies $\{s_j \mid \pi_i(s_j) > 0\}$.*
  - *A pure strategy is the special case of a mixed strategy whose support is a singleton.*
  - *A strategy π_i is fully mixed if it has full support, i.e., every pure strategy $s_j \in S_i$ occurs in it with nonzero probability.*

Nash Equilibrium in Mixed Strategies

Definition (Nash equilibrium in mixed strategies—continued)

3. A mixed strategy $\pi_i \in \Pi_i$ is *player i's best response to the mixed-strategy profile $\vec{\pi}_{-i} = (\pi_1, \ldots, \pi_{i-1}, \pi_{i+1}, \ldots, \pi_n) \in \Pi_{-i}$ of the other players* if for all mixed strategies $\pi'_i \in \Pi_i$,

$$G_i(\pi_1, \ldots, \pi_{i-1}, \pi_i, \pi_{i+1}, \ldots, \pi_n) \quad \geq \quad G_i(\pi_1, \ldots, \pi_{i-1}, \pi'_i, \pi_{i+1}, \ldots, \pi_n). \quad (1)$$

4. A profile $\vec{\pi} = (\pi_1, \pi_2, \ldots, \pi_n)$ of mixed strategies is in a *Nash equilibrium in mixed strategies* if π_i is a best response to $\vec{\pi}_{-i}$ for all players i.

Nash Equilibrium in Mixed Strategies

Remark

- *That is, a profile $\vec{\pi} = (\pi_1, \pi_2, \ldots, \pi_n)$ of mixed strategies is in a Nash equilibrium in mixed strategies*

  if and only if

  no player i has a mixed strategy $\pi'_i \in \Pi_i$ that would give her a higher profit than her mixed strategy π_i on S_i in response to the mixed strategies she expects the other players to choose.

- *For each player, one-sided deviation from their mixed strategies would thus be not beneficial (and might even be punished), assuming that the other players stick to their mixed strategies of the Nash equilibrium.*

Nash Equilibrium in Mixed Strategies

Remark (continued)

- *Consequently, for a Nash equilibrium in mixed strategies, every player is* **indifferent** *to each strategy she chooses with positive probability in her mixed strategy (i.e., to each strategy in her support).*
- *Also, the players' probability distributions in the profile of their mixed strategies are independent. (Compare: "correlated equilibrium.")*

Theorem

1. *Let $\vec{\pi} = (\pi_1, \pi_2, \ldots, \pi_n)$ be a profile of mixed strategies in a noncooperative game in normal form. A mixed strategy π_i is a best response to the mixed-strategy profile $\vec{\pi}_{-i}$ if and only if all pure strategies in its support are best responses.*
2. *Every pure Nash equilibrium is also a mixed Nash equilibrium.*

Nash Equilibrium in Mixed Strategies

1. Why?
   - For a contradiction, suppose that a best response mixed strategy contains in its support a pure strategy that itself is not a best response.
   - Then the player's expected utility would be improved by decreasing the probability of the worst such pure strategy (increasing proportionally the remaining nonzero probabilities to fill the gap).
   - This contradicts that the given mixed strategy was a best response.
   - The converse is immediate.
2. Exercise.

Nash Equilibrium in Mixed Strategies

Remark

- *As the following examples demonstrate, the converse is not necessarily true: The existence of a Nash equilibrium in mixed strategies does not imply the existence of a Nash equilibrium in pure strategies.*
- *That is, there can exist Nash equilibria in mixed strategies in addition to those in pure strategies.*
- *In particular, Nash equilibria in mixed strategies may exist in games that have no Nash equilibrium in pure strategies at all.*

Penalty Game: Mixed-Strategy Nash Equilibrium

Table: The penalty game

		Goalkeeper	
		Left	Right
Kicker	Left	$(-1,1)$	$(1,-1)$
	Right	$(1,-1)$	$(-1,1)$

- There is *no* Nash equilibrium in pure strategies.
- However, there is a Nash equilibrium in mixed strategies if the kicker K and the goalkeeper G both randomize uniformly:

$$\pi_K = (\pi_K(L), \pi_K(R)) = (1/2, 1/2) = (\pi_G(L), \pi_G(R)) = \pi_G.$$

Modified Penalty Game: Mixed-Strategy Nash Equilibrium

Table: The penalty game with a goalkeeper acting awkwardly on the left

		Goalkeeper	
		Left	Right
Kicker	Left	$(0,0)$	$(1,-1)$
	Right	$(1,-1)$	$(-1,1)$

- Again, there is *no* Nash equilibrium in pure strategies.
- However, there is a Nash equilibrium in mixed strategies:

$$(\pi_K, \pi_G) = ((2/3, 1/3), (2/3, 1/3)).$$

Why?

Modified Penalty Game: Mixed-Strategy Nash Equilibrium

For a mixed strategy profile to be in Nash equilibrium,

- the kicker has to find a mixed strategy π_K that makes the goalkeeper **indifferent** against each support strategy in π_G, and
- conversely the goalkeeper has to find a mixed strategy π_G that makes the kicker **indifferent** against each support strategy in π_K.
- If the kicker chooses the left side, his gain is $0 \cdot \pi_G(L) + \pi_G(R) = \pi_G(R)$.
- If he chooses the right side, however, his gain is $\pi_G(L) - \pi_G(R)$.

Thus, the kicker is made **indifferent** against a shot on goal to the left or to the right if the goalkeeper mixes his strategies such that

$$\pi_G(R) = \pi_G(L) - \pi_G(R).$$

Modified Penalty Game: Mixed-Strategy Nash Equilibrium

Since π_G is a probability distribution, we in addition have

$$\pi_G(L) + \pi_G(R) = 1,$$

so the goalkeeper achieves the kicker's desired indifference by choosing

$$\pi_G(L) = \frac{2}{3} \quad \text{and} \quad \pi_G(R) = \frac{1}{3}.$$

This can be interpreted as the goalkeeper trying to make up for this deficit on the left side by jumping there more often.
He thus anticipates the fact that the kicker is more likely to try to catch him wrongfooted on his weak side.

Modified Penalty Game: Mixed-Strategy Nash Equilibrium

- Conversely, the goalkeeper's gain for a jump to the left is $0 \cdot \pi_K(L) - \pi_K(R) = -\pi_K(R)$.
- If he jumps to the right, however, his gain is $-\pi_K(L) + \pi_K(R)$.

Thus, the goalkeeper is made **indifferent** against a jump to the left or to the right if the kicker mixes his strategies such that

$$-\pi_K(R) = -\pi_K(L) + \pi_K(R),$$

which together with

$$\pi_K(L) + \pi_K(R) = 1$$

gives the solution of

$$\pi_K(L) = \frac{2}{3} \quad \text{and} \quad \pi_K(R) = \frac{1}{3}$$

for the kicker as well.

Modified Penalty Game: Mixed-Strategy Nash Equilibrium

This mixed strategy reflects the above-mentioned fact that the kicker is more likely to challenge the goalkeeper's weak left side.

According to inequality (1),

$$(\pi_K, \pi_G) = ((2/3, 1/3), (2/3, 1/3))$$

is a Nash equilibrium in mixed strategies, and it is the only one.

Paper-Rock-Scissors: Mixed-Strategy Nash Equilibrium

Table: The paper-rock-scissors game

		Edgar		
		Rock	Scissors	Paper
David	Rock	$(0,0)$	$(1,-1)$	$(-1,1)$
	Scissors	$(-1,1)$	$(0,0)$	$(1,-1)$
	Paper	$(1,-1)$	$(-1,1)$	$(0,0)$

- Again, there is *no* Nash equilibrium in pure strategies.
- However, there is a Nash equilibrium in mixed strategies:

$$(\pi_D, \pi_E) = ((1/3, 1/3, 1/3), (1/3, 1/3, 1/3)).$$

Battle of the Sexes: Mixed-Strategy Nash Equilibria

Table: The battle of the sexes

		Helena	
		Soccer	Concert
George	Soccer	$(\mathbf{10},\mathbf{1})$	$(0,0)$
	Concert	$(0,0)$	$(\mathbf{1},\mathbf{10})$

- Nash equilibria in pure strategies:

$$(\text{Soccer},\text{Soccer}) \text{ and } (\text{Concert},\text{Concert}).$$

- In addition, there is also a third Nash equilibrium in mixed strategies:

$$(\pi_G,\pi_H) = ((^{10}/_{11},{}^{1}/_{11}),(^{1}/_{11},{}^{10}/_{11})).$$

Battle of the Sexes: Mixed-Strategy Nash Equilibria

To determine this third Nash equilibrium, it is again enough

- for George to find a mixed strategy π_G that makes Helena **indifferent** against her two strategies, while
- conversely Helena mixes her pure strategies in a way that also George is **indifferent** against his possible actions.
- If George chooses the soccer match (denoted by S) instead of the concert (denoted by C), his gain is $10 \cdot \pi_H(S) + 0 \cdot \pi_H(C) = 10 \cdot \pi_H(S)$.
- If he chooses the concert, however, then he gains $0 \cdot \pi_H(S) + \pi_H(C) = \pi_H(C)$.

To make him **indifferent** against these two actions, Helena must mix her strategies such that

$$10 \cdot \pi_H(S) = \pi_H(C).$$

Battle of the Sexes: Mixed-Strategy Nash Equilibria

Due to $\pi_H(S) + \pi_H(C) = 1$, we have

$$\pi_H(S) = 1/11 \quad \text{and} \quad \pi_H(C) = 10/11.$$

Since the gain vectors are symmetric for George and Helena, it follows that George's mixed strategy is analogously calculated to be

$$\pi_G(S) = 10/11 \quad \text{and} \quad \pi_G(C) = 1/11.$$

For this symmetric Nash equilibrium in mixed strategies,

$$(\pi_G, \pi_H) = ((10/11, 1/11), (1/11, 10/11)),$$

George and Helena would both stick to their own favorite ten times and give in only at the eleventh evening to finally fulfill their beloved one's desire.

Battle of the Sexes: Mixed-Strategy Nash Equilibria

This, however, obviously causes trouble.

In each round of the game, both have to commit themselves to one option, either the soccer game or the concert.

If both are stubborn on ten out of eleven of their anniversaries and are gentle only once, they will spend only two of these special days together (assuming they choose different anniversaries to give in), and there is nothing in it for either of them.

The reason for this lies in the intensity they each prefer their own favorite strategy over their partner's favorite strategy: Both valuate their own favorite ten times as much than their partner's!

Battle of the Sexes: Mixed-Strategy Nash Equilibria

How would one have to change the gain vectors of George and Helena to obtain a Nash equilibrium in mixed strategies having the form

$$(\pi'_G, \pi'_H) = ((1/2, 1/2), (1/2, 1/2))?$$

This Nash equilibrium would enable them to take turns in following his or her desire, and their relationship would have been saved.

As one can see, for a relation to work it is important that both partners are not too selfishly focused on their own preferences, but are open also for their partner's suggestions.

Chicken Game: Mixed-Strategy Nash Equilibria

Table: The chicken game

		Edgar	
		Swerve	Drive on
David	Swerve	$(2,2)$	$(\mathbf{1},\mathbf{3})$
	Drive on	$(\mathbf{3},\mathbf{1})$	$(0,0)$

- Nash equilibria in pure strategies:

$$(\text{Drive on}, \text{Swerve}) \text{ and } (\text{Swerve}, \text{Drive on}).$$

- Again, there is a third Nash equilibrium in mixed strategies:

$$(\pi_D, \pi_E) = ((1/2, 1/2), (1/2, 1/2)).$$

Chicken Game: Mixed-Strategy Nash Equilibria

Interpreting the three Nash equilibria in this game as recommendations for action, one could advice the players to do the following (and wish them good luck in evaluating their opponents well!):

1. If you expect your opponent to be a chicken, then you should definitely go all out and win heroically.

   This corresponds to one of the two Nash equilibria in pure strategies.

2. If you expect your opponent to be undaunted by death and risk it all, then you should be wise and swerve. You won't win, but you'll survive at least.

   This corresponds to the other one of the two pure Nash equilibria.

3. If you can't judge your opponent well and just have no idea of what he is up to do, then you should toss a coin and go all out with heads, but cautiously swerve with tails. Maybe you win; if not, maybe you survive—good luck!

   This corresponds to the additional Nash equilibrium in mixed strategies.

Prisoners' Dilemma: More Mixed-Strategy Nash Equilibria?

Table: The prisoners' dilemma

		Wesson	
		Confession	Silence
Smith	Confession	$(-\mathbf{4}, -\mathbf{4})$	$(0, -10)$
	Silence	$(-10, 0)$	$(-2, -2)$

- Nash equilibrium in pure strategies:

$$(\text{Confession}, \text{Confession}).$$

- There exists no additional Nash equilibrium in mixed strategies.

Different Interpretations of Mixed-Strategy Nash Equilibria

What does it mean to play a mixed strategy?

- Randomize to *confuse* your opponent:
  - Penalty game
  - Paper-Rock-Scissors game
- Randomize when you are *uncertain* about the other players' actions:
  - Battle of the sexes
  - Chicken game
- Mixed strategies describe what might happen in *repeated play*:
  - Number/frequency of pure strategies in the limit
- Mixed strategies describe *population dynamics*:
  - Some players chosen from a population of players, each with deterministic (i.e., pure) strategies
  - A mixed strategy is the probability of picking a player who will play one pure strategy or another

Properties of Some Two-Player Games

Table: Properties of some two-player games

	Prisoners' dilemma	Battle of the sexes	Chicken game	Penalty game	Paper-Rock-Scissors game
Dominant strategies?	yes	no	no	no	no
Strictly dominant strategies?	yes	no	no	no	no
Number of NE in pure str.	1	2	2	0	0
Number of NE in mixed str.	1	3	3	1	1
Number of PO	3	2	3	4	9
PO $=$ NE?	no	yes	no	no	no

Nash's Theorem

Theorem (Nash (1950; 1951))

For each noncooperative game in normal form with a finite number of players each having a finite set of strategies, there exists a Nash equilibrium in mixed strategies.

- Nash provided two proofs of his celebrated result.
- We sketch the first and give a more detailed outline of the second.

Sketch of First Proof of Nash's Theorem

- Let $\mathscr{S} = S_1 \times S_2 \times \cdots \times S_n$ be the set of strategy profiles of the n players in a noncooperative game in normal form.
  All sets S_i are here assumed to be finite.

- How can the abstract notion of "strategy" (in pure and in mixed form) be made accessible to mathematical or, specifically, topological arguments?

  > How can, for example, the very concrete strategy *Drive on* in the chicken game be compared with another concrete strategy from a different game, such as *Confession* in the prisoners' dilemma or *Left* in the penalty game?

- Nash views pure strategies as the unit vectors in an appropriate real vector space; every strategy from S_i is thus in $\mathbb{R}^{m_i}$.

Sketch of First Proof of Nash's Theorem

- Strategies can then be mixed using the common operations in vector spaces:
  - Every mixed strategy is the linear combination of pure strategies, each weighted by a certain probability, and
  - since a mixed strategy corresponds to a probability distribution on S_i, these probabilities sum up to 1.
- Mathematically speaking, mixed strategies over S_i are the points of a *simplex*, which can be viewed as a convex subset of $\mathbb{R}^{m_i}$.
- Such a subset is said to be *convex* if the direct connection between any two points of this subset completely lies within this subset.

Sketch of First Proof of Nash's Theorem

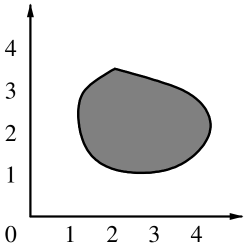

(a) a convex set

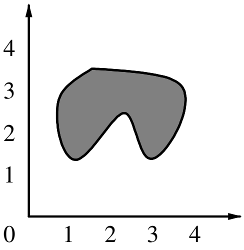

(b) a nonconvex set

Figure: A convex and a nonconvex set

Sketch of First Proof of Nash's Theorem

- In addition, the strategy sets are required to be *compact* (which is defined using the mathematical terms of closure and boundedness).
- Also the gain functions g_i, $1 \leq i \leq n$, mapping each strategy profile $\vec{s} = (s_1, s_2, \ldots, s_n) \in \mathscr{S}$ to a real number, must satisfy certain conditions so that known fixed point theorems from topology can be applied to them.
- To wit, it is required that the (multilinear) extensions of the functions g_i to the set of mixed strategies over $\mathscr{S}$ be *continuous* and *quasi-concave in s_j* for all j, $1 \leq j \leq n$.

  Continuity means that if there are only very small changes in the profiles of mixed strategies, then also the corresponding gains change only very little, i.e., there are no "jumps" (technically speaking, no points of discontinuity) in these gain functions.

Sketch of First Proof of Nash's Theorem

- A function $g : \mathbb{R} \to \mathbb{R}$ is said to be *quasi-convex* if all sets of the form

$$M_c = \{x \in \mathbb{R} \,\big|\, g(x) \leq c\}$$

  are convex, and

- a function $f : \mathbb{R} \to \mathbb{R}$ is said to be *quasi-concave* if its negation, $-f$, is quasi-convex. For example,
  - every monotonic function is both quasi-convex and quasi-concave, and
  - every function monotonically increasing up to a certain point and then monotonically decreasing is quasi-concave.

Sketch of First Proof of Nash's Theorem

- For concreteness, suppose that Anna and Belle play a two-player noncooperative game in normal form with
  - Anna having the pure strategies a_1, a_2, and a_3 and
  - Belle having the pure strategies b_1 and b_2.

Table: Anna's gain (left) and Belle's gain (right)

		Belle	
		Strategy b_1	Strategy b_2
Anna	Strategy a_1	$(1,2)$	$(1,2)$
	Strategy a_2	$(4,1)$	$(2,0)$
	Strategy a_3	$(3,0)$	$(4,3)$

Sketch of First Proof of Nash's Theorem

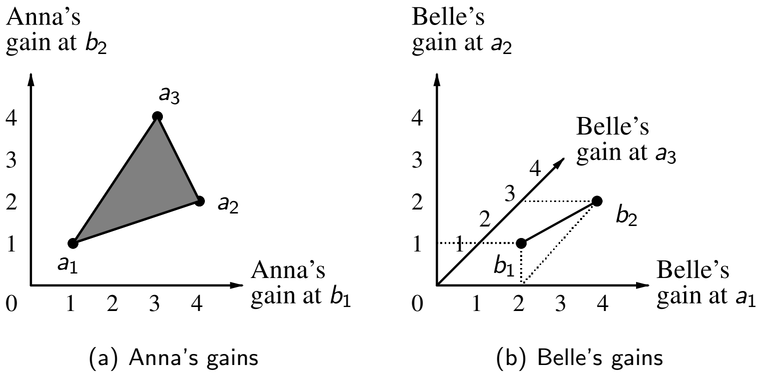

Figure: Convex gain sets for pure and mixed strategy sets

Sketch of First Proof of Nash's Theorem

- Since finite sets cannot be convex, the existence of a Nash equilibrium in *pure* strategies cannot be guaranteed by the proof of Nash's Theorem.
- The set of *mixed* strategies over $\mathscr{S}$ (including the pure strategies as special cases), however, is compact and convex and the extensions of the gain functions on these sets satisfy all required conditions, which makes certain fixed point theorems of topology applicable.
- It is then possible to define suitable transformations whose fixed points correspond to the Nash equilibria in mixed strategies.

Sketch of First Proof of Nash's Theorem

- For $\vec{\pi} = (\pi_1, \pi_2, \ldots, \pi_n) \in \Pi$ and every player i, a *best response correspondence $b_i(\vec{\pi}_{-i})$* is defined as a relation from the set of probability distributions Π_{-i} over the other players' strategies.
- Setting
$$b(\vec{\pi}) = b_1(\vec{\pi}_{-1}) \times b_2(\vec{\pi}_{-2}) \times \cdots \times b_n(\vec{\pi}_{-n})$$
and using the fixed point theorem of Kakutani, one can prove that b must have a fixed point under the hypotheses mentioned.
- That is, there exists a strategy profile $\vec{\pi}^*$ with $\vec{\pi}^* \in b(\vec{\pi}^*)$.

Sketch of First Proof of Nash's Theorem

- However, since $b(\vec{\pi})$ contains the best response strategies of all players to $\vec{\pi}$ by definition, this fixed point

$$\vec{\pi}^* \in b(\vec{\pi}^*)$$

  shows that the mixed strategies of all players in $\vec{\pi}^*$ are simultaneously in a Nash equilibrium in mixed strategies.

- No player has an incentive to deviate from her mixed strategy in $\vec{\pi}^*$, assuming that all other players stick to their strategies in $\vec{\pi}^*$ as well.

- This is the idea of the original proof of Nash's Theorem.

Nash's Second Proof: Some Basic Definitions

Definition

1. A set $X \subseteq \mathbb{R}^m$ is *convex* if for all $\vec{x}, \vec{y} \in X$ and for all real numbers $\lambda \in [0,1]$,
$$\lambda \cdot \vec{x} + (1-\lambda) \cdot \vec{y} \in X.$$

2. For vectors $\vec{x}_0, \vec{x}_1, \dots, \vec{x}_n \in \mathbb{R}^m$ and nonnegative scalars $\lambda_0, \lambda_1, \dots, \lambda_n$ satisfying $\sum_{i=0}^{n} \lambda_i = 1$, the vector
$$\sum_{i=0}^{n} \lambda_i \cdot \vec{x}_i$$
is a *convex combination of $\vec{x}_0, \vec{x}_1, \dots, \vec{x}_n$*.

3. A finite set $\{\vec{x}_0, \vec{x}_1, \dots, \vec{x}_n\}$ of vectors in $\mathbb{R}^m$ is said to be *affinely independent* if
$$\left(\sum_{i=0}^{n} \lambda_i \cdot \vec{x}_i = \vec{0} \text{ and } \sum_{i=0}^{n} \lambda_i = 0 \right) \quad \Rightarrow \quad \lambda_0 = \lambda_1 = \cdots = \lambda_n = 0.$$

Nash's Second Proof: Simplex

Definition

1. An *n-simplex* is the set of all convex combinations of the affinely independent set $\{\vec{x}_0, \vec{x}_1, \ldots, \vec{x}_n\}$ of vectors:

$$\vec{x}_0 \cdots \vec{x}_n = \left\{ \sum_{i=0}^{n} \lambda_i \cdot \vec{x}_i \,\middle|\, \lambda_i \geq 0 \text{ for each } i,\ 0 \leq i \leq n, \text{ and } \sum_{i=0}^{n} \lambda_i = 1 \right\}.$$

   (a) Every $\vec{x}_i$ is a *vertex of the n-simplex $\vec{x}_0 \cdots \vec{x}_n$*.
   (b) Every k-simplex $\vec{x}_{i_0} \cdots \vec{x}_{i_k}$, $i_0, \ldots i_k \in \{0, 1, \ldots, n\}$, is a *k-face of $\vec{x}_0 \cdots \vec{x}_n$*.

2. The *standard n-simplex Δ_n* is defined as

$$\Delta_n = \left\{ \vec{y} = (y_0, y_1, \ldots, y_n) \in \mathbb{R}^{n+1} \,\middle|\, y_i \geq 0, 0 \leq i \leq n, \text{ and } \sum_{i=0}^{n} y_i = 1 \right\}.$$

   That is, $\Delta_n = \vec{u}_0 \cdots \vec{u}_n$, where $\vec{u}_i$ denotes the i-th unit vector in $\mathbb{R}^{n+1}$.

Nash's Second Proof: Simplex

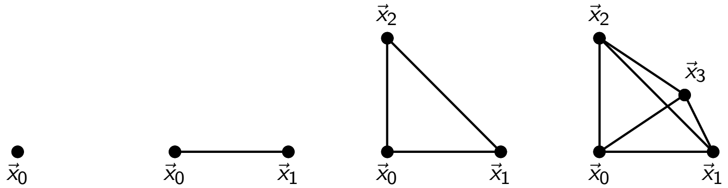

(a) 0-simplex (b) 1-simplex (c) 2-simplex (d) 3-simplex

Figure: n-simplexes for $0 \leq n \leq 3$

Nash's Second Proof: Simplicial Subdivision & Labeling

Definition

1. A *simplicial subdivision* of an n-simplex T is a finite set of simplexes $\{T_i \mid 1 \leq i \leq k\}$ such that
   (a) $\bigcup_{T_i \in T} T_i = T$ and
   (b) for each $T_i, T_j \in T$, $T_i \cap T_j$ is either empty or equal to a common face.

2.
   - Let $T = \vec{x}_0 \cdots \vec{x}_n$ be a simplicial subdivided n-simplex, and let V denote the set of all distinct vertices of all the subsimplexes.
   - For a point $\vec{y} \in T$, $\vec{y} = \sum_{i=0}^{n} \lambda_i \cdot \vec{x}_i$, let $\sigma(\vec{y}) = \{i \mid \lambda_i > 0\}$ be the set of vertices "involved" in $\vec{y}$.

   A function $\mathscr{L} : V \to \{0, 1, \ldots, n\}$ is a *proper labeling of a subdivision of T* if $\mathscr{L}(\vec{v}) \in \sigma(\vec{v})$.

3. A subsimplex of T is *completely labeled by $\mathscr{L}$* if $\mathscr{L}$ takes on all the values $0, 1, \ldots, n$ on its set of vertices.

Nash's Second Proof: Simplicial Subdivision & Labeling

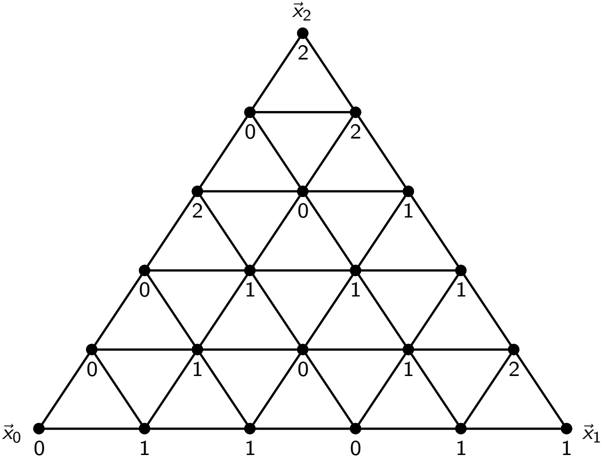

Figure: A properly labeled simplicial subdivision of a 2-simplex

Nash's Second Proof: Sperner's Lemma

Lemma (Sperner's Lemma)

Let $T = \vec{x}_0 \cdots \vec{x}_n$ be a simplicially subdivided n-simplex and let $\mathscr{L}$ be a proper labeling of the subdivision of T. There are an odd number of subsimplexes that are completely labeled by $\mathscr{L}$ in this subdivision of T.

Proof: The proof is by induction on n.

The base case, $n = 0$, holds trivially.

Indeed, the only simplicial subdivision of $T_0 = \vec{x}_0$ is $\{\vec{x}_0\}$, which can be labeled only by $\mathscr{L}(\vec{x}_0) = 0$, a proper labeling, so there is exactly one completely labeled 0-subsimplex of T_0, T_0 itself.

Nash's Second Proof: Sperner's Lemma

Suppose the claim holds true for $n-1$. We show that it also holds for n.

The given simplicial subdivision of the n-simplex $T_n = \vec{x}_0 \cdots \vec{x}_n$ induces a simplicial subdivision of its $(n-1)$-face

$$T_{n-1} = \vec{x}_0 \cdots \vec{x}_{n-1},$$

which is an $(n-1)$-simplex.

Furthermore, the labeling function $\mathscr{L}$ restricted to T_{n-1} is still proper.

By induction hypothesis, there are an odd number of $(n-1)$-subsimplexes in T_{n-1} with labels $0, 1, \ldots, n-1$.

Nash's Second Proof: Sperner's Lemma

Indeed, the 1-simplex $T_1 = \vec{x}_0\vec{x}_1$ has three 1-subsimplexes with labels 0 and 1 (and two 1-subsimplexes with labels 1 only):

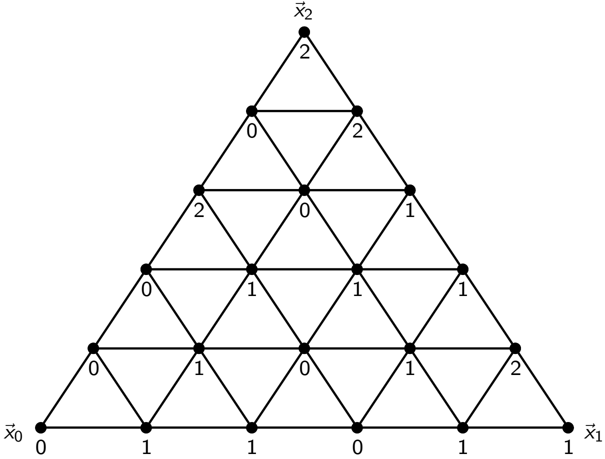

Nash's Second Proof: Sperner's Lemma

We now describe certain walks on T_n some of which will end in a completely labeled n-subsimplex of T_n.

The first type of walk we consider starts from T_{n-1}:

1. Start from any $(n-1)$-subsimplex in T_{n-1} with labels $0, 1, \ldots, n-1$. Call this $(n-1)$-subsimplex T'_{n-1}.
2. There is a unique n-subsimplex of T_n with $(n-1)$-face T'_{n-1}. Call this n-subsimplex T'_n. Walk into T'_n. Note that T'_n has the same vertices as T'_{n-1}, plus one additional vertex, say $\vec{z}$.

   Distinguish the following two cases.

   (2a) If $\mathscr{L}(\vec{z}) = n$, we have found a completely labeled n-subsimplex of T_n, namely T'_n, and the walk ends.

Nash's Second Proof: Sperner's Lemma

(2b) If $\mathscr{L}(\vec{z}) = j \neq n$, the $n+1$ vertices of T'_n have the labels $0, 1, \ldots, n-1$, so label j occurs twice and all other of these labels once.

We claim that in this case, T'_n has exactly one additional $(n-1)$-face, $T''_{n-1} \neq T'_{n-1}$, which is an $(n-1)$-subsimplex with labels $0, 1, \ldots, n-1$.

But this follows immediately from the fact that every $(n-1)$-face of T'_n has all vertices of T'_n except one.

Since only label j occurs twice, an $(n-1)$-face of T'_n has the labels $0, 1, \ldots, n-1$ if and only if one of the two vertices labeled by j is missing in it.

T'_{n-1} is one such $(n-1)$-face of T'_n, so there must be exactly another one, T''_{n-1}.

Continue the walk via T''_{n-1}. Again, distinguish the following two cases.

Nash's Second Proof: Sperner's Lemma

(2b.i) If T''_{n-1} belongs to an $(n-1)$-face of T_n, the walk ends.

(2b.ii) Otherwise, walk into the unique n-subsimplex of T_n having $(n-1)$-face T''_{n-1} with labels $0, 1, \ldots, n-1$.

Call this n-subsimplex T''_n and proceed as in the beginning of step 2, with T''_n and T''_{n-1} playing the roles of T'_n and T'_{n-1}, respectively.

The second type of walk we consider does not start from an $(n-1)$-subsimplex of T_{n-1} but from any completely labeled n-subsimplex of T_n, but otherwise follows the same rules, so only step 1 is skipped.

Walks of both types are uniquely and completely determined by their starting points:

- either on $(n-1)$-subsimplexes of T_{n-1} with labels $0, 1, \ldots, n-1$
- or on completely labeled n-subsimplexes of T_n.

Nash's Second Proof: Illustrating Sperner's Lemma

Two walks of the first type:

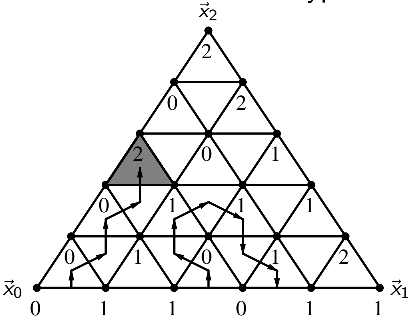

The walks end

- either in completely labeled n-subsimplexes of T_n: step (2a),
- or in $(n-1)$-subsimplexes of T_n's $(n-1)$-face T_{n-1}: step (2b.i).

They cannot end at another $(n-1)$-face of T_n because $\mathscr{L}$ is a proper labeling.

Each such walk can be reversed by essentially the same rules.

Nash's Second Proof: Illustrating Sperner's Lemma

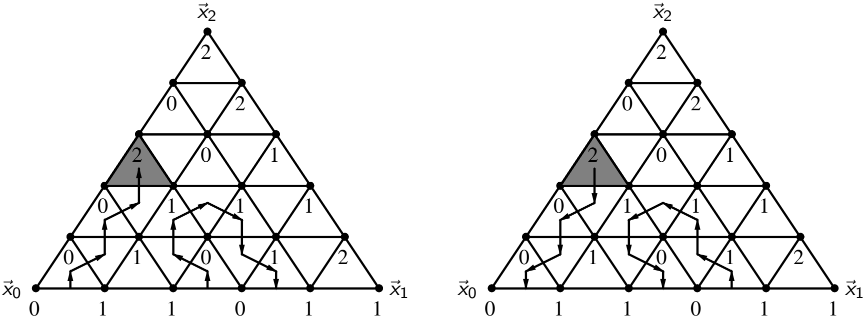

(a) Two walks of the first type (b) The same two walks reversed

Figure: Walking through the 2-simplex $T_2 = \vec{x}_0\vec{x}_1\vec{x}_2$

Nash's Second Proof: Sperner's Lemma

This implies that if a walk starts from an $(n-1)$-subsimplex T'_{n-1} on T_{n-1} and ends in an $(n-1)$-subsimplex T''_{n-1} on T_{n-1}, then $T'_{n-1} \neq T''_{n-1}$, for otherwise we could reverse this walk and would have two distinct walks with the same starting point, contradicting the uniqueness of walks.

Since the number of $(n-1)$-subsimplexes with labels $0, 1, \ldots, n-1$ on T_{n-1} is odd by the induction hypothesis, there are an odd number of walks starting from T_{n-1} and ending in a completely labeled n-subsimplex of T_n.

All these walks must end in *distinct* completely labeled n-subsimplexes of T_n, since otherwise they could be reversed, leading to distinct walks with the same starting point, again contradicting the uniqueness of walks.

Nash's Second Proof: Sperner's Lemma

Not all completely labeled n-subsimplexes of T_n can be reached by walks of the first type (i.e., by walks starting from an $(n-1)$-subsimplex of T_{n-1}).

However, such n-subsimplexes of T_n are connected by walks of the second type.

That is, all such completely labeled n-subsimplexes of T_n form pairs again, since (again by the reversal argument) they can neither be the starting points of walks leading to T_{n-1} (or any other $(n-1)$-face of T_n), nor the starting points of walks that return to themselves (forming cycles).

Nash's Second Proof: Illustrating Sperner's Lemma

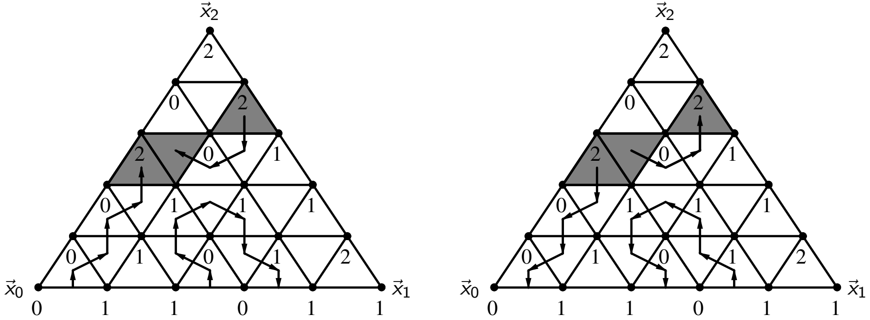

(a) Three walks, one of the second type (b) The same three walks reversed

Figure: All walks through the 2-simplex $T_2 = \vec{x}_0\vec{x}_1\vec{x}_2$

Summing up, we have shown that there are an odd number of completely labeled n-subsimplexes of T_n. ❑

Nash's Second Proof: Compact Set & Centroid

Definition (compact set)

A subset of $\mathbb{R}^m$ is *compact* if it is closed and bounded.

Remark

- *Δ_m is compact.*
- *A compact set has the property that every infinite sequence has a convergent subsequence.*

Definition (centroid)

The *centroid of an n-simplex $\vec{x}_0 \cdots \vec{x}_n$* is the "average" of its vertices:

$$\frac{1}{n+1} \sum_{i=0}^{n} \vec{x}_i .$$

Nash's Second Proof: Brouwer's Fixed Point Theorem

Theorem (Brouwer's Fixed Point Theorem)

Every continuous function $f : \Delta_m \rightarrow \Delta_m$ has a fixed point, i.e., there exists some $\vec{z} \in \Delta_m$ such that

$$f(\vec{z}) = \vec{z}.$$

Proof: The proof proceeds in two parts:

1. We construct a simplicial subdivision with a proper labeling function $\mathscr{L}$ for Δ_m so that Sperner's lemma can be applied, yielding at least one completely labeled m-subsimplex in this subdivision.
2. Making such subdivisions finer and finer, we show that this m-subsimplex contracts to a fixed point of f.

Nash's Second Proof: Brouwer's Fixed Point Theorem

For the first part:

- Fix an $\varepsilon > 0$.

- Subdivide Δ_m simplicially such that the *Euclidean distance* between any two points $\vec{x} = (x_0, \ldots, x_m)$ and $\vec{y} = (y_0, \ldots, y_m)$ in $\mathbb{R}^{m+1}$ in the same m-subsimplex of this subdivision is at most ε:

$$\sqrt{(x_0 - y_0)^2 + \cdots + (x_m - y_m)^2} \leq \varepsilon.$$

- We here assume that it is always possible to find such a simplicial subdivision of Δ_m, regardless of the dimension m, which is true, but not trivial to show.

Nash's Second Proof: Brouwer's Fixed Point Theorem

- Now define a labeling function $\mathscr{L} : V \to \{0, 1, \dots, m\}$ as follows.

  For each vertex $\vec{v} \in V$ of the m-subsimplexes in this subdivision, we choose a label $\mathscr{L}(\vec{v})$ from the set

$$\sigma(\vec{v}) \cap \{i \mid f_i(\vec{v}) \leq v_i\},$$

  where

  - $\vec{v} = (v_0, v_1, \dots, v_m)$ and $f(\vec{v}) = (f_0(\vec{v}), f_1(\vec{v}), \dots, f_m(\vec{v}))$ are points in Δ_m,
  - $\sigma(\vec{v}) = \{i \mid v_i > 0\}$ for $\vec{v} = \sum_{i=0}^{m} v_i \cdot \vec{u}_i$, since $\vec{u}_i$ is the ith unit vector in $\mathbb{R}^{m+1}$.

  That is, $\mathscr{L}(\vec{v}) = i$ means that $v_i > 0$ and $f_i(\vec{v}) \leq v_i$.

Nash's Second Proof: Brouwer's Fixed Point Theorem

- We have to show that this labeling function is well-defined, i.e., that

$$\sigma(\vec{v}) \cap \{i \mid f_i(\vec{v}) \leq v_i\} \neq \emptyset.$$

- Intuitively, this is true because
  - $\vec{v}$ and $f(\vec{v})$ are points in Δ_m, so their components each add up to one by definition of Δ_m.
  - Thus there exists an i such that $f_i(\vec{v}) \leq v_i$, and this holds true even when restricted to $\sigma(\vec{v})$, so $v_i > 0$.

Nash's Second Proof: Brouwer's Fixed Point Theorem

- Formally, for a contradiction suppose that $\sigma(\vec{v}) \cap \{i \mid f_i(\vec{v}) \leq v_i\} = \emptyset$.
- Since $\vec{v}$ is a point in Δ_m (i.e., $\sum_{i=0}^{m} v_i = 1$) and $v_j > 0$ exactly if $j \in \sigma(\vec{v})$, we have

$$\sum_{j \in \sigma(\vec{v})} v_j = \sum_{i=0}^{m} v_i = 1.$$

- From our assumption we know that $f_j(\vec{v}) > v_j$ for each $j \in \sigma(\vec{v})$, which implies

$$\sum_{j \in \sigma(\vec{v})} f_j(\vec{v}) > \sum_{j \in \sigma(\vec{v})} v_j = 1. \qquad (2)$$

- However, since $f(\vec{v})$ is a point in Δ_m as well, we have

$$\sum_{j \in \sigma(\vec{v})} f_j(\vec{v}) \leq \sum_{i=0}^{m} f_i(\vec{v}) = 1,$$

contradicting (2). Thus $\mathscr{L}$ is well-defined.

Nash's Second Proof: Brouwer's Fixed Point Theorem

- By construction,

$$\mathscr{L}(\vec{v}) \in \sigma(\vec{v})$$

  for each $\vec{v} \in V$.

- Thus $\mathscr{L}$ is also proper.

- By Sperner's lemma, in this simplicial subdivision of Δ_m there exists at least one m-subsimplex T_m^{ε} that depends on ε and is completely labeled by $\mathscr{L}$.

Nash's Second Proof: Brouwer's Fixed Point Theorem

In the second part of the proof:

- We will show that when ε goes to zero, the resulting m-subsimplex

$$T_m^{\varepsilon} = \vec{t}_0 \cdots \vec{t}_m$$

  contracts to a fixed point of f.

- T_m^{ε} is completely labeled; without loss of generality, we may assume that $\mathscr{L}(\vec{t}_i) = i$. (Otherwise, we simply rename the labels accordingly.)

- Furthermore, by construction of $\mathscr{L}$, we have

$$f_i(\vec{t}_i) \leq (\vec{t}_i)_i \qquad (3)$$

  for each i, $0 \leq i \leq m$, where $(\vec{t}_i)_i$ denotes the ith component of $\vec{t}_i$.

Nash's Second Proof: Brouwer's Fixed Point Theorem

- For ε going to zero, we consider the infinite sequence of centroids in these completely labeled m-subsimplexes T_m^ε.
- Since Δ_m is compact, there exists a convergent subsequence with limit $\vec{z}$.
- The vertices of these m-subsimplexes T_m^ε then move toward $\vec{z}$ with ε going to zero, that is, $\vec{t}_i \underset{\varepsilon \to 0}{\longrightarrow} \vec{z}$ for each i, $0 \leq i \leq m$.

Nash's Second Proof: Brouwer's Fixed Point Theorem

- Since f is continuous, it follows from (3) that

$$f_i(\vec{z}) \leq \vec{z}_i$$

  for each i, $0 \leq i \leq m$.

- This implies that $f(\vec{z}) = \vec{z}$, as desired, since otherwise, by the same argument as used in the first part of this proof, we would have

$$1 = \sum_{i=0}^{m} f_i(\vec{z}) < \sum_{i=0}^{m} \vec{z}_i = 1,$$

  a contradiction. ❑

Nash's Second Proof: Brouwer's Fixed Point Theorem

Reminder: What we have shown is

Theorem (Brouwer's Fixed Point Theorem)

Every continuous function $f : \Delta_m \to \Delta_m$ has a fixed point, i.e., there exists some $\vec{z} \in \Delta_m$ such that

$$f(\vec{z}) = \vec{z}.$$

Corollary (Brouwer's Fixed Point Theorem, applied to simplotopes)

Let $K = \prod_{j=1}^{k} \Delta_{m_j}$ be a simplotope (i.e., a Cartesian product of simplexes). Every continuous function $f : K \to K$ has a fixed point. **without proof**

Nash's Second Proof

Theorem (Nash (1950; 1951))

For each noncooperative game in normal form with a finite number of players each having a finite set of strategies, there exists a Nash equilibrium in mixed strategies.

> *"A proof of this existence theorem based on Kakutani's generalized fixed point theorem was published in Proc. Nat. Acad. Sci. U.S.A., 36, pp. 48–49. The proof given here is a considerable improvement over that earlier version and is based directly on the Brouwer theorem."*
>
> *John F. Nash (1951)*

Nash's Second Proof

Proof:

- Let $\vec{\pi} = (\pi_1, \pi_2, \dots, \pi_n) \in \Pi$ be a profile of mixed strategies with the expected gain functions $G_i(\vec{\pi})$.
- Let $\mathscr{S} = S_1 \times S_2 \times \cdots \times S_n$ be the underlying set of pure strategy profiles, where each S_i is finite.
- For each pure strategy s_j of each player i, let $G_i(\vec{\pi}_{-i}, s_j)$ be i's gain when switching one-sidedly from π_i to s_j.
- Define the functions

$$\varphi_{ij}(\vec{\pi}) = \max(0, G_i(\vec{\pi}_{-i}, s_j) - G_i(\vec{\pi}))$$

for each i and j with $1 \leq i \leq n$ and $1 \leq j \leq \|S_i\|$.

Nash's Second Proof

- Since the expected gain functions are continuous, so is each function φ_{ij}.

- Now, define the function $f : \Pi \to \Pi$ by $f(\vec{\pi}) = \vec{\pi}' = (\pi'_1, \pi'_2, \dots, \pi'_n)$, where the modifications π'_i of π_i are defined by

$$\pi'_i(s_j) = \frac{\pi_i(s_j) + \varphi_{ij}(\vec{\pi})}{\sum_{s_k \in S_i} (\pi_i(s_k) + \varphi_{ik}(\vec{\pi}))} = \frac{\pi_i(s_j) + \varphi_{ij}(\vec{\pi})}{1 + \sum_{s_k \in S_i} \varphi_{ik}(\vec{\pi})}. \tag{4}$$

- Intuitively, $\vec{\pi}'$ puts more probability weight π'_i on those pure strategies of each player i that are "better" responses to the other players' mixed strategies $\vec{\pi}'_{-i}$.

Nash's Second Proof

- Since every function φ_{ij} is continuous, so is f.
- Since Π, as a simplotope, is convex and compact and since $f : \Pi \rightarrow \Pi$ is continuous, f has at least one fixed point by Brouwer's fixed point theorem for simplotopes.
- It remains to show that $\vec{\pi}$ is a Nash equilibrium in mixed strategies if and only if $f(\vec{\pi}) = \vec{\pi}$.
- From left to right, if $\vec{\pi}$ is a Nash equilibrium in mixed strategies, we have $\varphi_{ij}(\vec{\pi}) = 0$ for all i and j.
- Hence, $f(\vec{\pi}) = \vec{\pi}' = \vec{\pi}$, so $\vec{\pi}$ is a fixed point.

Nash's Second Proof

- From right to left, suppose $f(\vec{\pi}) = \vec{\pi}$.
- Consider player i.
- Since G_i is linear in its ith component, there exists at least one pure strategy s_j in the support of π_i (i.e., $\pi_i(s_j) > 0$) such that

$$G_i(\vec{\pi}_{-i}, s_j) \leq G_i(\vec{\pi}).$$

In other words, by linearity of G_i in its ith component, we see that the situation where for each pure strategy s_k (in the support of π_i) it holds that $G_i(\vec{\pi}_{-i}, s_k) > G_i(\vec{\pi})$ is impossible.

Nash's Second Proof

- By definition of φ_{ij}, it follows that $\varphi_{ij}(\vec{\pi}_{-i}, s_j) = 0$.
- Since $f(\vec{\pi}) = \vec{\pi}$, this enforces that

$$\pi'_i(s_j) = \pi_i(s_j).$$

That is, the enumerator in (4) simplifies to $\pi_i(s_j)$ (due to $\varphi_{ij}(\vec{\pi}_{-i}, s_j) = 0$) and it is positive because s_j is in the support of π_i.

- This implies, by simple arithmetic, that the denominator in (4) must be one. Consequently,

$$\sum_{s_k \in S_i} \varphi_{ik}(\vec{\pi}) = 0.$$

Nash's Second Proof

- This holds true for each player i and, in effect, for all i and k, we have $\varphi_{ik}(\vec{\pi}) = 0$.
- That is, no player i can increase her gain by moving one-sidedly from her mixed strategy π_i to some pure strategy.
- However, since we know that

$$\max_{\pi_i' \in \Pi_i} G_i(\vec{\pi}_{-1}, \pi_i') = \max_{s_j \in \Pi_i} G_i(\vec{\pi}_{-1}, s_j)$$

from the theorem on slide 8, this means that $\vec{\pi}$ is a Nash equilibrium in mixed strategies. ❑

Nash's Theorem

Nash has won numerous prizes and awards and has been loaded with the highest academic honors for his superb insights and pathbreaking ideas, such as

- the 1978 *John von Neumann Theory Prize* for inventing the equilibria in noncooperative games named after him and
- the 1994 *Nobel Prize in Economics* (jointly with the game theoreticians Reinhard Selten and John Harsanyi).

"That's trivial, you know.
That's just a fixed point theorem."

John von Neumann (1950)