Algorithmic game theory

Martin Balko

2nd lecture

October 7th 2025

Proof of Nash's Theorem

Nash equilibria in normal-form games

	Rock	Paper	Scissors
Rock	(0,0)	(-1,1)	(1,-1)
Paper	(1,-1)	(0,0)	(-1,1)
Scissors	(-1,1)	(1,-1)	(0,0)

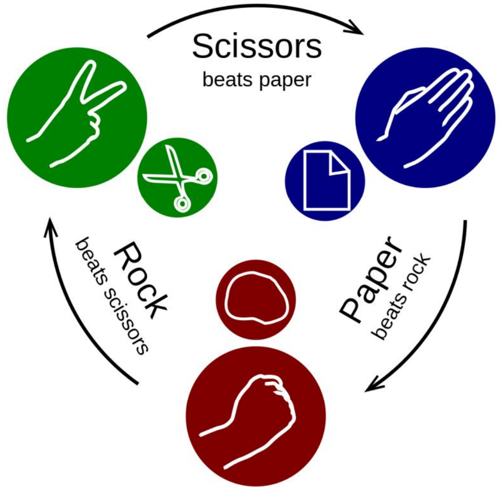

Sources: https://en.wikipedia.org/

- We introduced perhaps the most influential solution concept, which captures a notion of stability.
- The best response of player i to a strategy profile s_{-i} is a mixed strategy s_i^* such that $u_i(s_i^*; s_{-i}) \geq u_i(s_i'; s_{-i})$ for each $s_i' \in S_i$.
- For a normal-form game $G = (P, A, u)$ of n players, a Nash equilibrium (NE) in G is a strategy profile $(s_1, \ldots, s_n)$ such that s_i is a best response of player i to s_{-i} for every $i \in P$.
- Amazingly, every normal-form game has a Nash equilibrium.

Nash's Theorem

Nash's Theorem (Theorem 2.16)

Every normal-form game has a Nash equilibrium.

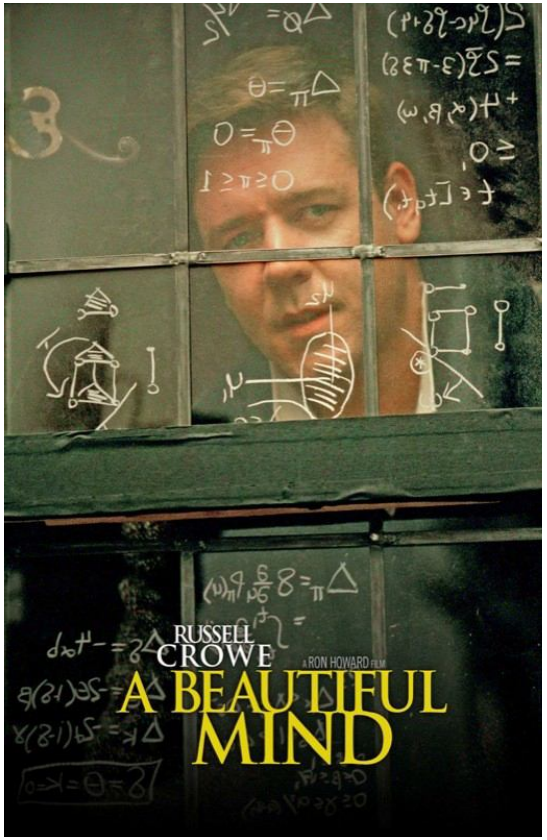

Figure: John Forbes Nash Jr. (1928–2015) and his depiction in the movie A Beautiful mind.

Sources: https://britannica.com and https://medium.com

Preparations for the proof of Nash's theorem

- The proof is essentially topological, as its main ingredient is a fixed-point theorem. We use a theorem due to Brouwer.

- For $d \in \mathbb{N}$, a subset X of $\mathbb{R}^d$ is compact if X is closed and bounded.
- We say that a subset Y of $\mathbb{R}^d$ is convex if every line segment containing two points from Y is fully contained in Y. Formally: for all x, y from Y, $tx + (1-t)y \in Y$ for every $t \in [0, 1]$.
- For n affinely independent points $x_1, \ldots, x_n \in \mathbb{R}^d$, an $(n-1)$-simplex Δ_n on $x_1, \ldots, x_n$ is the set of convex combinations of the points $x_1, \ldots, x_n$. Each simplex is a compact convex set in $\mathbb{R}^d$.

Lemma (Lemma 2.18)

For $n, d_1, \ldots, d_n \in \mathbb{N}$, let $K_1, \ldots, K_n$ be compact sets, each K_i lying in $\mathbb{R}^{d_i}$. Then, $K_1 \times \cdots \times K_n$ is a compact set in $\mathbb{R}^{d_1+\cdots+d_n}$.

Brouwer's Fixed Point Theorem

- For each $d \in \mathbb{N}$, let K be a non-empty compact convex set in $\mathbb{R}^d$ and $f : K \to K$ be a continuous mapping. Then, there exists a fixed point $x_0 \in K$ for f, that is, $f(x_0) = x_0$.

Figure: L. E. J. Brouwer (1881–1966).

Source: https://arxiv.org/pdf/1612.06820.pdf

- https://www.youtube.com/watch?v=csInNn6pfT4&t=268s&ab_

Proof of Nash's Theorem I

- Let $G = (P, A, u)$ be a normal-form game of n players. Recall that S_i is the set of mixed strategies of player i.
- We want to apply Brouwer's theorem, thus we need to find a suitable compact convex body K and a continuous mapping $f\colon K \to K$ whose fixed points are NE in G.

- We start with K. Let $K = S_1 \times \cdots \times S_n$ be the set of all mixed strategies.
  - We verify that K is compact and convex.
  - By definition, each S_i is, a simplex which is compact and convex.
  - By Lemma 2.18, the set $K = S_1 \times \cdots \times S_n$ is compact.
  - For any strategy profiles $s = (s_1, \ldots, s_n), s' = (s'_1, \ldots, s'_n) \in K$ and a number $t \in [0, 1]$, the point

$$ts + (1-t)s' = (ts_1 + (1-t)s'_1, \ldots, ts_n + (1-t)s'_n)$$

is also a mixed-strategy profile in K. Thus, K is convex.

Proof of Nash's Theorem II

- We now find the continuous mapping $f\colon K \to K$.
- For every player $i \in P$ and action $a_i \in A_i$, we define a mapping $\varphi_{i,a_i}\colon K \to \mathbb{R}$ by setting

$$\varphi_{i,a_i}(s) = \max\{0, u_i(a_i; s_{-i}) - u_i(s)\}.$$

  - $\varphi_{i,a_i}(s) > 0$ iff i can improve his payoff by using a_i instead of s_i.
  - By the definition of u_i, this mapping is continuous.
- Given $s \in K$, we define a new "improved" strategy profile $s' \in K$ as

$$s'_i(a_i) = \frac{s_i(a_i) + \varphi_{i,a_i}(s)}{\sum_{b_i \in A_i}(s_i(b_i) + \varphi_{i,b_i}(s))} = \frac{s_i(a_i) + \varphi_{i,a_i}(s)}{1 + \sum_{b_i \in A_i}\varphi_{i,b_i}(s)}.$$

  - "Increase probability at actions that are better responses to s_{-i}."
  - $s' \in K$ as each $s'_i(a_i)$ lies in $[0,1]$ and $\sum_{a_i \in A_i} s'_i(a_i) = 1$.
- We then define f by setting $f(s) = s'$.

Proof of Nash's Theorem III

- Then, f is continuous, since the mappings φ_{i,a_i} are.
- It remains to show that fixed points of f are exactly NE in G. Then, Brouwer's theorem gives us a fixed point of f, which is NE in G.
- First, if s is NE, then all functions φ_{i,a_i} are constant zero functions and thus $f(s) = s$. So s is a fixed point for f.
- Second, assume that $s = (s_1, \ldots, s_n) \in K$ is a fixed point for f.
  - For any player i, there is $a_i' \in A_i$ with $s_i(a_i) > 0$ such that $u_i(a_i'; s_{-i}) \leq u_i(s)$. Otherwise, $u_i(s) < \sum_{a_i \in A_i} s_i(a_i) u_i(a_i; s_{-i})$, which is impossible by the linearity of the expected payoff.
  - Then, $\varphi_{i,a_i'}(s) = 0$ and we get $s_i'(a_i') = \frac{s_i(a_i')}{1+\sum_{b_i \in A_i} \varphi_{i,b_i}(s)}$.
  - Since s is a fixed point, we get $s_i'(a_i') = s_i(a_i')$ and, since $s_i(a_i') > 0$, the denominator in the denominator is 1. This means that $\varphi_{i,b_i}(s) = 0$ for every $b_i \in A_i$. It follows that s is NE as

$$u_i(s_i''; s_{-i}) = \sum_{b_i \in A_i} s_i''(b_i) u_i(b_i; s_{-i}) \leq \sum_{b_i \in A_i} s_i''(b_i) u_i(s) = u_i(s).$$

$\square$

Nash's Theorem: remarks

- Two pages worth of Nobel prize!

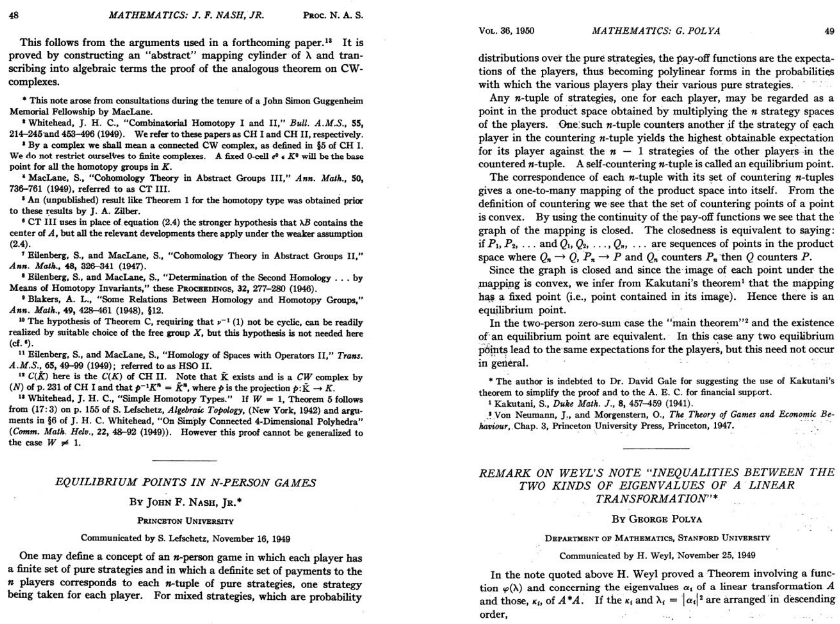

48 *MATHEMATICS: J. F. NASH, JR.* PROC. N. A. S.

This follows from the arguments used in a forthcoming paper.[13] It is proved by constructing an "abstract" mapping cylinder of λ and transcribing into algebraic terms the proof of the analogous theorem on CW-complexes.

* This note arose from consultations during the tenure of a John Simon Guggenheim Memorial Fellowship by MacLane.

[2] Whitehead, J. H. C., "Combinatorial Homotopy I and II," *Bull. A.M.S.*, **55**, 214–245 and 453–496 (1949). We refer to these papers as CH I and CH II, respectively.

[3] By a complex we shall mean a connected CW complex, as defined in §5 of CH I. We do not restrict ourselves to finite complexes. A fixed 0-cell $e^0 \in K^0$ will be the base point for all the homotopy groups in K.

[4] MacLane, S., "Cohomology Theory in Abstract Groups III," *Ann. Math.*, **50**, 736–761 (1949), referred to as CT III.

[5] An (unpublished) result like Theorem 1 for the homotopy type was obtained prior to these results by J. A. Zilber.

[6] CT III uses in place of equation (2.4) the stronger hypothesis that λB contains the center of A, but all the relevant developments there apply under the weaker assumption (2.4).

[7] Eilenberg, S., and MacLane, S., "Cohomology Theory in Abstract Groups II," *Ann. Math.*, **48**, 326–341 (1947).

[8] Eilenberg, S., and MacLane, S., "Determination of the Second Homology . . . by Means of Homotopy Invariants," these PROCEEDINGS, **32**, 277–280 (1946).

[9] Blakers, A. L., "Some Relations Between Homology and Homotopy Groups," *Ann. Math.*, **49**, 428–461 (1948), §12.

[10] The hypothesis of Theorem C, requiring that $\nu^{-1}(1)$ not be cyclic, can be readily realized by suitable choice of the free group X, but this hypothesis is not needed here (cf. [6]).

[11] Eilenberg, S., and MacLane, S., "Homology of Spaces with Operators II," *Trans. A.M.S.*, **65**, 49–99 (1949); referred to as HSO II.

[12] $C(\tilde{K})$ here is the $C(K)$ of CH II. Note that $\tilde{K}$ exists and is a CW complex by (N) of p. 231 of CH I and that $p^{-1}K^n = \tilde{K}^n$, where p is the projection $p:\tilde{K} \to K$.

[13] Whitehead, J. H. C., "Simple Homotopy Types." If $W = 1$, Theorem 5 follows from (17:3) on p. 155 of S. Lefschetz, *Algebraic Topology*, (New York, 1942) and arguments in §6 of J. H. C. Whitehead, "On Simply Connected 4-Dimensional Polyhedra" (*Comm. Math. Helv.*, **22**, 48–92 (1949)). However this proof cannot be generalized to the case $W \neq 1$.

EQUILIBRIUM POINTS IN N-PERSON GAMES

BY JOHN F. NASH, JR.*

PRINCETON UNIVERSITY

Communicated by S. Lefschetz, November 16, 1949

One may define a concept of an n-person game in which each player has a finite set of pure strategies and in which a definite set of payments to the n players corresponds to each n-tuple of pure strategies, one strategy being taken for each player. For mixed strategies, which are probability

VOL. 36, 1950 *MATHEMATICS: G. POLYA* 49

distributions over the pure strategies, the pay-off functions are the expectations of the players, thus becoming polylinear forms in the probabilities with which the various players play their various pure strategies.

Any n-tuple of strategies, one for each player, may be regarded as a point in the product space obtained by multiplying the n strategy spaces of the players. One such n-tuple counters another if the strategy of each player in the countering n-tuple yields the highest obtainable expectation for its player against the $n - 1$ strategies of the other players in the countered n-tuple. A self-countering n-tuple is called an equilibrium point.

The correspondence of each n-tuple with its set of countering n-tuples gives a one-to-many mapping of the product space into itself. From the definition of countering we see that the set of countering points of a point is convex. By using the continuity of the pay-off functions we see that the graph of the mapping is closed. The closedness is equivalent to saying: if $P_1, P_2, \ldots$ and $Q_1, Q_2, \ldots, Q_n, \ldots$ are sequences of points in the product space where $Q_n \to Q$, $P_n \to P$ and Q_n counters P_n then Q counters P.

Since the graph is closed and since the image of each point under the mapping is convex, we infer from Kakutani's theorem[1] that the mapping has a fixed point (i.e., point contained in its image). Hence there is an equilibrium point.

In the two-person zero-sum case the "main theorem"[2] and the existence of an equilibrium point are equivalent. In this case any two equilibrium points lead to the same expectations for the players, but this need not occur in general.

* The author is indebted to Dr. David Gale for suggesting the use of Kakutani's theorem to simplify the proof and to the A. E. C. for financial support.

[1] Kakutani, S., *Duke Math. J.*, **8**, 457–459 (1941).

[2] Von Neumann, J., and Morgenstern, O., *The Theory of Games and Economic Behaviour*, Chap. 3, Princeton University Press, Princeton, 1947.

*REMARK ON WEYL'S NOTE "INEQUALITIES BETWEEN THE TWO KINDS OF EIGENVALUES OF A LINEAR TRANSFORMATION"**

BY GEORGE POLYA

DEPARTMENT OF MATHEMATICS, STANFORD UNIVERSITY

Communicated by H. Weyl, November 25, 1949

In the note quoted above H. Weyl proved a Theorem involving a function $\varphi(\lambda)$ and concerning the eigenvalues α_i of a linear transformation A and those, κ_i, of A^*A. If the κ_i and $\lambda_i = |\alpha_i|^2$ are arranged in descending order,

Sources: J. F. Nash: Equilibrium points in n-person games (1950).

- Requires finite numbers of players and actions, both assumptions are necessary. (Consider 2-player game "who guesses larger number wins".)
- The proof is non-constructive. How to find NE efficiently?

Pareto optimality

Pareto optimality

- A brief detour: another example of an interesting solution concept, other than NE.
- We want to capture "the best" state of a game. Might be difficult, consider the Battle of sexes.
- A strategy profile s in G Pareto dominates s', written $s' \prec s$, if, for every player i, $u_i(s) \geq u_i(s')$, and there exists a player j such that $u_j(s) > u_j(s')$.
  - The relation $\prec$ is a partial ordering of the set S of all strategy profiles of G.
  - The outcomes of G that are considered best are the maximal elements of S in $\prec$.
- A strategy profile $s \in S$ is Pareto optimal if there does not exist another strategy profile $s' \in S$ that Pareto dominates s.
  - In zero-sum games, all strategy profiles are Pareto-optimal.
  - Not all NE are Pareto-optimal (the NE in Prisoner's dilemma)

Vilfredo Pareto

- an Italian engineer, sociologist, economist, political scientist, and philosopher.

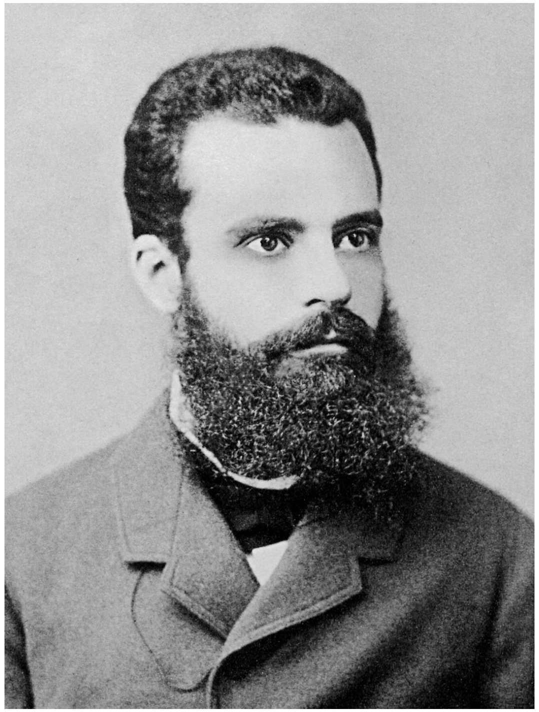

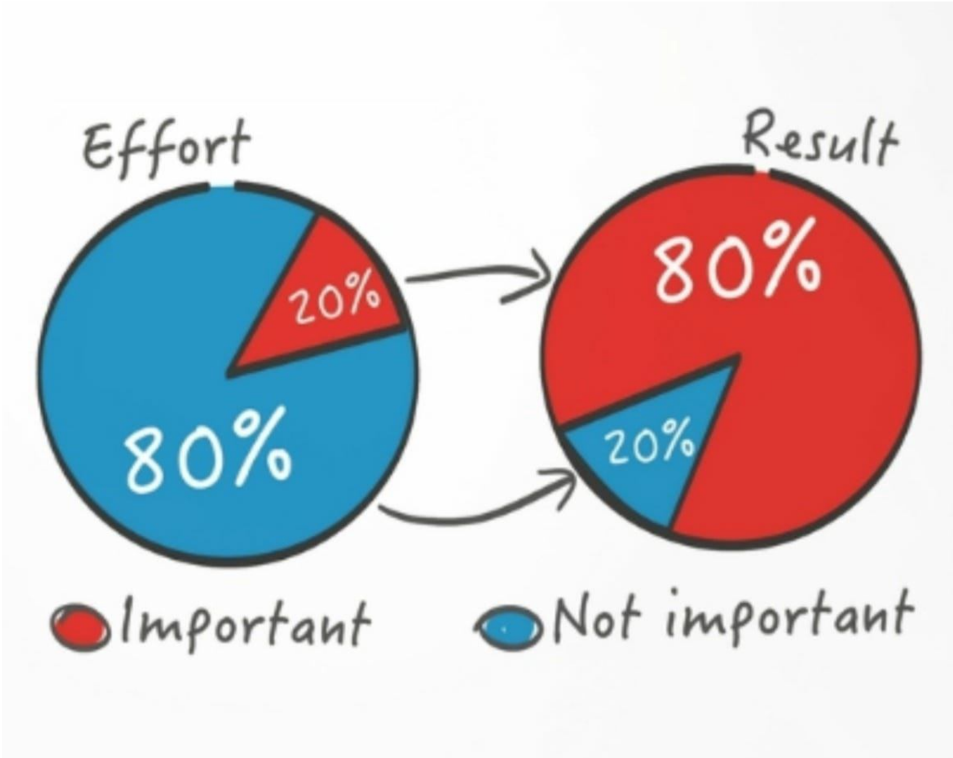

Figure: Vilfredo Pareto (1848–1923).

Sources: https://en.wikipedia.org and https://medium.com/

- Pareto principle: for many outcomes roughly 80% of consequences come from 20% of the causes.

Finding Nash equilibria

- We know that NE exist in every normal-form game (Nash's theorem).
- However, we do not have any algorithm for how to find them yet.
- We start with a simple class of 2-player games, so-called zero-sum games.
- We show that we can find NE efficiently in this case. In fact, we show that NE "solves" zero-sum games completely.
- Historically, zero-sum games were considered first in game game theory (by Morgenstern and Von Neumann in the 1940s).

Zero-sum games

- Two-player games (P, A, u) where $u_1(a) = -u_2(a)$ for every $a \in A$.

	Rock	Paper	Scissors
Rock	(0,0)	(-1,1)	(1,-1)
Paper	(1,-1)	(0,0)	(-1,1)
Scissors	(-1,1)	(1,-1)	(0,0)

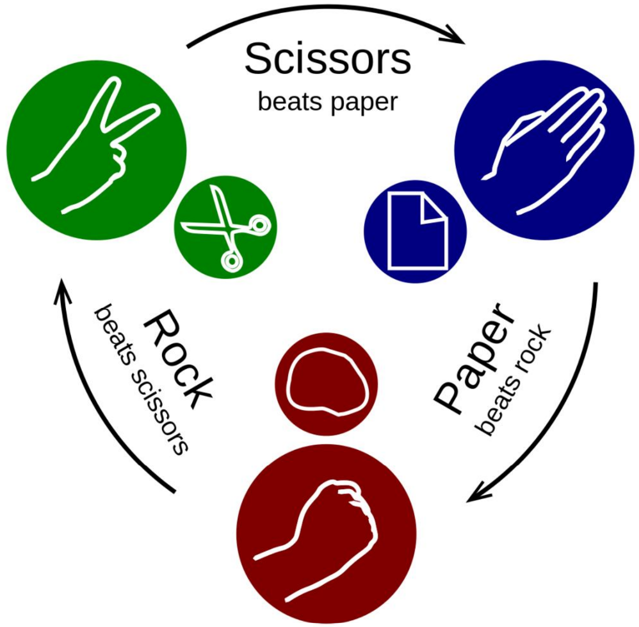

Sources: https://en.wikipedia.org/

Zero-sum games examples: chess

Source: https://edition.cnn.com/

Zero-sum games examples: table tennis

Source: https://www.reddit.com/

Zero-sum games examples: derivative trading

Source: https://www.linkedin.com/

Zero-sum games examples: elections

Source: https://youtube.com/

Zero-sum games examples: many more

Source: https://lhongtortai.com/collection/what-is-a-non-zero-sum-game

Representing zero-sum games

- With zero-sum games, our notation simplifies.
- Let $G = (P, A = A_1 \times A_2, u)$ be a zero-sum game. That is, $u_1(a) + u_2(a) = 0$ for every $a \in A$.
- If $A_1 = \{1, \ldots, m\}$ and $A_2 = \{1, \ldots, n\}$, then G can be represented with an $m \times n$ payoff matrix M where $M_{i,j} = u_1(i,j) = -u_2(i,j)$.
- For a strategy profile (s_1, s_2), we write $x_i = s_1(i)$ and $y_j = s_2(j)$, representing (s_1, s_2) with mixed strategy vectors $x = (x_1, \ldots, x_m)$ and $y = (y_1, \ldots, y_n)$ that satisfy $\sum_{i=1}^{m} x_i = 1$ and $\sum_{j=1}^{n} y_j = 1$.
- The expected payoff of player 1 then equals

$$u_1(s) = \sum_{a=(i,j)\in A} u_1(a) s_1(i) s_2(j) = \sum_{i=1}^{m} \sum_{j=1}^{n} M_{i,j} x_i y_j = x^\top M y = -u_2(s).$$

Worst-case optimal strategies

- Thus, player's 2 best response to a strategy x of 1, is a vector $y \in S_2$ that minimizes $x^\top My$. Player's 1 best response to a strategy y of 2 is $x \in S_1$ that maximizes $x^\top My$.
- Let $\beta(x) = \min_{y \in S_2} x^\top My$ be the best expected payoff of 2 against x. Let $\alpha(y) = \max_{x \in S_1} x^\top My$ be the best expected payoff of 1 to y.
- A strategy profile (x, y) is then a NE if and only if it satisfies $\beta(x) = x^\top My = \alpha(y)$.
- Assume player 1 expects player 2 to select a best response to every strategy x he can come up with. Player 1 then chooses a mixed strategy $\overline{x}$ from S_1 that maximizes his expected payoff under this, rather pessimistic, assumption.
- This worst-case optimal strategy for 1 satisfies $\beta(\overline{x}) = \max_{x \in S_1} \beta(x)$. The worst-case optimal strategy for 2 is a mixed strategy $\overline{y} \in S_2$ that satisfies $\alpha(\overline{y}) = \min_{y \in S_2} \alpha(y)$.

Worst-case optimal strategies and NE

- To achieve NE in a zero-sum game, both players must select their worst-case optimal strategies.

Lemma 2.20

(a) For all $x \in S_1$ and $y \in S_2$, we have $\beta(x) \leq x^\top My \leq \alpha(y)$.

(b) If a strategy profile (x^*, y^*) is NE, then both strategies x^* and y^* are worst-case optimal.

(c) Any strategies $x^* \in S_1$ and $y^* \in S_2$ satisfying $\beta(x^*) = \alpha(y^*)$ form NE (x^*, y^*).

(a) This follows immediately from the definitions of β and α.

(b) Part (a) implies that $\beta(x) \leq \alpha(y^*)$ for every $x \in S_1$. Since (x^*, y^*) is NE, we have $\beta(x^*) = \alpha(y^*)$ and thus $\beta(x) \leq \beta(x^*)$ for every $x \in S_1$. Thus, x^* is a worst-case optimal for 1. Analogously for player 2.

(c) If $\beta(x^*) = \alpha(y^*)$, then (a) implies $\beta(x^*) = (x^*)^\top My^* = \alpha(y^*)$. $\square$

The Minimax Theorem

The Minimax Theorem (Theorem 2.21)

For every zero-sum game, worst-case optimal strategies for both players exist and can be efficiently computed. There is a number v such that, for any worst-case optimal strategies x^* and y^*, the strategy profile (x^*, y^*) is a Nash equilibrium and $\beta(x^*) = (x^*)^\top M y^* = \alpha(y^*) = v$.

Figure: John von Neumann (1903–1957) and Oskar Morgenstern (1902–1977).

Sources: https://en.wikiquote.org and https://austriainusa.org

The Minimax Theorem: remarks

- It was a starting point of game theory.
- Proved by Von Neumann in 1928 (predates Nash's Theorem).
- *"As far as I can see, there could be no theory of games ... without that theorem ... I thought there was nothing worth publishing until the Minimax Theorem was proved."* (Von Neumann).
- The Minimax theorem tells us everything about zero-sum games: there is NE and it can be found efficiently. Moreover, there is a unique value of the game $v = (x^*)^\top M(y^*)$ of the payoff attained in any NE (x^*, y^*).
- There are no secrets in zero-sum games: strategies known in advance change nothing, each player can choose a worst-case optimal strategy and get payoff $\geq v$. If the opponent chooses his worst-case optimal strategy, then his payoff is always $\leq v$.
- The name: the expanded equality $\beta(x^*) = v = \alpha(y^*)$ becomes

$$\max_{x \in S_1} \min_{y \in S_2} x^\top My = v = \min_{y \in S_2} \max_{x \in S_1} x^\top My.$$

- Original proof uses Brouwer's theorem. We will use linear programming.

Source: https://czthomas.files.wordpress.com

Thank you for your attention.