Introduction to Game Theory I

Presented To:

2014 Summer REU

Penn State Department of Mathematics

Presented By:

Christopher Griffin

griffin@math.psu.edu

Types of Game Theory

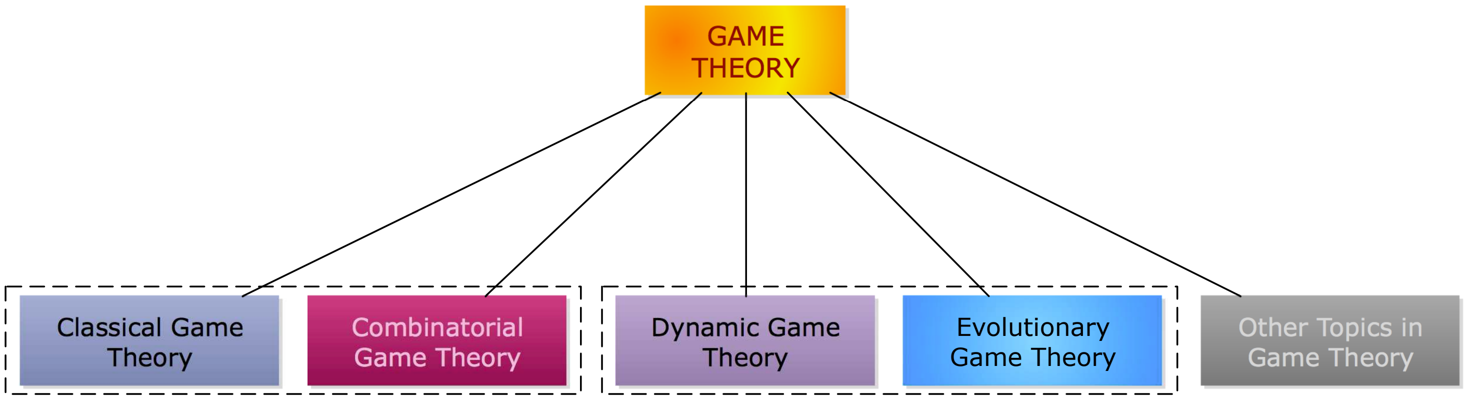

No time, doesn't contain differential equations

Classical Game Theory

Games with finite or infinite strategy space, but no time.

Games with probability (either induced by the player or the game).

Games with coalitions or negotiations.

Examples:
Poker, Strategic Military Decision Making, Negotiations.

Combinatorial Game Theory

Games with **no** chance.

Generally two player strategic games played on boards.

Moves change the structure of a game board.

Examples:
Chess, Checkers, Go, Nim.

Notion of time, may contain differential equations

Dynamic Game Theory

Games with time.

Games with motion or a dynamic component.

Examples:
Optimal play in a dog fight. Chasing your brother across a room.

Evolutionary Game Theory

Modeling population dynamics under competition

Modeling the evolution of strategies in a population.

Examples:
The evolution of behaviors in a group of animals.

Equilibrium of human behavior in social networks.

Other Topics in Game Theory

Experimental / Behavioral Game Theory

Voting Theory

Examples:
Determining why altruism is present in human society.

Determining if there's a fair voting system.

Games Against the House

The *games* we often see on television fall into this category. TV Game Shows (that do not pit players against each other in knowledge tests) often require a single player (who is, in a sense, playing against *The House*) to make a decision that will affect only his life.

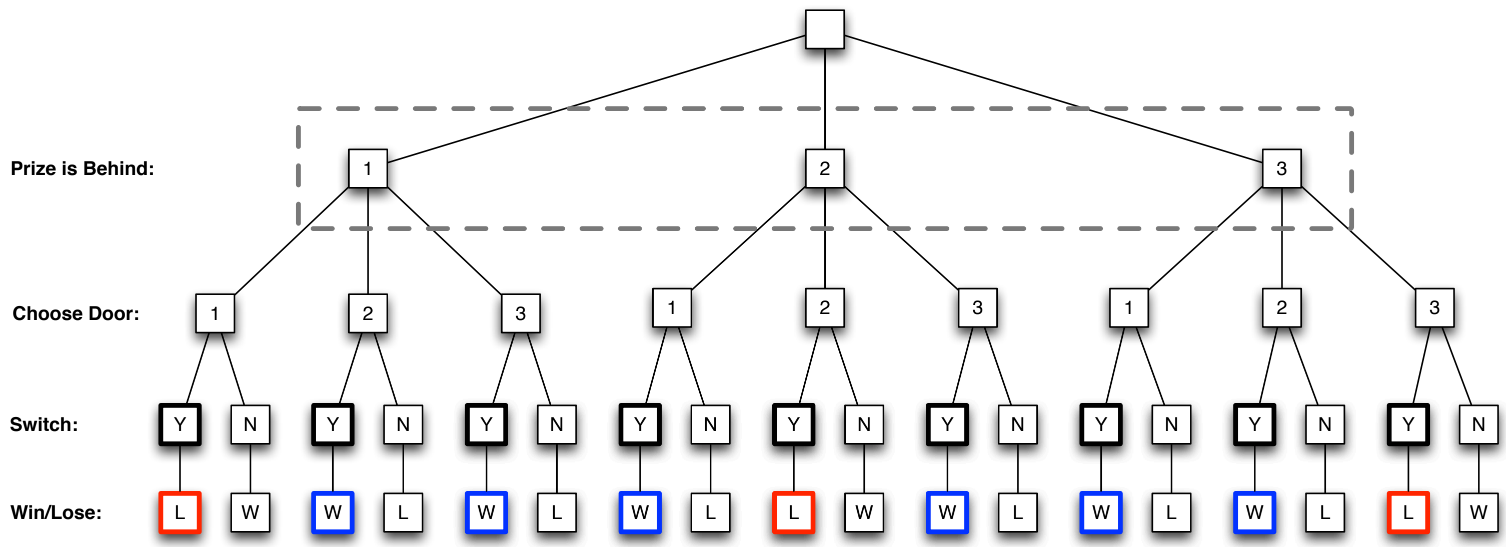

The Monty Hall Problem

Other games against the house:

- Blackjack
- The Price is Right

Assumptions of Multi-Player Games

1. There are a finite set of Players: $\mathbf{P} = \{P_1, \ldots, P_N\}$
2. Each player has a knowledge of the *rules of the game* (the rules under which the game state evolves) and the rules are fixed.
3. At any time $t \in \mathbb{R}_+$ during game play, the player has a finite set of *moves* or choices to make. These choices will affect the evolution of the game. The set of all available moves will be denoted $\mathcal{S}$.
4. The game ends after some finite number of moves.
5. At the end of the game, each player receives a reward or pays a penalty (negative reward).

In addition to these assumptions, some games may incorporate two other components:

1. At certain points, there may be chance moves which advance the game in a non-deterministic way. This only occurs in games of chance. (This occurs, e.g., in poker when the cards are dealt.)
2. In some games the players will know the *entire* history of moves that have been made at all times. (This occurs, e.g., in Tic-Tac-Toe and Chess, but not e.g., in Poker.)

Game Trees (1)

Play(s) in a game can often be described by a tree:

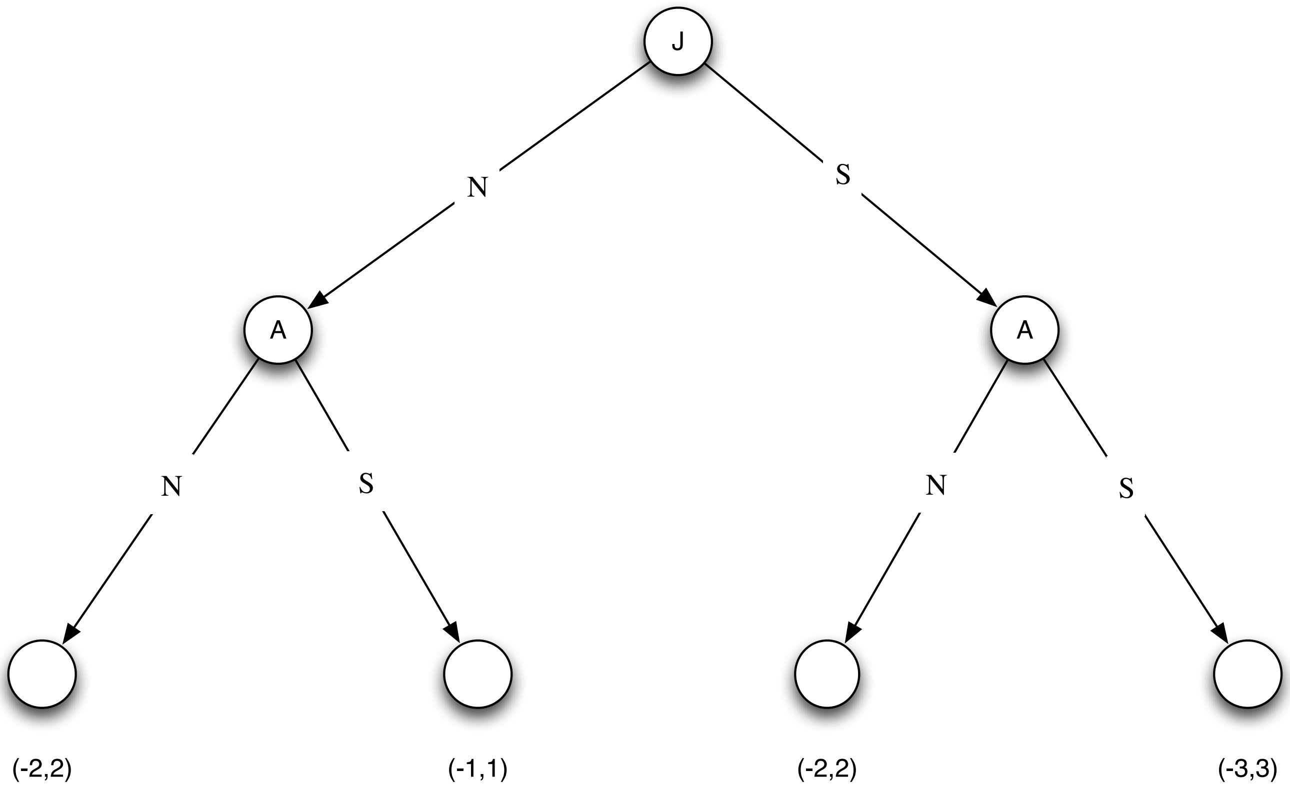

Game Trees (2)

Strategies proscribe actions for players at each vertex of the tree.

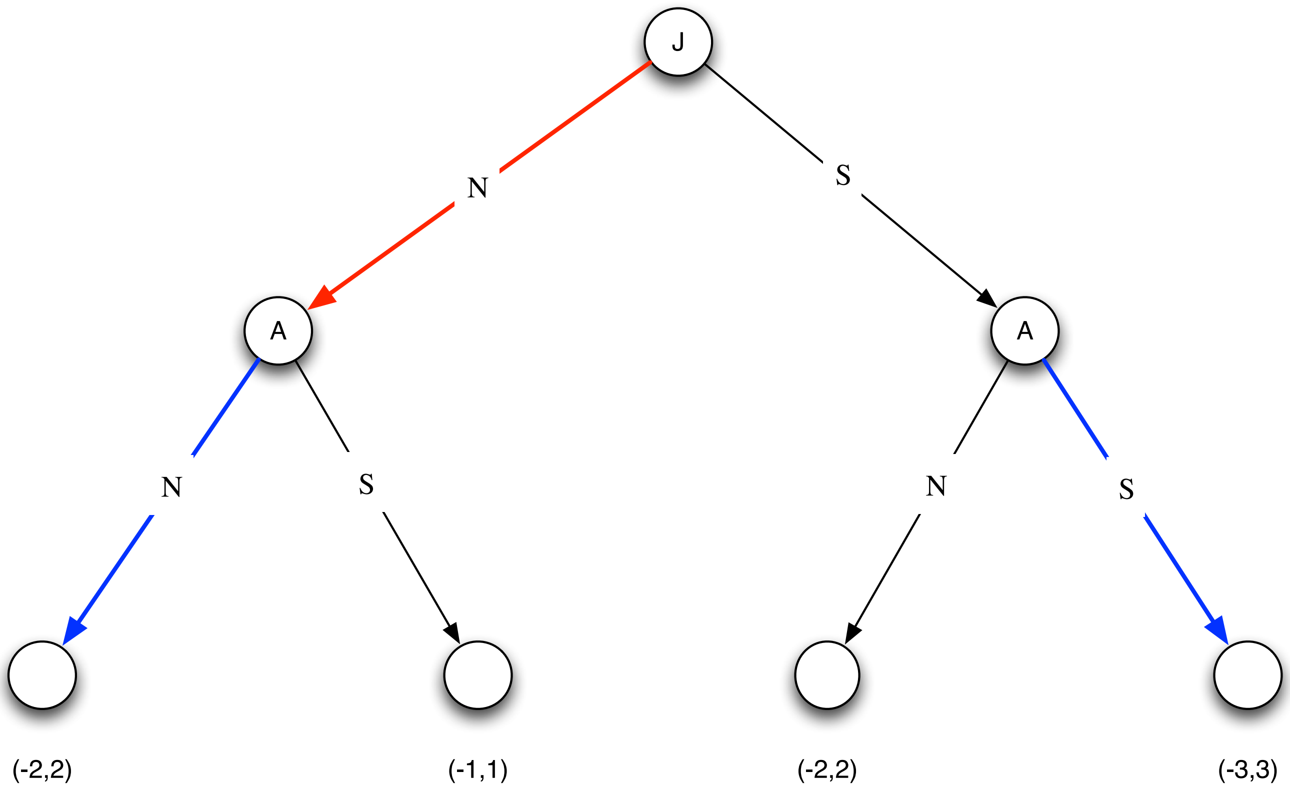

Game Trees (3)

Information sets allow us to limit the information that players know before they make a decision.

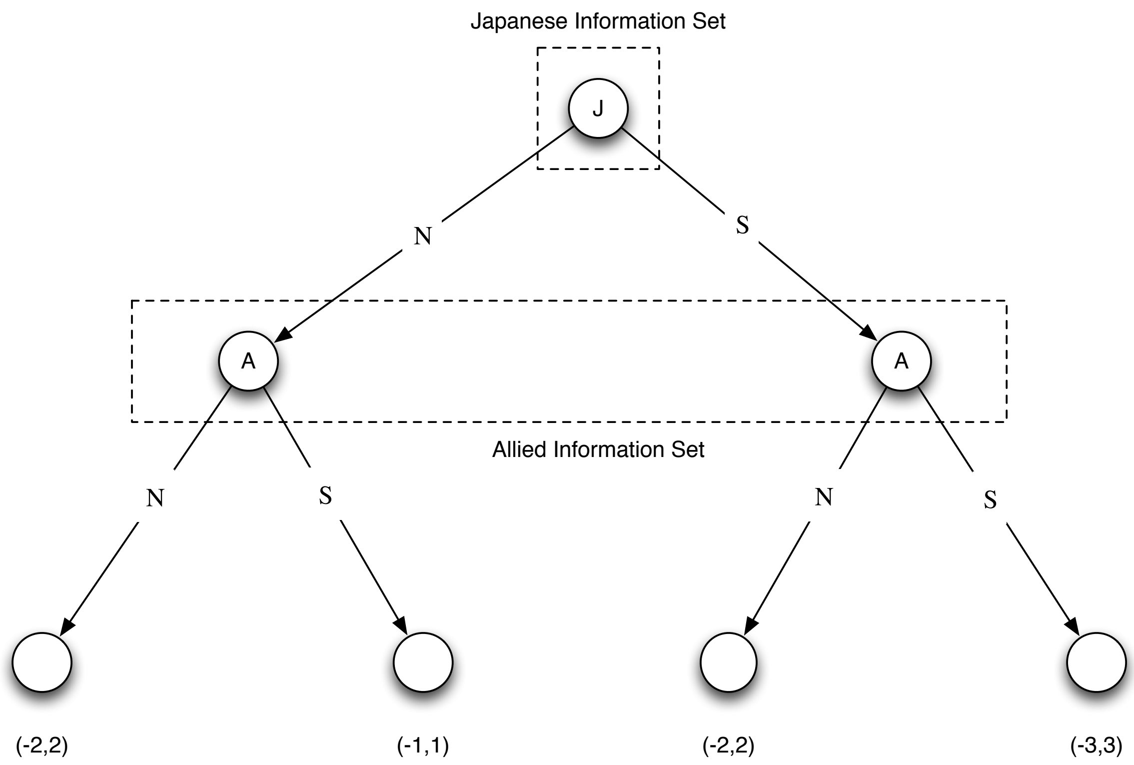

Game Trees (4)

Games of chance can also be modeled in this way:

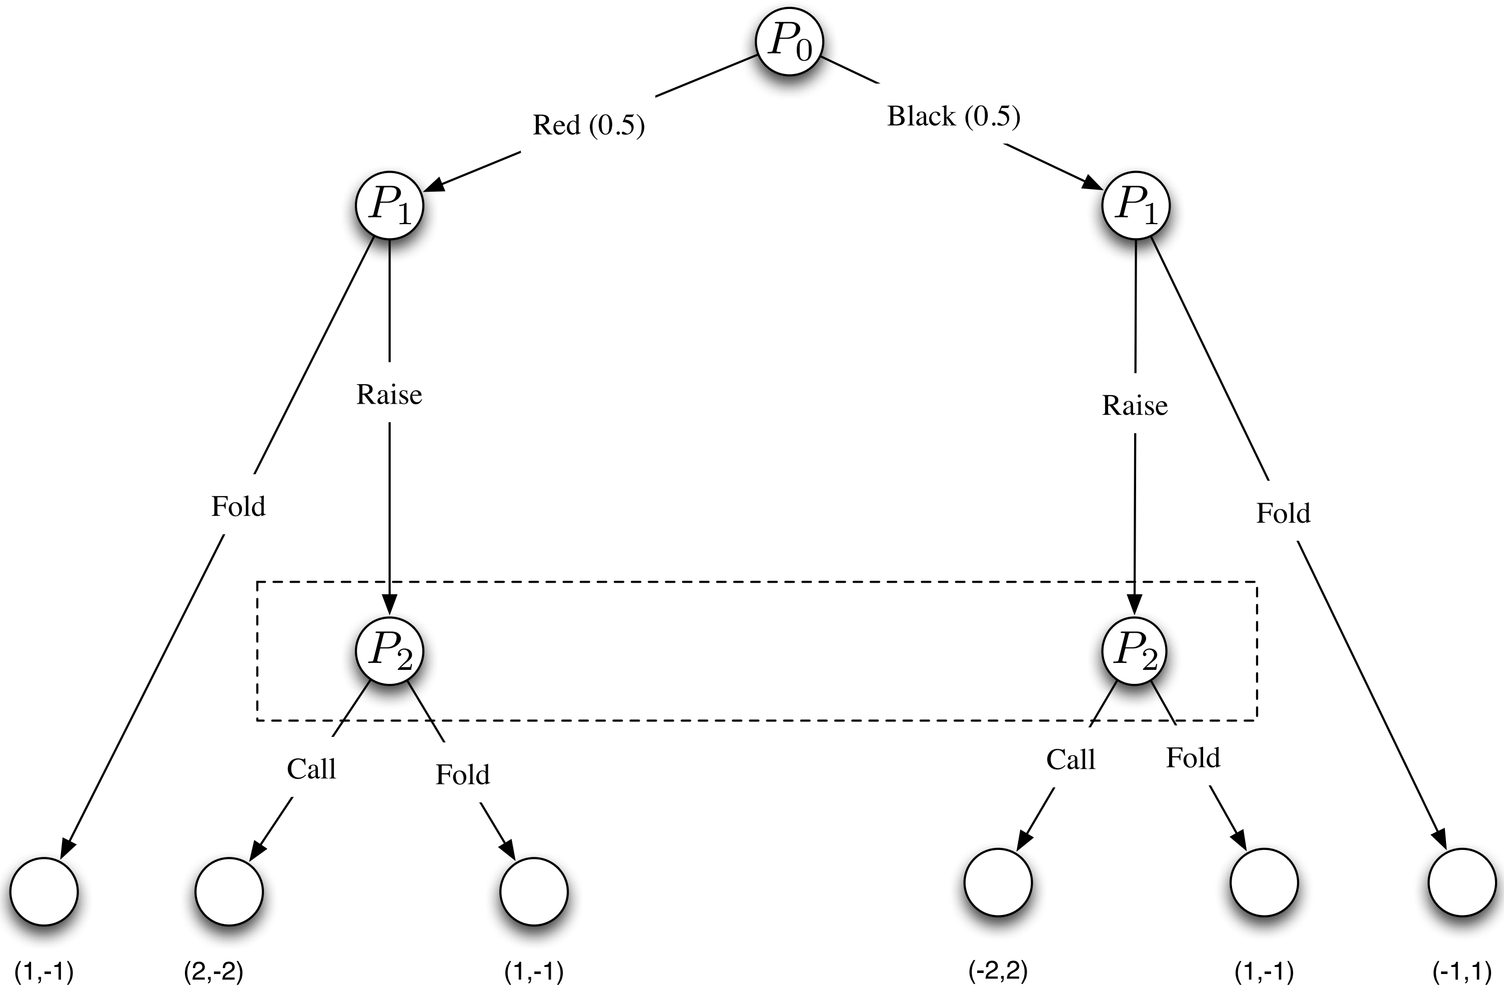

Game Trees (5)

Definition 1 (Game Tree)**.** Let $T = (V, E)$ be a finite directed tree, let $F \subseteq V$ be the terminal vertices and let $D = V \setminus F$ be the intermediate (or decision) vertices. Let $\mathbf{P} = \{P_0, P_1, \ldots, P_n\}$ be a set of players including P_0 the chance player. Let $\mathcal{S}$ be a set of moves for the players. Let $\nu : D \rightarrow \mathbf{P}$ be a player vertex assignment function and $\mu : E \rightarrow \mathcal{S}$ be a move assignment function. Let

$$\mathcal{P} = \{p_v : \nu(v) = P_0 \text{ and } p_v \text{ is the moves of Player 0}\}$$

Let $\pi : F \rightarrow \mathbb{R}^n$ be a payoff function. Let $\mathcal{I} \subseteq 2^D$ be the set of **information sets**.
A *game tree* is a tuple $\mathcal{G} = (T, \mathbf{P}, \mathcal{S}, \nu, \mu, \pi, \mathcal{I}, \mathcal{P})$. In this form, the game defined by the game tree $\mathcal{G}$ is said to be in *extensive* form.

Two Theorems

Theorem 2. *Let $\mathcal{G} = (T, \mathbf{P}, \mathcal{S}, \nu, \mu, \pi, \mathcal{I}, \mathcal{P})$. Let $\sigma_1, \ldots, \sigma_N$ be a collection of strategies for Players 1 through n. Then these strategies determine a discrete probability space $(\Omega, \mathcal{F}, P)$ where Ω is a set of paths leading from the root of the tree to a subset of the terminal nodes and if $\omega \in \Omega$, then $P(\omega)$ is the product of the probabilities of the chance moves defined by the path ω.*

All this theorem says is the combinations of strategies from players defines a set of paths through the game tree from the root to the leafs and each path can be assigned a probability. In the event that there are no probabilistic moves, the path is unique.

Theorem 3. *Let $\mathcal{G} = (T, \mathbf{P}, \mathcal{S}, \nu, \mu, \pi, \mathcal{I})$ be a game with no chance. Let $\sigma_1, \ldots, \sigma_N$ be set of strategies for Players 1 through n. Then these strategies determine a unique path through the game tree.*

Some Definitions

Definition 4 (Strategy Space)**.** Let Σ_i be the set of all strategies for Player i in a game tree $\mathcal{G}$. Then the entire *strategy space* is $\Sigma = \Sigma_1 \times \Sigma_2 \times \cdots \times \Sigma_n$.

Definition 5 (Strategy Payoff Function)**.** Let $\mathcal{G}$ be a game tree with *no chance moves*. The *strategy payoff function* is a mapping $\pi : \Sigma \to \mathbb{R}^n$. If $\sigma_1, \ldots, \sigma_N$ are strategies for Players 1 through n, then $\pi(\sigma_1, \ldots, \sigma_N)$ is the vector of payoffs assigned to the terminal node of the path determined by the strategies $\sigma_1, \ldots, \sigma_N$ in game tree $\mathcal{G}$. For each $i = 1, \ldots, N$ $\pi_i(\sigma_1, \ldots, \sigma_N)$ is the payoff to Player i in $\pi_i(\sigma_1, \ldots, \sigma_N)$.

Consider the Battle of the Bismark Sea game (derived from events in World War II). Then there are four distinct strategies in Σ with the following payoffs:

$$\pi\left(\text{Sail North, Search North}\right) = (-2, 2)$$
$$\pi\left(\text{Sail South, Search North}\right) = (-2, 2)$$
$$\pi\left(\text{Sail North, Search South}\right) = (-1, 1)$$
$$\pi\left(\text{Sail South, Search South}\right) = (-3, 3)$$

Expected Payoff

Definition 6 (Expected Strategy Payoff Function)**.** Let $\mathcal{G}$ be a game tree *with chance moves*. The *expected strategy payoff function* is a mapping $\pi : \Sigma \to \mathbb{R}^n$ defined as follows: If $\sigma_1, \dots, \sigma_N$ are strategies for Players 1 through n, then let $(\Omega, \mathcal{F}, P)$ be the probability space over the paths constructed by these strategies as given in Theorem 2. Let Π_i be a random variable that maps $\omega \in \Omega$ to the payoff for Player i at the terminal node in path ω. Let:

$$\pi_i(\sigma_1, \dots, \sigma_N) = \mathbb{E}(\Pi_i)$$

Then:

$$\pi(\sigma_1, \dots, \sigma_N) = \langle \pi_1(\sigma_1, \dots, \sigma_N), \dots, \pi_N(\sigma_1, \dots, \sigma_N) \rangle$$

As before, $\pi_i(\sigma_1, \dots, \sigma_N)$ is the expected payoff to Player i in $\pi(\sigma_1, \dots, \sigma_N)$.

Further analysis of these general game trees leads to sub-game perfection.

Equilibrium

Definition 7 (Equilibrium)**.** A strategy $(\sigma_1^*, \ldots, \sigma_N^*) \in \Sigma$ is an equilibrium if for all i.

$$\pi_i(\sigma_1^*, \ldots, \sigma_i^*, \ldots, \sigma_N^*) \geq \pi_i(\sigma_1^*, \ldots, \sigma_i, \ldots, \sigma_N^*)$$

where $\sigma_i \in \Sigma_i$.

Consider the Battle of the Bismark Sea. We can show that (Sail North, Search North) is an equilibrium strategy. Recall that:

$$\pi\,(\text{Sail North, Search North}) = (-2, 2)$$

Now, suppose that the Japanese deviate from this strategy and decide to sail south. Then the new payoff is:

$$\pi\,(\text{Sail South, Search North}) = (-2, 2)$$

Thus:

$$\pi_1\,(\text{Sail North, Search North}) \geq \pi_1\,(\text{Sail South, Search North})$$

PENNSTATE

Equilibrium (2)

Now suppose that the Allies deviate from the strategy and decide to search south. Then the new payoff is:

$$\pi\left(\text{Sail North, Search South}\right) = (-1, 1)$$

Thus:

$$\pi_2\left(\text{Sail North, Search North}\right) > \pi_2\left(\text{Sail North, Search South}\right)$$

Theorem 8. *Let $\mathcal{G} = (T, \mathbf{P}, \mathcal{S}, \nu, \mu, \pi, \mathcal{I}, \mathcal{P})$ be a game tree with complete information. Then there is an equilibrium strategy $(\sigma_1^*, \ldots, \sigma_N^*) \in \Sigma$.*

Corollary 9. *This strategy when restricted to any sub-tree is still an equilibrium.*

Such a strategy is called *sub-game perfect*.

Zermelo's Theorem

Corollary 10 (Zermelo's Theorem)**.** *Let $\mathcal{G} = (T, \mathbf{P}, \mathcal{S}, \nu, \mu, \pi)$ be a two-player game with complete information and no chance. Assume that the payoff is such that:*

1. *The only payoffs are $+1$ (win), -1 (lose).*
2. *Player 1 wins $+1$ if and only if Player 2 wins -1.*
3. *Player 2 wins $+1$ if and only if Player 1 wins -1.*

Finally, assume that the players alternate turns. Then one of the two players must have a strategy to obtain $+1$ always.

A variation on this theorem has applications to Chess (Checkers, Go etc.). It tells us that for combinatorial games like Chess, either there is a strategy so that white always wins, or a strategy so that black always wins or ties or a strategy so that black always wins or ties.

Further analysis in this direction leads to Combinatorial Game Theory.

Strategic Form Games

Definition 11 (Normal Form)**.** Let $\mathbf{P}$ be a set of players, $\Sigma = \Sigma_1 \times \Sigma_2 \times \cdots \times \Sigma_N$ be a strategy space and let $\pi : \Sigma \to \mathbb{R}^N$ be a strategy payoff function. Then the triple: $\mathcal{G} = (\mathbf{P}, \Sigma, \pi)$ is a game in *normal form*.

Definition 12 (Strategic Form–2 Player Games)**.** $\mathcal{G} = (\mathbf{P}, \Sigma, \pi)$ be a normal form game with $\mathbf{P} = \{P_1, P_2\}$ and $\Sigma = \Sigma_1 \times \Sigma_2$. If the strategies in Σ_i ($i = 1, 2$) are ordered so that $\Sigma_i = \{\sigma_1^i, \ldots, \sigma_{n_i}^i\}$ ($i = 1, 2$). Then for each player there is a matrix $\mathbf{A}_i \in \mathbb{R}^{n_1 \times n_2}$ so that element (r, c) of $\mathbf{A}_i$ is given by $\pi_i(\sigma_r^1, \sigma_c^2)$. Then the tuple $\mathcal{G} = (\mathbf{P}, \Sigma, \mathbf{A}_1, \mathbf{A}_2)$ is a two-player game in *strategic form*.

Strategic Example

Consider the two-player game defined in the Battle of the Bismark Sea. If we assume that the strategies for the players are:

$$\Sigma_1 = \{\text{Sail North, Sail South}\}$$
$$\Sigma_2 = \{\text{Search North, Search South}\}$$

Then the payoff matrices for the two players are:

$$\mathbf{A}_1 = \begin{bmatrix} -2 & -1 \\ -2 & -3 \end{bmatrix}$$
$$\mathbf{A}_2 = \begin{bmatrix} 2 & 1 \\ 2 & 3 \end{bmatrix}$$

Here, the *rows* represent the different strategies of Player 1 and the *columns* represent the strategies of Player 2.

For historic reasons we usually write $\mathbf{A} = \mathbf{A}_1$ and $\mathbf{B} = \mathbf{A}_2$.

Types of Games

Definition 13 (Symmetric Game)**.** Let $\mathcal{G} = (\mathbf{P}, \Sigma, \mathbf{A}, \mathbf{B})$. If $\mathbf{A} = \mathbf{B}^T$ then $\mathcal{G}$ is called a *symmetric game*.

Definition 14 (Constant / General Sum Game)**.** Let $\mathcal{G} = (\mathbf{P}, \Sigma, \pi)$ be a game in normal form. If there is a constant $C \in \mathbb{R}$ so that for all tuples $(\sigma_1, \ldots, \sigma_N) \in \Sigma$ we have:

$$\sum_{i=1}^{N} \pi_i(\sigma_1, \ldots, \sigma_N) = C \tag{1}$$

then $\mathcal{G}$ is called a *constant sum game*. If $C = 0$, then $\mathcal{G}$ is called a *zero sum game*. Any game that is *not* constant sum is called *general sum*.

The Battle of the Bismark Sea is a zero sum game.

Chicken

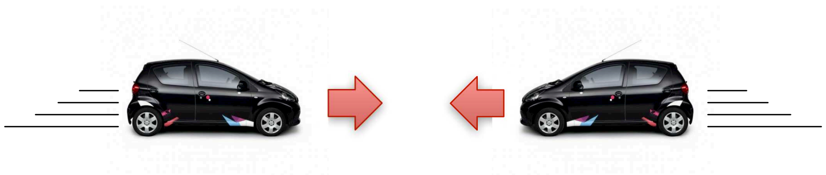

Player 1

	Swerve	Don't Swerve
Swerve	0	-1
Don't Swerve	1	-10

Player 2

	Swerve	Don't Swerve
Swerve	0	1
Don't Swerve	-1	-10

From this we can see the matrices are:

$$\mathbf{A}_1 = \begin{bmatrix} 0 & -1 \\ 1 & -10 \end{bmatrix} \quad \mathbf{A}_2 = \begin{bmatrix} 0 & 1 \\ -1 & -10 \end{bmatrix}$$

Note that the Game of Chicken is **not** a zero-sum game, i.e. it is a general sum game.

PENNSTATE

Relation Back to Game Trees

Proposition 15. *Let* $\mathcal{G} = (\mathbf{P}, \Sigma, \mathbf{A}, \mathbf{B})$ *be a two-player game in strategic form with* $\Sigma_1 = \{\sigma_1^1, \ldots, \sigma_m^1\}$ *and* $\Sigma_2 = \{\sigma_1^2, \ldots, \sigma_n^2\}$*. If Player* P_1 *chooses strategy* σ_r^1 *and Player* P_2 *chooses strategy* σ_c^2*, then:*

$$\pi_1(\sigma_r^1, \sigma_c^2) = \mathbf{e}_r^T \mathbf{A} \mathbf{e}_c \tag{2}$$

$$\pi_2(\sigma_r^1, \sigma_c^2) = \mathbf{e}_r^T \mathbf{B} \mathbf{e}_c \tag{3}$$

Proposition 16 (Equilibrium)**.** *Let* $\mathcal{G} = (\mathbf{P}, \Sigma, \mathbf{A}, \mathbf{B})$ *be a two-player game in strategic form with* $\Sigma = \Sigma_1 \times \Sigma_2$*. The expressions*

$$\mathbf{e}_i^T \mathbf{A} \mathbf{e}_j \geq \mathbf{e}_k^T \mathbf{A} \mathbf{e}_j \quad \forall k \neq i \tag{4}$$

and

$$\mathbf{e}_i^T \mathbf{B} \mathbf{e}_j \geq \mathbf{e}_i^T \mathbf{B} \mathbf{e}_l \quad \forall l \neq j \tag{5}$$

hold if and only if $(\sigma_i^1, \sigma_j^2) \in \Sigma_1 \times \Sigma_2$ *is an equilibrium strategy.*

It is not the case than an equilibrium in pure strategies exists for all games.

Mixed Strategies

Definition 17 (Mixed Strategy Vector). Let $\mathcal{G} = (\mathbf{P}, \Sigma, \pi)$ be a game in normal form with $\mathbf{P} = \{P_1, \dots, P_N\}$. Let $\Sigma_i = \{\sigma_1^i, \dots, \sigma_{n_i}^i\}$. To any mixed strategy for Player P_i we may associate a vector $\mathbf{x}^i = [x_1^i, \dots, x_{n_i}^i]^T$ provided that it satisfies the properties:

1. $x_j^i \geq 0$ for $j = 1, \dots, n_i$
2. $\sum_{j=1}^{n_i} x_j^i = 1$

These two properties ensure we are defining a mathematically correct probability distribution over the strategies set Σ_i.

Definition 18 (Player Mixed Strategy Space). Let $\mathcal{G} = (\mathbf{P}, \Sigma, \pi)$ be a game in normal form with $\mathbf{P} = \{P_1, \dots, P_N\}$. Let $\Sigma_i = \{\sigma_1^i, \dots, \sigma_{n_i}^i\}$. Then the set:

$$\Delta_{n_i} = \left\{ [x_1, \dots, x_{n_i}]^T \in \mathbb{R}^{n \times 1} : \sum_{i=1}^{n_i} x_j = 1; x_j \geq 0, j = 1, \dots, n_i \right\}$$

is the *mixed strategy space* in n_i dimensions for Player P_i.

Strategy Space

There is a pleasant geometry to the space Δ_n (sometimes called a *simplex*). In three dimensions, for example, the space is the face of a tetrahedron. (See Figure 1.)

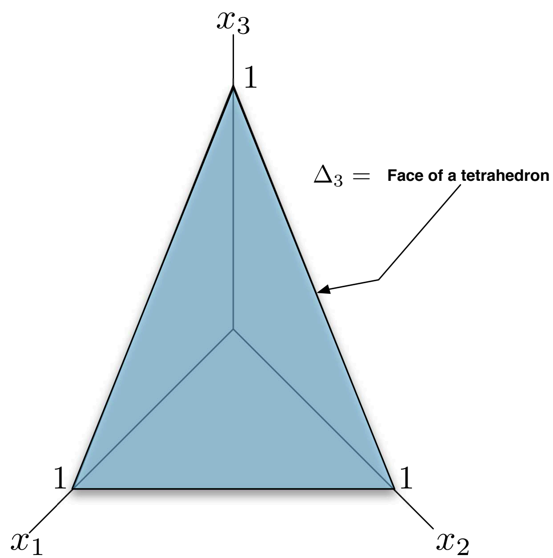

Figure 1: In three dimensional space Δ_3 is the face of a tetrahedron. In four dimensional space, it would be a tetrahedron, which would itself be the face of a four dimensional object.

Mixed Strategies (3)

Definition 19 (Mixed Strategy Payoff Function)**.** Let $\mathcal{G} = (\mathbf{P}, \Sigma, \pi)$ be a game in normal form with $\mathbf{P} = \{P_1, \dots, P_N\}$. Let $\Sigma_i = \{\sigma_1^i, \dots, \sigma_{n_i}^i\}$. The expected payoff can be written in terms of a tuple of mixed strategy vectors $(\mathbf{x}^1, \dots, \mathbf{x}^N)$:

$$u_i(\mathbf{x}^1, \dots, \mathbf{x}^N) = \sum_{i_1=1}^{n_1} \sum_{i_2=1}^{n_2} \cdots \sum_{i_N=1}^{n_N} \pi_i(\sigma_{i_1}^1, \dots, \sigma_{i_N}^n)\mathbf{x}_{i_1}^1 \mathbf{x}_{i_2}^2 \cdots \mathbf{x}_{i_N}^N \quad (6)$$

Here $\mathbf{x}_i^j$ is the i^{th} element of vector $\mathbf{x}^j$. The function $u_i : \Delta \to \mathbb{R}$ defined in Equation 6 is the *mixed strategy payoff function* for Player P_i. (Note: This notation is adapted from Weibull's book on Evolutionary Game Theory.)

Mixed Strategies (2)

Definition 20 (Nash Equilibrium)**.** Let $\mathcal{G} = (\mathbf{P}, \Sigma, \pi)$ be a game in normal form with $\mathbf{P} = \{P_1, \dots, P_N\}$. Let $\Sigma_i = \{\sigma_1^i, \dots, \sigma_{n_i}^i\}$. A *Nash equilibrium* is a tuple of mixed strategies $(\mathbf{x}^{1^*}, \dots, \mathbf{x}^{N^*}) \in \Delta$ so that for all $i = 1, \dots, N$:

$$u_i(\mathbf{x}^{1^*}, \dots, \mathbf{x}^{i^*}, \dots, \mathbf{x}^{N^*}) \geq u_i(\mathbf{x}^{1^*}, \dots, \mathbf{x}^{i}, \dots, \mathbf{x}^{N^*}) \tag{7}$$

for all $\mathbf{x}^i \in \Delta_{n_i}$

Proposition 21. *Let $\mathcal{G} = (\mathbf{P}, \Sigma, \mathbf{A}, \mathbf{B})$ be a two-player matrix game. Let $\Sigma = \Sigma_1 \times \Sigma_2$ where $\Sigma_1 = \{\sigma_1^1, \dots, \sigma_m^1\}$ and $\Sigma_2 = \{\sigma_1^2, \dots, \sigma_n^2\}$. Let $\mathbf{x} \in \Delta_m$ and $\mathbf{y} \in \Delta_n$ be mixed strategies for Players 1 and 2 respectively. Then:*

$$u_1(\mathbf{x}, \mathbf{y}) = \mathbf{x}^T \mathbf{A} \mathbf{y} \tag{8}$$

$$u_2(\mathbf{x}, \mathbf{y}) = \mathbf{x}^T \mathbf{B} \mathbf{y} \tag{9}$$

If you're Player 1 and you don't know $\mathbf{B}$ exactly, but know it can be drawn from a certain probability distribution, this leads to Bayesian games.

Minimax Theorem

Theorem 22. *Let $\mathcal{G} = (\mathbf{P}, \Sigma, \mathbf{A})$ be a zero-sum game with $\mathbf{A} \in \mathbb{R}^{m \times n}$. Then the following are equivalent:*

1. *There is a Nash equilibrium $(\mathbf{x}^*, \mathbf{y}^*)$ for $\mathcal{G}$*
2. *The following equation holds:*

$$v_1 = \max_{\mathbf{x}} \min_{\mathbf{y}} \mathbf{x}^T \mathbf{A} \mathbf{y} = \min_{\mathbf{y}} \max_{\mathbf{x}} \mathbf{x}^T \mathbf{A} \mathbf{y} = v_2 \tag{10}$$

3. *There exists a real number v and $\mathbf{x}^* \in \Delta_m$ and $\mathbf{y}^* \in \Delta_n$ so that:*
   - *(a)* $\sum_i \mathbf{A}_{ij} \mathbf{x}_i^* \geq v$ *for* $j = 1, \ldots, n$ *and*
   - *(b)* $\sum_j \mathbf{A}_{ij} \mathbf{y}_j^* \leq v$ *for* $i = 1, \ldots, m$

Theorem 23 (Minimax Theorem)**.** *Let $\mathcal{G} = (\mathbf{P}, \Sigma, \mathbf{A})$ be a zero-sum game with $\mathbf{A} \in \mathbb{R}^{m \times n}$. Then there is a Nash equilibrium $(\mathbf{x}^*, \mathbf{y}^*)$.*

Nash's Theorem

Definition 24 (Player Best Response)**.** Let $\mathcal{G} = (\mathbf{P}, \Sigma, \pi)$ be an N player game in normal form with $\Sigma_i = \{\sigma_1^i, \ldots, \sigma_{n_i}^i\}$ and let Δ be the mixed strategy space for this game. If $\mathbf{y} \in \Delta$ is a mixed strategy for all players, then the *best reply* for Player P_i is the set:

$$B_i(\mathbf{y}) = \left\{\mathbf{x}^i \in \Delta_{n_i} : u_i(\mathbf{x}^i, \mathbf{y}^{-i}) \geq u_i(\mathbf{z}^i, \mathbf{y}^{-i}) \quad \forall \mathbf{z}^i \in \Delta_{n_i}\right\} \tag{11}$$

Recall $\mathbf{y}^{-i} = (\mathbf{y}^1, \ldots, \mathbf{y}^{i-1}, \mathbf{y}^{i+1}, \ldots, \mathbf{y}^N)$.

Theorem 25. *Let $\mathcal{G} = (\mathbf{P}, \Sigma, \pi)$ be an N player game in normal form with $\Sigma_i = \{\sigma_1^i, \ldots, \sigma_{n_i}^i\}$ and let Δ be the mixed strategy space for this game. The strategy $\mathbf{x}^* \in \Delta$ is a Nash equilibrium for $\mathcal{G}$ if and only if $\mathbf{x}^* \in B(\mathbf{x}^*)$.*

Theorem 26 (Existence of Nash Equilibria)**.** *Let $\mathcal{G} = (\mathbf{P}, \Sigma, \pi)$ be an N player game in normal form. Then $\mathcal{G}$ has at least one Nash equilibrium.*

The original proof was based on Kakutani's Fixed Point Theorem, which is not satisfying.

Finding Nash Equilibria

Nash modified and extended his proof to use Brouwer's Fixed Point Theorem by defining:

$$J_k^i(\mathbf{x}) = \max\left\{0, u_i(\mathbf{e}_k, \mathbf{x}^{-i}) - u_i(\mathbf{x}^i, \mathbf{x}^{-i})\right\} \tag{12}$$

We can now define:

$$\mathbf{x}_j^{i\,\prime} = \frac{\mathbf{x}_j^i + J_j^i(\mathbf{x})}{1 + \sum_{k=1}^{n_i} J_k^i(\mathbf{x})} \tag{13}$$

Using this equation, we can construct a mapping $T : \Delta \to \Delta$ and show that every fixed point of T is a Nash Equilibrium. Using the Brouwer fixed point theorem, it then follows that a Nash equilibrium exists. Unfortunately, this is still not a very useful way to construct a Nash equilibrium.

Equation 13 can help lead to Evolutionary Game Theory.

Optimization for Games

Consider a game in normal form $\mathcal{G} = (\mathbf{P}, \Sigma, \pi)$. We'll assume that $\mathbf{P} = \{P_1, \dots, P_N$ and $\Sigma_i = \{\sigma_1^i, \dots, \sigma_{n_i}^i\}$. If we assume a fixed mixed strategy $\mathbf{x} \in \Delta$, Player P_i's objective when choosing a response $\mathbf{x}^i \in \Delta_{n_i}$ is to solve the following problem:

$$\text{Player } P_i : \begin{cases} \max & u_i(\mathbf{x}^i, \mathbf{x}^{-i}) \\ s.t. & \mathbf{x}_1^i + \cdots + \mathbf{x}_{n_i}^i = 1 \\ & \mathbf{x}_j^i \geq 0 \quad j = 1, \dots, n_i \end{cases} \tag{14}$$

The interesting part (and the part that makes Game Theory hard) is that each player is solving this problem *simultaneously*. Thus an equilibrium solution is a simultaneous solution to:

$$\forall i : \begin{cases} \max & u_i(\mathbf{x}^i, \mathbf{x}^{-i}) \\ s.t. & \mathbf{x}_1^i + \cdots + \mathbf{x}_{n_i}^i = 1 \\ & \mathbf{x}_j^i \geq 0 \quad j = 1, \dots, n_i \end{cases} \tag{15}$$

This leads to an incredibly rich class of problems in mathematical programming.

Linear Programming

Theorem 27. *Let* $\mathcal{G} = (\mathbf{P}, \Sigma, \mathbf{A})$ *be a zero-sum two player game with* $\mathbf{A} \in \mathbb{R}^{m \times n}$*. Then a Nash equilibrium solution for Player 1 is an optimal solution to:*

$$
\begin{aligned}
\max \quad & v \\
s.t. \quad & \mathbf{A}_{11}x_1 + \cdots + \mathbf{A}_{m1}x_m - v \geq 0 \\
& \mathbf{A}_{12}x_1 + \cdots + \mathbf{A}_{m2}x_m - v \geq 0 \\
& \vdots \\
& \mathbf{A}_{1n}x_1 + \cdots + \mathbf{A}_{mn}x_m - v \geq 0 \\
& x_1 + \cdots + x_m - 1 = 0 \\
& x_i \geq 0 \quad i = 1, \dots, m
\end{aligned}
$$

Linear Programming (2)

Theorem 28. *Let* $\mathcal{G} = (\mathbf{P}, \Sigma, \mathbf{A})$ *be a zero-sum two player game with* $\mathbf{A} \in \mathbb{R}^{m \times n}$*. Then a Nash equilibrium solution for for Player 2 is an optimal solution to:*

$$
\begin{aligned}
\min\ & \nu \\
s.t.\ & \mathbf{A}_{11}y_1 + \cdots + \mathbf{A}_{1n}y_n - \nu \leq 0 \\
& \mathbf{A}_{21}y_1 + \cdots + \mathbf{A}_{2n}y_n - \nu \leq 0 \\
& \vdots \\
& \mathbf{A}_{m1}y_1 + \cdots + \mathbf{A}_{mn}y_n - \nu \leq 0 \\
& y_1 + \cdots + y_n - 1 = 0 \\
& y_i \geq 0 \quad i = 1, \dots, m
\end{aligned}
$$

Quadratic Programming

Theorem 29. *Let $\mathcal{G} = (\mathbf{P}, \Sigma, \mathbf{A}, \mathbf{B})$ be a general sum two-player matrix game with $\mathbf{A}, \mathbf{B} \in \mathbb{R}^{m \times n}$. A point $(x^*, y^*) \in \Delta$ is a Nash equilibrium if and only if there are reals α^* and β^* so that $(\mathbf{x}^*, \mathbf{y}^*, \alpha^*, \beta^*)$, is a global maximizer for the quadratic programming problem:*

$$\begin{aligned}
\max \quad & \mathbf{x}^T(\mathbf{A} + \mathbf{B})\mathbf{y} - \alpha - \beta \\
s.t. \quad & \mathbf{A}\mathbf{y} - \alpha \mathbf{1}_m \leq \mathbf{0} \\
& \mathbf{x}^T\mathbf{B} - \beta \mathbf{1}_n^T \leq \mathbf{0} \\
& \mathbf{1}_m^T\mathbf{x} - 1 = 0 \\
& \mathbf{1}_n^T\mathbf{y} - 1 = 0 \\
& \mathbf{x} \geq \mathbf{0} \\
& \mathbf{y} \geq \mathbf{0}
\end{aligned} \tag{16}$$

Example

We can find a third Nash equilibrium for the Chicken game using this approach. Recall we have:

$$\mathbf{A} = \begin{bmatrix} 0 & -1 \\ 1 & -10 \end{bmatrix} \quad \mathbf{B} = \begin{bmatrix} 0 & 1 \\ -1 & -10 \end{bmatrix}$$

This yields the QP:

$$\begin{cases} \max & -20x_2y_2 - \alpha - \beta \\ s.t. & -y_2 - \alpha \leq 0 \\ & y_1 - 10y_2 - \alpha \leq 0 \\ & -x_2 - \beta \leq 0 \\ & x_1 - 10x_2 - \beta \leq 0 \\ & x_1 + x_2 = 1 \\ & y_1 + y_2 = 1 \\ & x_1, x_2, y_1, y_2 \geq 0 \end{cases} \tag{17}$$

An optimal solution to this problem is $x_1 = 0.9$, $x_2 = 0.1$, $y_1 = 0.9$, $y_2 = 0.1$. This is a third Nash equilibrium in mixed strategies for this instance of Chicken.

PENNSTATE

Linear Complementarity

It turns out, we can generalize this entire framework into something called a *Linear Complementarity Problem* in which we try and find vectors $\mathbf{w}$ and $\mathbf{z}$ so that for some matrix $\mathbf{M}$:

$$\begin{cases} \mathbf{w} - \mathbf{M}\mathbf{z} = -\mathbf{1} \\ \mathbf{w}^T\mathbf{z} = 0 \\ \mathbf{w}, \mathbf{z} \geq 0 \end{cases} \tag{18}$$

Here the matrix $\mathbf{M}$ is defined by the $\mathbf{A}$ and $\mathbf{B}$ matrices and the $\mathbf{w}$ and $\mathbf{z}$ vectors can be used to extract strategy vectors $\mathbf{x}^*$ and $\mathbf{y}^*$.
Lemke and Howson proved this result in 1964 and they also proved:

Theorem 30. *Let $\mathcal{G} = (\mathbf{P}, \Sigma, \mathbf{A}, \mathbf{B})$ be a general sum two-player matrix game. If the game is non-degenerate, then there are an odd number of Nash equilibria.*

This theorem was generalized by Wilson in 1971. “Well behaved” games have an odd number of Nash equilibria. The study of the computational complexity of finding Nash equilibria starts here.

Trembling Hand Perfection

Consider a game $\mathcal{G} = (\mathbf{P}, \Sigma, \pi)$ with $\Sigma = \Sigma_1 \times \cdots \times \Sigma_N$. To each pure Player i strategy σ_j^i assign a (small) value $\mu_{ij} > 0$ so that in the corresponding mixed strategy space we require $x_j^i \geq \mu_{ij}$. If $\boldsymbol{\mu}_i$ is the vector of these μ_{ij} for Player i, then we may define:

$$\Delta_{n_i}(\boldsymbol{\mu}_i) = \left\{ [x_1, \dots, x_{n_i}]^T \in \mathbb{R}^{n \times 1} : \sum_{j=1}^{n_i} x_j = 1; x_j \geq \mu_{ij}, j = 1, \dots, n_i \right\} \tag{19}$$

Define the game $\mathcal{G}(\boldsymbol{\mu})$ to be $\mathcal{G}$ were we require all mixed strategies to be chosen from $\Delta(\boldsymbol{\mu})$.

Definition 31. If $\mathbf{x}^{\boldsymbol{\mu}}$ is a Nash equilibrium in $G(\boldsymbol{\mu})$ and $\mathbf{x}^*$ is a Nash equilibrium for $\mathcal{G}$ and:

$$\lim_{\boldsymbol{\mu} \to \mathbf{0}} \mathbf{x}^{\boldsymbol{\mu}} \to \mathbf{x}^* \tag{20}$$

then $\mathbf{x}^*$ is a *trembling hand perfect equilibrium*.