Constant Sum Games

Paper - in lieu of final exam; choose one of three types:
- Original topic applying game theory
- Critique paper that uses game theory
- Critical survey of some area we don't cover in class (review basic literature; explain why they did what they did and determine if it's sensible)

Focus of Course - application of game theory (not theoretical mathematical foundations); learn enough to read articles that apply game theory and to use it in our own papers

Review

Describing Game - 2 ways
- **Strategic (Normal) Form** - lists all strategies and puts payoffs in table
- **Extensive Form** - draw game three

Strategy - instructions for what player would do in <u>every</u> contingency of the game (not just the equilibrium path); point is, no matter what happens in the game, an agent with a strategy can play a response for anything the opponent does
- **Problem** - strategies can get very complicated in multi-stage game; e.g., chess is a two player, constant sum game with perfect information; equilibrium should be "easy", but strategy space is very complex

Preferences - course will assume Von Neumann/Morgenstern (VNM) utility functions (unique up to a positive affine transformation: $au(\mathbf{x})+b \quad (a>0)$

Continuous Strategy Space - can use to approximate discrete game with many strategies (and vice versa); sometimes easier to find mixed strategy equilibrium

Constant Sum Games

Two Person Constant Sum Games - solved and set benchmark for quality of solution/equilibria for other classes of games

Constant Sum - sum of payoffs in strategic form is constant
- **Utilities** - constant sum refers to payoffs in utility terms; there's no discussion of what the utility functions must look like for this to be possible; the only way the game will also be constant sum with respect to monetary payoffs is if the players are risk neutral
- **Zero Sum** - sum of payoffs equals zero; same as constant sum because of VNM
  - <u>Proof</u>: assume constant sum so $u^1(\mathbf{x})+u^2(\mathbf{x})=k\neq 0$
    - Can transform utility functions because of VNM: $a_1u^1(\mathbf{x})+b_1+a_2u^2(\mathbf{x})+b_2=0$
    - From original equation sub $u^2(\mathbf{x})=k-u^1(\mathbf{x})$: $a_1u^1(\mathbf{x})+a_2(k-u^1(\mathbf{x}))+b_1+b_2=0$
    - Can set $a_1=a_2=1$: $k+b_1+b_2=0$
    - $\therefore$ if we use $b_1+b_2=k$, we can convert a constant sum game into a zero sum game
- **Flexibility** - can ensure all payoffs are positive (as long as no payoff is -∞)... benefit for mathematical techniques

Strictly Competitive - used synonymously with constant sum by Luce and Raiffa, but others use it as a more general form: players' ordinal preferences over pure strategy outcomes are opposites

Example - if player 1 prefers outcome 1 (O^1) to outcome 2 (O^2) then player 2 prefers outcome 2 to outcome 1; in notation:

$O^1 \; \mathrm{P}_1 \; O^2 \;\Rightarrow\; O^2 \; \mathrm{P}_2 \; O^1$ or $u^1(O^1) > u^1(O^2) \;\Rightarrow\; u^2(O^2) > u^2(O^1)$

Mixed Strategies - with new definition, strictly competitive is not the same as constant sum because in strictly competitive players can agree on preferences for a lottery (mixed strategy) which they can't do in a constant sum game with opposite ordinal preferences over pure strategy outcomes

Proof: Assume $u^1(O^1) > u^1(O^2) > u^1(O^3)$ (so $u^2(O^3) > u^2(O^2) > u^2(O^1)$)

Pick lottery of O^1 and O^3 where O^1 occurs with probability p (so O^3 occurs with probability $1-p$)

From VNM players value lottery by expected payoff so for both players to prefer the lottery we must have:

(Player 1) $pu^1(O^1) + (1-p)u^1(O^3) > u^1(O^2)$

(Player 2) $pu^2(O^1) + (1-p)u^2(O^3) > u^2(O^2)$

For player 1 this simplifies to $p > \dfrac{u^1(O^2) - u^1(O^3)}{u^1(O^1) - u^1(O^3)}$

Now we can use a VNM transformation:

$v^1(O^3) = au^1(O^3) + b = 0$

$v^1(O^1) = au^1(O^3) + b = 1$

(2 equations with 2 unknowns: a and b and results in $0 < v^1(O^2) < 1$)

Now we can say $p > \dfrac{v^1(O^2) - 0}{1-0}$ or simply $p > v^1(O^2)$

For player 2 to prefer the lottery we need $p > \dfrac{u^2(O^2) - u^2(O^3)}{u^2(O^1) - u^2(O^3)}$

Use a similar VNM transformation for player 2 (i.e., $v^2(O^3) = 1$ and $v^2(O^1) = 0$)

$$p > \frac{v^2(O^2) - 1}{0-1} \;\Rightarrow\; 1-p > v^2(O^2)$$

It is feasible that these two conditions could hold and the players can both prefer the lottery over outcome 2 for certain

For a constant sum game (with the same preferences), however, this is not possible. After the transformation, the constant sum payoff is 1 (just pick outcome 1 or 3 and this should be obvious). Therefore, we must have $v^1(O^2) + v^2(O^2) = 1$. We just showed that for both players to prefer the lottery we must have (a) $p > v^1(O^2)$ and (b) $1-p > v^2(O^2)$. If we add these together, we get $1 > v^1(O^2) + v^2(O^2)$ which is not possible in a constant sum game.

Why Strictly Competitive - so constant sum and strictly competitive (with the new definition) are slightly different, but why do we care? Many results for strictly competitive games are the same as constant sum games, but they differ in mixed strategies so some game theorists like to distinguish the two

Real World - constant sum games are hard to find; strictly competitive is more realistic
Poker - constant sum in terms of money, but not necessarily in utility (unless all players are risk neutral); it will, however, always be strictly competitive over pure strategies
Deviations - poker books are usually written for risk neutral players; players following "the book" can take advantage of players who are risk averse, risk seeking, or simply make a mistake; the bad part is once a player deviates from the equilibrium path (say to take advantage of another player's risk aversion or mistake), the other player can then take advantage of the one who deviates to take advantage of the original deviation
Why Play - if all players are risk neutral and play by "the book", all players have expected winnings of zero; the fact that people play indicates that not everyone is risk neutral (or people think they can decipher another player's level of risk aversion and take advantage of deviations from equilibrium play from "the book")

Matrix Game - implies a 2 person (rows and columns), constant sum (1 number per cell which corresponds to the row player's utility), finite game
Bimatrix Game - 2 person, nonconstant sum game (lists payoffs for both players in each cell)

Equilibrium - 2 types in constant sum games

Nash Equilibrium - defined as a pair of strategies $(\mathbf{x}_i^*, \mathbf{y}_j^*)$ such that utility to the row (column) player is at least as much as the utility to the row (column) player for all other strategies:

$$\boxed{\begin{aligned} u^R(\mathbf{x}_i^*, \mathbf{y}_j^*) &\geq u^R(\mathbf{x}_i, \mathbf{y}_j^*) \ \forall \ \mathbf{x}_i \\ u^R(\mathbf{x}_i^*, \mathbf{y}_j^*) &\leq u^R(\mathbf{x}_i^*, \mathbf{y}_j) \ \forall \ \mathbf{y}_j \end{aligned}}$$

(this is based on row player's utility)

Alternative - in each column $\mathbf{x}_i^*$ gives maximum payoff (i.e., it's the row player's best reply to that strategy played by the column player); in each row $\mathbf{y}_j^*$ is the minimum payoff (i.e., it's the column player's best reply to that strategy played by the row player)
Saddle Point - the equilibrium point describes a saddle point; maximizing in one direction and minimizing in another

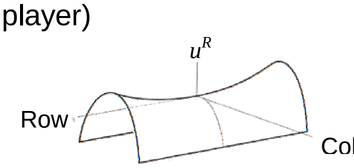

Cournot - Nash equilibrium is similar to Cournot because it assumes simultaneous play

Guarantee Levels - VNM equilibrium
Intuition - don't know what the other guy is going to do so look at each strategy available and determine what is the worst that could happen (based on what opponent does); then pick the strategy with the best guarantee
Math - guarantee for row strategy *i*: $G^R(\mathbf{x}_i) \equiv \min_{\mathbf{y}_j} u^R(\mathbf{x}_i, \mathbf{y}_j)$

This is the worst that can happen to the row player when he/she plays strategy *i*; minimize payoffs across the columns for row *i*

Guarantee for column strategy *j*: $G^C(\mathbf{y}_j) \equiv \max_{\mathbf{x}_i} u^R(\mathbf{x}_i, \mathbf{y}_j)$

"Best" strategies (based on guarantee levels) are determined by:

$$\boxed{G^R \equiv \max_{\mathbf{x}_i} \min_{\mathbf{y}_j} u^R(\mathbf{x}_i, \mathbf{y}_j) \ \ \& \ \ G^C \equiv \min_{\mathbf{y}_j} \max_{\mathbf{x}_i} u^R(\mathbf{x}_i, \mathbf{y}_j)}$$

Pure Strategy Equilibrium - if $G^R = G^C$ there is a pure strategy equilibrium (and a common guarantee for both players)

Mixed Strategy - VNM proved if $G^R \neq G^C$ for pure strategies, there will be a mixed strategy where $G^R = G^C$

Mixed Extension - same basic idea except use $u^R(\mathbf{s}^R, \mathbf{s}^C)$, where the arguments are probability vectors (or they could be strategy vectors for continuous games)

Not Nash - not taking opponent's strategy as fixed

Stakelberg - playing against opponent's best reply and choosing your best reply as if it were a sequential game (like Stakelberg)... this only applies in constant sum games

Paranoia - it makes sense to assume opponent wants to hurt you in constant sum game because the only way your opponent does better is by making you do worse; Slutsky: "In constant sum games, paranoia is rational"

Equivalence - for matrix game, Nash equilibrium and guarantee level are the same

Timing - doesn't matter for matrix game (that's why the two equilibria are the same)

Real World - for nonconstant (bimatrix) games, Nash and guarantee level are not necessarily the same thing which raises a **problem**: which do we use?

Nash vs. Guarantee - we usually use Nash equilibrium, but there are cases when guarantee levels look more appealing

Example - assume nonconstant sum game, but only look at payoffs for row player; also assume $(\mathbf{x}_1, \mathbf{y}_2)$ is the Nash equilibrium; in real world, row player would probably pick $\mathbf{x}_2$ rather than $\mathbf{x}_1$

	$\mathbf{y}_1$	$\mathbf{y}_2$	$\mathbf{y}_3$	$\mathbf{y}_4$
$\mathbf{x}_1$	0	10	0	0
$\mathbf{x}_2$	9.99	9.99	9.99	9.99
$\mathbf{x}_3$				

Counter - given VNM preferences, we can change payoffs so it doesn't look so bad; don't be misled by looking at small differences

	$\mathbf{y}_1$	$\mathbf{y}_2$	$\mathbf{y}_3$	$\mathbf{y}_4$
$\mathbf{x}_1$	0	10	0	0
$\mathbf{x}_2$	0.01	0.01	0.01	0.01
$\mathbf{x}_3$				

Multiple Equilibria - another problem occurs when there are multiple equilibrium (Coordination Problem)

Refinements - try to eliminate those equilibria that are "unreasonable" but usually don't get whittle the list down to one (some of the "unreasonable" ones are better than what's left)

Constant Sum Games - this problem doesn't occur in matrix games; any equilibrium strategy of the row player combined with any equilibrium strategy for the column player also defines an equilibrium and all equilibria are equivalent

	$\mathbf{y}_1$	$\mathbf{y}_2$	$\mathbf{y}_3$	$\mathbf{y}_4$
$\mathbf{x}_1$				
$\mathbf{x}_2$				
$\mathbf{x}_3$				

Example - if $(\mathbf{x}_1, \mathbf{y}_1)$ and $(\mathbf{x}_3, \mathbf{y}_3)$ are equilibria, then $(\mathbf{x}_1, \mathbf{y}_3)$ and $(\mathbf{x}_3, \mathbf{y}_1)$ are too

Equilibrium vs. Efficiency - prisoners' dilemma

Weakly Dominated - prisoners' dilemma uses strictly dominated strategies, but even with weakly dominated strategies, refinements may remove a Pareto optimal equilibrium which doesn't make sense

Example - (10, 10) and (2, 2) are Nash equilibria, but (2, 2) would be the only one if refinements eliminate weakly dominated strategies

	$\mathbf{y}_1$	$\mathbf{y}_2$
$\mathbf{x}_1$	0, 10	10, 10
$\mathbf{x}_2$	2, 2	10, 0

Constant Sum Games - issue of efficiency is non existent; sum of

payoffs is always the same so all results are Pareto optimal
Good - don't have to worry about it
Bad - not realistic; can't use constant s um game to analyze real world situations

Details - closer look at some of the statements made above

Mixed Strategy - VNM proved if $G^R \neq G^C$ for pure strategies, there will be a mixed strategy where $G^R = G^C$

Proof: VNM used a fixed point theorem; later proven with separating hyperplane theorem; Owen proved it with induction argument (no math, but hardest proof to understand)
We'll sketch the second version... needs lots of support material it's long (goes to p.8!)

Notation - matrix game can be represented by the payoff matrix $\mathbf{A}^{m \times n}$

Mixed Strategies - use probability vectors:

$$\mathbf{x} \in R^m \text{ with } \mathbf{x} \geq \mathbf{0} \text{ and } \sum_{j=1}^{m} x_j = 1 \text{ (row vector)}$$

$$\mathbf{y} \in R^n \text{ with } \mathbf{y} \geq \mathbf{0} \text{ and } \sum_{k=1}^{n} y_k = 1 \text{ (column vector)}$$

Semipositive Vector - vector with some elements = 0 and others > 0; more general than a probability vector, but all probability vectors are semipositive

Expected Payoff - to row player is $\mathbf{xAy}$ $(1\text{x}m)(m\text{x}n)(n\text{x}1) = (1\text{x}1)$

Guarantee Levels - we could rewrite $u^R(\mathbf{x},\mathbf{y}) \equiv \mathbf{xAy}$ which allows us to rewrite the guarantee levels: $G^R \equiv \max_{\mathbf{x}} \min_{\mathbf{y}} \mathbf{xAy}$ & $G^C \equiv \min_{\mathbf{y}} \max_{\mathbf{x}} \mathbf{xAy}$

Note: in general $\max_{\mathbf{x}} \min_{\mathbf{y}} \leq \min_{\mathbf{y}} \max_{\mathbf{x}}$

(from ECO 7504 notes)

Matching Pennies - has no pure strategy Nash equilibrium; play with mixed strategy (e.g., player 1 picks H with probably p and T with probability 1 - p)

		Player 2: q H	$1 - q$ T	$G^R(x_i)$
Player 1	p H	$\underline{1}$, -1	-1, $\underline{1}$	-1
	$1 - p$ T	-1, $\underline{1}$	$\underline{1}$, -1	-1
	$G^C(y_j)$	1	1	

Best Reply for Player 1 - look at boundary points first

$q = 1 \Rightarrow$ best reply is H ($p = 1$)
$q = 0 \Rightarrow$ best reply is T ($p = 0$)

In Between - compare expected payoffs for player 1 using H and T

EV_H^1 = 1q - 1(1 - q) = 2q - 1
EV_T^1 = -1q + 1(1 - q) = 1 - 2q

At q near 0, $EV_T^1 > EV_H^1$ so player 1 should play T; At q near 1, $EV_H^1 > EV_T^1$ so player 1 should play H

EV_T^1 and EV_H^1 converge at q = 0.5 at which point player 1 is indifferent between H and T

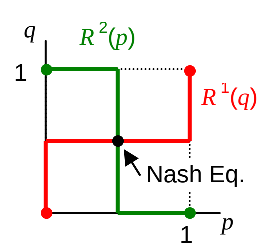

Best Reply for Player 2 - similar argument to player 1

$p = 1 \Rightarrow$ best reply is T ($q = 0$)
$p = 0 \Rightarrow$ best reply is H ($q = 1$)

Nash Equilibrium - best replies intersect at $p = q = 0.5$

(new notes)

Guarantee Levels - $G^R = \max_{\mathbf{x}}\left[\min_{\mathbf{y}}(-1,-1)\right] = -1$, $G^C = \min_{\mathbf{y}}\left[\max_{\mathbf{x}}(1,1)\right] = 1$

$\therefore\ G^R < G^C$ (no pure strategy equilibrium)

With mixed strategies ($p = q = 0.5$): $G^R = G^C = 0$

Symmetric Game - matrix game $\mathbf{A}$ is symmetric if $m = n$ and $A_{ij} = -A_{ji}$

	$\mathbf{y}_1$	$\mathbf{y}_2$
$\mathbf{x}_1$	0	-1
$\mathbf{x}_2$	1	0

Zero Diagonal - if players play same strategy, each gets payoff of zero (all elements of $\mathbf{A}$'s diagonal are zero)

Mirror Payoffs - if players swap strategies, they also swap payoffs; $u^R(\mathbf{x},\mathbf{y}) = -u^R(\mathbf{y},\mathbf{x}) = u^C(\mathbf{y},\mathbf{x})$

"Symmetrizing" - any finite game G based on payoff matrix $\mathbf{A}^{m\times n}$ can be "symmetrized"...

i.e., $\exists$ a symmetric game $\hat{G}$ which is equivalent to G

Why - we can limit the proof of the VNM result to looking at symmetric games only

How - denote row strategies with $\mathbf{x}$ and column strategies with $\mathbf{y}$; each player then chooses both a row and column strategy: player 1 chooses $\boldsymbol{\alpha} = (\mathbf{x}_\alpha, \mathbf{y}_\alpha)$ & player 2 chooses $\boldsymbol{\beta} = (\mathbf{x}_\beta, \mathbf{y}_\beta)$

Payoff - to player 1 is $\phi(\boldsymbol{\alpha},\boldsymbol{\beta}) \equiv u^R(x_\alpha, y_\beta) - u^R(x_\beta, y_\alpha)$

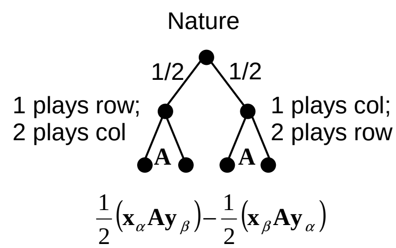

Note: we subtract the second payoff because in the second game, player 2 is the row player and payoffs are stated in terms of the row player

Symmetry - satisfies the two properties:

Zero diagonal... $\phi(\boldsymbol{\alpha},\boldsymbol{\alpha}) \equiv u^R(\mathbf{x}_\alpha, \mathbf{y}_\alpha) - u^R(\mathbf{x}_\alpha, \mathbf{y}_\alpha) = 0$

Mirror payoffs... $\phi(\boldsymbol{\beta},\boldsymbol{\alpha}) \equiv u^R(\mathbf{x}_\beta, \mathbf{y}_\alpha) - u^R(\mathbf{x}_\alpha, \mathbf{y}_\beta) = -\phi(\boldsymbol{\alpha},\boldsymbol{\beta})$

Analogy - schoolyard games (baseball, football, tennis) are not symmetric; players may not even face same strategy space; usually flip coin to create symmetry (to negate possible advantages of first/second mover)

Symmetrizing Theorem - there exists an equilibrium in game G *iff* there exists an equilibrium in $\hat{G}$ (the symmetrized version of G), specifically:

(a) $(\bar{\mathbf{x}},\bar{\mathbf{y}})$ is a Nash equilibrium in $G \Rightarrow \bar{\boldsymbol{\alpha}} = \bar{\boldsymbol{\beta}} = (\bar{\mathbf{x}},\bar{\mathbf{y}})$ is a Nash equilibrium in $\hat{G}$

(b) $(\bar{\boldsymbol{\alpha}},\bar{\boldsymbol{\beta}})$ is Nash equilibrium in $\hat{G} \Rightarrow (\bar{\mathbf{x}}_\alpha,\bar{\mathbf{y}}_\alpha)$ and $(\bar{\mathbf{x}}_\beta,\bar{\mathbf{y}}_\beta)$ are NE in G

<u>Proof</u>: (part a) assume $(\bar{\mathbf{x}},\bar{\mathbf{y}})$ is a Nash equilibrium in G

That means $\underset{\text{i}}{u^R(\bar{\mathbf{x}},\mathbf{y})} \geq \underset{\text{ii}}{u^R(\bar{\mathbf{x}},\bar{\mathbf{y}})} \geq \underset{\text{iii}}{u^R(\mathbf{x},\bar{\mathbf{y}})}\ \forall\ \mathbf{x}, \mathbf{y}$

i & ii say $\bar{\mathbf{y}}$ minimizes u^R (maxes u^C) when row player plays $\bar{\mathbf{x}}$ (i.e., $\bar{\mathbf{y}}$ is a best reply to $\bar{\mathbf{x}}$)

ii & iii say $\bar{\mathbf{x}}$ maximizes u^R when column player plays $\bar{\mathbf{y}}$ (i.e., $\bar{\mathbf{x}}$ is a best reply to $\bar{\mathbf{y}}$)

Using i & iii we have $u^R(\bar{\mathbf{x}},\mathbf{y}) - u^R(\mathbf{x},\bar{\mathbf{y}}) \geq 0\ \forall\ \mathbf{x}, \mathbf{y}$

In terms of $\hat{G}$ this is written $\phi(\bar{\boldsymbol{\alpha}},\boldsymbol{\beta}) \geq 0\ \forall\ \boldsymbol{\beta}$ (where $\bar{\boldsymbol{\alpha}} = (\bar{\mathbf{x}},\bar{\mathbf{y}})$ & $\boldsymbol{\beta} = (\mathbf{x},\mathbf{y})$)

Because $\hat{G}$ is symmetric, we know $\phi(\bar{\boldsymbol{\alpha}},\bar{\boldsymbol{\alpha}}) = 0$

Combine these: $\phi(\bar{\boldsymbol{\alpha}},\boldsymbol{\beta}) \ge \phi(\bar{\boldsymbol{\alpha}},\bar{\boldsymbol{\alpha}})\ \forall\ \boldsymbol{\beta}$

That says that $\bar{\boldsymbol{\alpha}} = (\bar{\mathbf{x}},\bar{\mathbf{y}})$ is player 2's best reply to $\bar{\boldsymbol{\alpha}}$ played by player 1

Using i & iii in reverse we have $u^R(\mathbf{x},\bar{\mathbf{y}}) - u^R(\bar{\mathbf{x}},\mathbf{y}) \le 0\ \forall\ \mathbf{x}, \mathbf{y}$

In terms of $\hat{G}$ this is written $\phi(\boldsymbol{\alpha},\bar{\boldsymbol{\alpha}}) \le 0\ \forall\ \boldsymbol{\alpha}$ (where $\boldsymbol{\alpha} = (\mathbf{x},\mathbf{y})$ & $\bar{\boldsymbol{\alpha}} = (\bar{\mathbf{x}},\bar{\mathbf{y}})$)

Because $\hat{G}$ is symmetric, we know $\phi(\bar{\boldsymbol{\alpha}},\bar{\boldsymbol{\alpha}}) = 0$

Combine these: $\phi(\boldsymbol{\alpha},\bar{\boldsymbol{\alpha}}) \le \phi(\bar{\boldsymbol{\alpha}},\bar{\boldsymbol{\alpha}})\ \forall\ \boldsymbol{\alpha}$

That says that $\bar{\boldsymbol{\alpha}} = (\bar{\mathbf{x}},\bar{\mathbf{y}})$ is player 1's best reply to $\bar{\boldsymbol{\alpha}}$ played by player 2

$\therefore\ (\bar{\boldsymbol{\alpha}},\bar{\boldsymbol{\alpha}})$ is a Nash equilibrium in $\hat{G}$

(part b) Assume $(\bar{\boldsymbol{\alpha}},\bar{\boldsymbol{\beta}})$ is Nash equilibrium in $\hat{G}$

Because $\hat{G}$ is symmetric, the payoff must be zero: $\phi(\bar{\boldsymbol{\alpha}},\bar{\boldsymbol{\beta}}) = 0$

Because it's a Nash equilibrium: $\phi(\bar{\boldsymbol{\alpha}},\boldsymbol{\beta}) \ge \phi(\bar{\boldsymbol{\alpha}},\bar{\boldsymbol{\beta}}) \ge \phi(\boldsymbol{\alpha},\bar{\boldsymbol{\beta}})\ \forall\ \boldsymbol{\alpha},\ \boldsymbol{\beta}$

(similar to case a: $\bar{\boldsymbol{\beta}}$ is best reply to $\bar{\boldsymbol{\alpha}}$ and vice versa)

Since the middle term is 0, we have $\phi(\bar{\boldsymbol{\alpha}},\boldsymbol{\beta}) \ge 0$

Expand that: $\phi(\bar{\boldsymbol{\alpha}},\boldsymbol{\beta}) = u^R(\bar{\mathbf{x}}_\alpha,\mathbf{y}_\beta) - u^R(\mathbf{x}_\beta,\bar{\mathbf{y}}_\alpha) \ge 0\ \forall\ \mathbf{x}_\beta,\ \mathbf{y}_\beta$

Since we can use any $\mathbf{y}_\beta$, use $\bar{\mathbf{y}}_\alpha$: $u^R(\bar{\mathbf{x}}_\alpha,\bar{\mathbf{y}}_\alpha) \ge u^R(\mathbf{x}_\beta,\bar{\mathbf{y}}_\alpha)\ \forall\ \mathbf{x}_\beta$

That means $\bar{\mathbf{x}}_\alpha$ is a best reply to $\bar{\mathbf{y}}_\alpha$

Since we can use any $\mathbf{x}_\beta$, use $\bar{\mathbf{x}}_\alpha$: $u^R(\bar{\mathbf{x}}_\alpha,\mathbf{y}_\beta) \ge u^R(\bar{\mathbf{x}}_\alpha,\bar{\mathbf{y}}_\alpha)\ \forall\ \mathbf{y}_\beta$

That means $\bar{\mathbf{y}}_\alpha$ is a best reply to $\bar{\mathbf{x}}_\alpha$

$\therefore\ (\bar{\mathbf{x}}_\alpha,\bar{\mathbf{y}}_\alpha)$ is a Nash equilibrium in game G

Repeat this procedure for $0 \ge \phi(\boldsymbol{\alpha},\bar{\boldsymbol{\beta}})$ and we'd show $(\bar{\mathbf{x}}_\beta,\bar{\mathbf{y}}_\beta)$ is NE in G

Harsanye - proposed a similar tool to turn an incomplete information game into a complete, but imperfect information game

Symmetric Game Lemma - a symmetric matrix game has a solution (Nash equilibrium) *iff* $\exists$ a semipositive vector $\bar{\mathbf{x}}$ (i.e., a mixed strategy) with $\bar{\mathbf{x}}\mathbf{A}\mathbf{y} \ge 0\ \forall$ semipositive $\mathbf{y}$

Proof: (Sufficient) Assume $(\bar{\mathbf{x}},\bar{\mathbf{y}})$ is a solution

That means $\bar{\mathbf{x}}\mathbf{A}\mathbf{y} \ge \bar{\mathbf{x}}\mathbf{A}\bar{\mathbf{y}}\ \forall\ \mathbf{y}$ (i.e., $\bar{\mathbf{y}}$ is a best reply to $\bar{\mathbf{x}}$) and $\bar{\mathbf{x}}\mathbf{A}\bar{\mathbf{y}} \ge \mathbf{x}\mathbf{A}\bar{\mathbf{y}}\ \forall\ \mathbf{x}$ (i.e., $\bar{\mathbf{x}}$ is a best reply to $\bar{\mathbf{y}}$)

Combine these: $\bar{\mathbf{x}}\mathbf{A}\mathbf{y} \ge \bar{\mathbf{x}}\mathbf{A}\bar{\mathbf{y}} \ge \mathbf{x}\mathbf{A}\bar{\mathbf{y}}\ \forall\ \mathbf{x},\ \mathbf{y}$

Since we can use any $\mathbf{x}$, pick $\mathbf{x} = \bar{\mathbf{y}}$ (we can do this because it's a symmetric game and $\mathbf{x}$ and $\mathbf{y}$ have the same dimensions)

From symmetry $\bar{\mathbf{y}}\mathbf{A}\bar{\mathbf{y}} = 0\ \therefore\ \bar{\mathbf{x}}\mathbf{A}\mathbf{y} \ge \bar{\mathbf{y}}\mathbf{A}\bar{\mathbf{y}} = 0$

(Necessary) Assume $\exists$ semipositive $\bar{\mathbf{x}}$ with $\bar{\mathbf{x}}\mathbf{A}\mathbf{y} \ge 0\ \forall$ semipositive $\mathbf{y}$

We know $\bar{\mathbf{x}}\mathbf{A}\bar{\mathbf{x}} = 0$ from symmetry

That means $\bar{\mathbf{x}}\mathbf{A}\mathbf{y} \ge \bar{\mathbf{x}}\mathbf{A}\bar{\mathbf{x}}\ \forall\ \mathbf{y}$

From symmetry, reverse strategies have mirrored payoffs: $\mathbf{y}\mathbf{A}\bar{\mathbf{x}} = -\bar{\mathbf{x}}\mathbf{A}\mathbf{y}\ \forall\ \mathbf{y}$

Combine that with previous inequality: $\mathbf{y}\mathbf{A}\bar{\mathbf{x}} \le \bar{\mathbf{x}}\mathbf{A}\bar{\mathbf{x}} = 0\ \forall\ \mathbf{y}$

Combine all the inequalities: $\bar{\mathbf{x}}\mathbf{A}\mathbf{y} \ge \bar{\mathbf{x}}\mathbf{A}\bar{\mathbf{x}} \ge \mathbf{y}\mathbf{A}\bar{\mathbf{x}}$

That means $(\bar{\mathbf{x}},\bar{\mathbf{x}})$ is a Nash equilibrium so the symmetric game has a solution

Simplifying $\bar{\mathbf{x}}\mathbf{A}\mathbf{y} \geq 0$ - since $\mathbf{y}$ is a semipositive vector, we know $y_i \geq 0$; $\therefore$ since $\bar{\mathbf{x}}\mathbf{A}\mathbf{y} \geq 0$, we know that $\bar{\mathbf{x}}\mathbf{A}$ can't have any negative components... so we can rewrite the second party of the symmetric game lemma to say $\exists$ a semipositive vector $\bar{\mathbf{x}}$ with $\bar{\mathbf{x}}\mathbf{A} \geq 0$

Theorem of the Alternative - if $\bar{\mathbf{x}}\mathbf{A} \geq 0$ has no semipositive solutions, then the system $\mathbf{A}\mathbf{y} < 0$ has a semipositive solution

Proof: need separating hyperplane theorem to prove this ("it's hard")

Existence Theorem - for any symmetric game with payoff matrix $\mathbf{A}$ $\exists$ a semipositive vector $\bar{\mathbf{x}}$ with $\bar{\mathbf{x}}\mathbf{A} \geq 0$ (Note: combined with the symmetric game lemma, $\bar{\mathbf{x}}\mathbf{A} \geq 0$ guarantees a solution exists for the symmetric game; then combined with the "symmetrizing" theorem there's a solution to the original game)

Proof: Assume $\bar{\mathbf{x}}\mathbf{A} \geq 0$ does not have a semipositive solution

Theorem of the Alternative says $\exists$ a semipositive $\mathbf{y}$ such that $\mathbf{A}\mathbf{y} < 0$

From symmetry, reverse strategies have mirrored payoffs so $\mathbf{y}\mathbf{A} > 0$

That's a contradiction so $\bar{\mathbf{x}}\mathbf{A} \geq 0$ does have a semipositive solution

Where Are We? - just showed that $\exists$ a Nash equilibrium in any two person, finite, constant sum game

Guarantee Level Same As Nash Eq. - general outline of proof; since there is limited applicability of constant sum games we won't spend more time on this proof; what we're trying to prove is if $G^R = G^C$ then there exists a Nash equilibrium; there are three possible relationships for G^R and G^C (guarantee levels):

$G^R > G^C$ - not possible; maxmin will always be $\leq$ minmax

$G^R = G^C$ - assume Nash equilibrium doesn't exist then get a contradiction

$G^R < G^C$ - assume Nash equilibrium does exist then get a contradiction

Properties of Nash Equilibrium (for constant sum games)

1) Same Payoffs - if $(\mathbf{x}^*, \mathbf{y}^*)$ and $(\hat{\mathbf{x}}, \hat{\mathbf{y}})$ are Nash equilibria, then $\mathbf{x}^*\mathbf{A}\mathbf{y}^* = \hat{\mathbf{x}}\mathbf{A}\hat{\mathbf{y}}$

Proof: Since $(\mathbf{x}^*, \mathbf{y}^*)$ is NE, we know $\mathbf{x}^*\mathbf{A}\mathbf{y} \geq \mathbf{x}^*\mathbf{A}\mathbf{y}^*$ $\forall$ $\mathbf{y}$

We can pick any $\mathbf{y}$, so let $\mathbf{y} = \hat{\mathbf{y}}$: $\mathbf{x}^*\mathbf{A}\hat{\mathbf{y}} \geq \mathbf{x}^*\mathbf{A}\mathbf{y}^*$

Since $(\hat{\mathbf{x}}, \hat{\mathbf{y}})$ is NE, we know $\hat{\mathbf{x}}\mathbf{A}\hat{\mathbf{y}} \geq \mathbf{x}\mathbf{A}\hat{\mathbf{y}}$ $\forall$ $\mathbf{x}$

We can pick any $\mathbf{x}$, so let $\mathbf{x} = \mathbf{x}^*$: $\hat{\mathbf{x}}\mathbf{A}\hat{\mathbf{y}} \geq \mathbf{x}^*\mathbf{A}\hat{\mathbf{y}}$

Chain the two inequalities together: $\hat{\mathbf{x}}\mathbf{A}\hat{\mathbf{y}} \geq \mathbf{x}^*\mathbf{A}\hat{\mathbf{y}} \geq \mathbf{x}^*\mathbf{A}\mathbf{y}^*$

Repeat this procedure starting with $(\hat{\mathbf{x}}, \hat{\mathbf{y}})$ and using $\mathbf{y} = \mathbf{y}^*$, then using $(\mathbf{x}^*, \mathbf{y}^*)$ and using $\mathbf{x} = \hat{\mathbf{x}}$ and we'd get $\mathbf{x}^*\mathbf{A}\mathbf{y}^* \geq \mathbf{x}^*\mathbf{A}\hat{\mathbf{y}} \geq \hat{\mathbf{x}}\mathbf{A}\hat{\mathbf{y}}$

In order for both of those strings of inequalities to be true, we must have $\mathbf{x}^*\mathbf{A}\mathbf{y}^* = \hat{\mathbf{x}}\mathbf{A}\hat{\mathbf{y}}$

Significance - players don't care which NE they pick, but this could lead to a coordination problem

Reverse Not True - having same payoff as the NE doesn't make a set of strategies a NE; example below has NE with expected payoff of 1; (x_3, y_3) also has expected payoff of 1, but it is not a Nash equilibrium

(min in row)

	y_1	y_2	y_3	$G^R(x_i) = \min_{y_j} u^R(x_i, y_j)$
x_1	1	-1	-2	-2
x_2	3	1	2	1
x_3	-4	0	1	-4
$\max_{x_i} u^R(x_i, y_j) = G^C(y_j)$	3	1	2	

$\left.\begin{matrix}-2\\1\\-4\end{matrix}\right\} G^R = \max G^R(x_i) = 1$

(max in column)

$\underbrace{3 \quad 1 \quad 2}_{G^C = \min G^C(y_j) = 1}$

2) No Coordination Problem - if $(\mathbf{x}^*, \mathbf{y}^*)$ and $(\hat{\mathbf{x}}, \hat{\mathbf{y}})$ are Nash equilibria, then $(\mathbf{x}^*, \hat{\mathbf{y}})$ and $(\hat{\mathbf{x}}, \mathbf{y}^*)$ are also Nash equilibria

Proof: From previous proof, we know $\mathbf{x}^*\mathbf{A}\mathbf{y}^* = \hat{\mathbf{x}}\mathbf{A}\hat{\mathbf{y}} = \mathbf{x}^*\mathbf{A}\hat{\mathbf{y}} = \hat{\mathbf{x}}\mathbf{A}\mathbf{y}^*$

Since $(\mathbf{x}^*, \mathbf{y}^*)$ is a NE $\Rightarrow$ $\mathbf{x}\mathbf{A}\mathbf{y}^* \le \mathbf{x}^*\mathbf{A}\mathbf{y}^* \ \forall\ \mathbf{x}$ ($\mathbf{x}^*$ is best reply to $\mathbf{y}^*$)

Sub $\mathbf{x}^*\mathbf{A}\mathbf{y}^* = \hat{\mathbf{x}}\mathbf{A}\mathbf{y}^*$: $\mathbf{x}\mathbf{A}\mathbf{y}^* \le \hat{\mathbf{x}}\mathbf{A}\mathbf{y}^* \ \forall\ \mathbf{x}$ ($\hat{\mathbf{x}}$ is best reply to $\mathbf{y}^*$)

Also since $(\hat{\mathbf{x}}, \hat{\mathbf{y}})$ is a NE $\Rightarrow$ $\hat{\mathbf{x}}\mathbf{A}\hat{\mathbf{y}} \le \hat{\mathbf{x}}\mathbf{A}\mathbf{y} \ \forall\ \mathbf{y}$ ($\hat{\mathbf{y}}$ is best reply to $\hat{\mathbf{x}}$)

Sub $\hat{\mathbf{x}}\mathbf{A}\hat{\mathbf{y}} = \hat{\mathbf{x}}\mathbf{A}\mathbf{y}^*$: $\hat{\mathbf{x}}\mathbf{A}\mathbf{y}^* \le \hat{\mathbf{x}}\mathbf{A}\mathbf{y} \ \forall\ \mathbf{y}$ ($\mathbf{y}^*$ is best reply to $\hat{\mathbf{x}}$)

$\therefore$ $(\hat{\mathbf{x}}, \mathbf{y}^*)$ is NE

(similar argument for $(\mathbf{x}^*, \hat{\mathbf{y}})$)

Significance - solves coordination problem; players can pick any equilibrium strategy without worrying about which equilibrium strategy opponent picks

3) Convex - the set of Nash equilibria is convex: if $(\mathbf{x}^*, \mathbf{y}^*)$ and $(\hat{\mathbf{x}}, \hat{\mathbf{y}})$ are Nash equilibria, then $(\lambda\mathbf{x}^* + (1-\lambda)\hat{\mathbf{x}}, \lambda\mathbf{y}^* + (1-\lambda)\hat{\mathbf{y}})$ ($\lambda \in [0,1]$) is also a Nash equilibrium

Proof: To simplify notation let $\mathbf{x}_\lambda = \lambda\mathbf{x}^* + (1-\lambda)\hat{\mathbf{x}}$ and $\mathbf{y}_\lambda = \lambda\mathbf{y}^* + (1-\lambda)\hat{\mathbf{y}}$

$\mathbf{x}_\lambda\mathbf{A}\mathbf{y}_\lambda = (\lambda\mathbf{x}^* + (1-\lambda)\hat{\mathbf{x}})\mathbf{A}\mathbf{y}_\lambda$

Break this up: $\lambda\mathbf{x}^*\mathbf{A}\mathbf{y}_\lambda + (1-\lambda)\hat{\mathbf{x}}\mathbf{A}\mathbf{y}_\lambda$

Because $(\mathbf{x}^*, \mathbf{y}^*)$ is NE, $\mathbf{x}^*\mathbf{A}\mathbf{y}_\lambda \ge \mathbf{x}^*\mathbf{A}\mathbf{y}^*$ ($\mathbf{y}^*$ is best reply to $\mathbf{x}^*$)

Because $(\hat{\mathbf{x}}, \hat{\mathbf{y}})$ is NE, $\hat{\mathbf{x}}\mathbf{A}\mathbf{y}_\lambda \ge \hat{\mathbf{x}}\mathbf{A}\hat{\mathbf{y}}$ ($\hat{\mathbf{y}}$ is best reply to $\hat{\mathbf{x}}$)

We know $\mathbf{x}^*\mathbf{A}\mathbf{y}^* = \hat{\mathbf{x}}\mathbf{A}\hat{\mathbf{y}}$ so that last inequality can be rewritten: $\hat{\mathbf{x}}\mathbf{A}\mathbf{y}_\lambda \ge \mathbf{x}^*\mathbf{A}\mathbf{y}^*$

Putting that all together, we have $\lambda\mathbf{x}^*\mathbf{A}\mathbf{y}_\lambda + (1-\lambda)\hat{\mathbf{x}}\mathbf{A}\mathbf{y}_\lambda \ge \mathbf{x}^*\mathbf{A}\mathbf{y}^*$

(convex combination of two numbers that are greater than a third is greater than the third number... i.e., if $a \ge c$ and $b \ge c$, then $\lambda a + (1-\lambda b) \ge c$ ($\lambda \in [0,1]$))

$\mathbf{x}_\lambda\mathbf{A}\mathbf{y}_\lambda = \mathbf{x}_\lambda\mathbf{A}(\lambda\mathbf{y}^* + (1-\lambda)\hat{\mathbf{y}})$

Break this up: $\lambda\mathbf{x}_\lambda\mathbf{A}\mathbf{y}^* + (1-\lambda)\mathbf{x}_\lambda\mathbf{A}\hat{\mathbf{y}}$

Because $(\mathbf{x}^*, \mathbf{y}^*)$ is NE, $\mathbf{x}_\lambda\mathbf{A}\mathbf{y}^* \le \mathbf{x}^*\mathbf{A}\mathbf{y}^*$ ($\mathbf{x}^*$ is best reply to $\mathbf{y}^*$)

Because $(\hat{\mathbf{x}}, \hat{\mathbf{y}})$ is NE, $\mathbf{x}_\lambda\mathbf{A}\hat{\mathbf{y}} \le \hat{\mathbf{x}}\mathbf{A}\hat{\mathbf{y}}$ ($\hat{\mathbf{x}}$ is best reply to $\hat{\mathbf{y}}$)

We know $\mathbf{x}^*\mathbf{A}\mathbf{y}^* = \hat{\mathbf{x}}\mathbf{A}\hat{\mathbf{y}}$ so that last inequality can be rewritten: $\mathbf{x}_\lambda\mathbf{A}\hat{\mathbf{y}} \le \mathbf{x}^*\mathbf{A}\mathbf{y}^*$

Putting that all together, we have $\lambda\mathbf{x}_\lambda\mathbf{A}\mathbf{y}^* + (1-\lambda)\mathbf{x}_\lambda\mathbf{A}\hat{\mathbf{y}} \le \mathbf{x}^*\mathbf{A}\mathbf{y}^*$

So that means $\mathbf{x}_\lambda \mathbf{A}\mathbf{y}_\lambda = \mathbf{x}^*\mathbf{A}\mathbf{y}^*$ (the convex combination of NE strategies has the same expected payoff as the NE strategies... but as shown on the previous page, this is a necessary condition for NE, not a sufficient condition)

Assume $(\mathbf{x}_\lambda, \mathbf{y}_\lambda)$ is not a Nash equilibrium

That means either $\exists\ \tilde{\mathbf{x}}$ with $\tilde{\mathbf{x}}\mathbf{A}\mathbf{y}_\lambda \geq \mathbf{x}_\lambda \mathbf{A}\mathbf{y}_\lambda$ (i.e., $\mathbf{x}_\lambda$ is not the best reply to $\mathbf{y}_\lambda$) or $\exists\ \tilde{\mathbf{y}}$ with $\mathbf{x}_\lambda \mathbf{A}\mathbf{y}_\lambda \geq \mathbf{x}_\lambda \mathbf{A}\tilde{\mathbf{y}}$ (i.e., $\mathbf{y}_\lambda$ is not the best reply to $\mathbf{x}_\lambda$)

Pick the first case and write out $\mathbf{y}_\lambda$: $\tilde{\mathbf{x}}\mathbf{A}\mathbf{y}_\lambda = \tilde{\mathbf{x}}\mathbf{A}(\lambda\mathbf{y}^* + (1-\lambda)\hat{\mathbf{y}})$

Now break that up: $\tilde{\mathbf{x}}\mathbf{A}\mathbf{y}_\lambda = \lambda\tilde{\mathbf{x}}\mathbf{A}\mathbf{y}^* + (1-\lambda)\tilde{\mathbf{x}}\mathbf{A}\hat{\mathbf{y}}$

Combine this with the first case inequality above: $\lambda\tilde{\mathbf{x}}\mathbf{A}\mathbf{y}^* + (1-\lambda)\tilde{\mathbf{x}}\mathbf{A}\hat{\mathbf{y}} \geq \mathbf{x}_\lambda \mathbf{A}\mathbf{y}_\lambda$

We showed earlier that $\mathbf{x}_\lambda \mathbf{A}\mathbf{y}_\lambda = \mathbf{x}^*\mathbf{A}\mathbf{y}^*$, so $\lambda\tilde{\mathbf{x}}\mathbf{A}\mathbf{y}^* + (1-\lambda)\tilde{\mathbf{x}}\mathbf{A}\hat{\mathbf{y}} \geq \mathbf{x}^*\mathbf{A}\mathbf{y}^*$

For this to be true, we must have either $\tilde{\mathbf{x}}\mathbf{A}\mathbf{y}^* \geq \mathbf{x}^*\mathbf{A}\mathbf{y}^*$ ($\mathbf{x}^*$ is <u>not</u> best reply to $\mathbf{y}^*$) or $\tilde{\mathbf{x}}\mathbf{A}\hat{\mathbf{y}} \geq \mathbf{x}^*\mathbf{A}\mathbf{y}^* = \hat{\mathbf{x}}\mathbf{A}\hat{\mathbf{y}}$ ($\hat{\mathbf{x}}$ is <u>not</u> best reply to $\hat{\mathbf{y}}$)... both cases contradict $(\mathbf{x}^*, \mathbf{y}^*)$ and $(\hat{\mathbf{x}}, \hat{\mathbf{y}})$ being NE

Similar reasoning with the second case ($\exists\ \tilde{\mathbf{y}}$ with $\mathbf{x}_\lambda \mathbf{A}\mathbf{y}_\lambda \geq \mathbf{x}_\lambda \mathbf{A}\tilde{\mathbf{y}}$)

$\therefore\ (\mathbf{x}_\lambda, \mathbf{y}_\lambda)$ is also a NE

Significance - same reason we like convex sets for consumer preferences: we can apply a fixed point theorem (which requires the mapping to be convex valued)

(from ECO 7504 notes)

1. If two (or more) strategies are played with positive probability by a player, then the expected value of those strategies are equal to each other
2. The expected payoff of strategies with positive probability must be at least as great as (≥) the expected value from strategies with zero probability

Player 2

			q_1	q_2	q_3	q_4
			W	**X**	**Y**	**Z**
Player 1	p_1	**A**	a_1, w_1	a_2, x_1	a_3, y_1	a_4, z_1
	p_2	**B**	b_1, w_2	b_2, x_2	b_3, y_2	b_4, z_2
	p_3	**C**	c_1, w_2	c_2, x_3	c_3, y_3	c_4, z_3
	p_4	**D**	d_1, w_4	d_2, x_4	d_3, y_4	d_4, z_4

Assume p_1 & $p_2 > 0$; p_3 & $p_4 = 0$

$$\overbrace{\sum_{j=1}^{4} a_j q_j = \sum_{j=1}^{4} b_j q_j}^{\text{Criteria 1}} \overbrace{\geq \sum_{j=1}^{4} c_j q_j}^{\text{Criteria 2}}$$

$$\sum_{j=1}^{4} a_j q_j = \sum_{j=1}^{4} b_j q_j \geq \sum_{j=1}^{4} d_j q_j$$

(new notes)

4a) Played Strategy Payoffs - assume $(\mathbf{x}^*, \mathbf{y}^*)$ is a Nash equilibrium; if $x_j^* > 0$ and $x_k^* > 0$ (i.e., these strategies are played with positive probability), then $\mathbf{A}_j\mathbf{y}^* = \mathbf{A}_k\mathbf{y}^*$ (i.e., the expected payoffs from those pure strategies are equal against $\mathbf{y}^*$ so the row player is indifferent between them when facing the mixed strategy $\mathbf{y}^*$ from the column player)

Notation - $\mathbf{A}_j$ is the j^{th} row of matrix $\mathbf{A}$; $\mathbf{A}^j$ is the j^{th} column of matrix $\mathbf{A}$

<u>Proof</u>: assume $\mathbf{A}_j\mathbf{y}^* = \mathbf{A}_k\mathbf{y}^*$ didn't hold; the row player would put more weight on the pure strategy with the higher expected payoff which means $\mathbf{x}^*$ is <u>not</u> best reply to $\mathbf{y}^*$

Column Player - if $y_j^* > 0$ and $y_k^* > 0$, then $\mathbf{x}^*\mathbf{A}^j = \mathbf{x}^*\mathbf{A}^k$

4b) Non-played Strategy Payoffs - assume $(\mathbf{x}^*, \mathbf{y}^*)$ is a Nash equilibrium; if $x_j^* > 0$ and $x_l^* = 0$ (i.e., the row player puts zero probability on using this strategy), then $\mathbf{A}_l \mathbf{y}^* \le \mathbf{A}_j \mathbf{y}^*$ (i.e., the pure strategy with zero probability has no greater expected payoff than those strategies with positive probability)

Generalizes - 4 a&b are the only properties that carry over to general (non constant sum) games

Problem - the counter intuitive result (e.g., change row player payoff doesn't change row player's mixed strategy, but changes column player's mixed strategy) causes some game theorists to say mixed strategies are not compelling; they argue if a player is indifferent between x_j^* and x_k^*, the assigned probability shouldn't matter, yet mixed strategy equilibria requires the player to chose the probabilities in order to make his opponent indifferent

Counter - point of analysis is to get non-intuitive results

5) Order Doesn't Matter - with pure strategy equilibria (in constant sum games), timing doesn't matter; for mixed strategies timing doesn't matter unless realization (actual result) occurs between decisions

Second Mover Advantage - if realization occurs before second player has to choose, the second player has an optimal pure strategy (vs. equilibrium mixed strategy) so timing matters

War Game Example - espionage to learn opponent's plans not good unless it learns the realization of the strategy: knowing aggressor will attack left 50% or right 50% does no good (reply by spreading defenses), but if espionage learns the realization (i.e., the result of the coin flip: left or right), then the defender has an advantage

Problems with Nash Equilibrium

Mixed Strategies - already covered this above

"Purification" of Nash - use nature to determine mixed strategy; basically playing a pure strategy conditional on nature; argument is that this reduces the "cost" of playing a mixed strategy; Slutsky: this "hasn't accomplished a huge amount"

Strong Nash - subgroup of players can't collude to improve their own payoffs

Problem - existence proofs don't exist and finding a strong Nash equilibrium is hard to do

Truth as Dominant Strategy - setup incentives to ensure truth is played; yields a Nash equilibrium, but it's not a strong Nash equilibrium

Strict Nash - unique best reply; not the same as strong, but just as hard to find

Summary

Guarantee Levels - $G^R \equiv \max_{\mathbf{x}_i} \min_{\mathbf{y}_j} u^R(\mathbf{x}_i, \mathbf{y}_j)$ & $G^C \equiv \min_{\mathbf{y}_j} \max_{\mathbf{x}_i} u^R(\mathbf{x}_i, \mathbf{y}_j)$

Nash Equilibrium - $u^R(\mathbf{x}_i^*, \mathbf{y}_j^*) \ge u^R(\mathbf{x}_i, \mathbf{y}_j^*)\ \forall\ \mathbf{x}_i$ and $u^R(\mathbf{x}_i^*, \mathbf{y}_j^*) \le u^R(\mathbf{x}_i^*, \mathbf{y}_j)\ \forall\ \mathbf{y}_j$

1) Same payoffs
2) No coordination problem
3) Convex
4) Played vs. non-played (expected value of strategies played are equal; non-played ≤)
5) Order doesn't matter