

A Few Games and Geometric Insights

Brian Powers

Arizona State University


brpowers@asu.edu

January 20, 2017

- 1 Nash/Correlated Equilibria
 - A Motivating Example
 - The Nash Equilibrium
 - Correlated Equilibrium
 - Relationship between Nash and Correlated Equilibria
- 2 Multi-Issue Final-Offer Arbitration
 - Single-Issue Case
 - Normal Distribution
 - Uniform Distribution
- 3 N-Player Final-Offer Arbitration

For Example

Two cars approach an intersection...



Each driver can choose to drive (“Dare”) or stop (“Chicken”).

The Game of "Chicken"

We have four possible outcomes of our game

	<i>2 : Dare</i>	<i>2 : Chicken</i>
<i>1 : Dare</i>	Collision	2 Passes
<i>1 : Chicken</i>	1 Passes	Both Stop

The Game of “Chicken”

The game with ordered pairs assigning rewards to players:

$$\begin{array}{cc} & D & C \\ D & (0, 0) & (7, 2) \\ C & (2, 7) & (6, 6) \end{array}$$

Each player secretly commits to a strategy, then both act simultaneously.

Finally each receives his component reward.

A General Bimatrix Games

In general, a 2-player finite game where players have m and n strategies respectively, we may summarize the game as an $m \times n$ “bimatrix”

$$(A, B) = \begin{matrix} & \begin{matrix} 1 & 2 & \cdots & n \end{matrix} \\ \begin{matrix} 1 \\ 2 \\ \vdots \\ m \end{matrix} & \left(\begin{array}{cccc} (a_{11}, b_{11}) & (a_{12}, b_{12}) & \cdots & (a_{1n}, b_{1n}) \\ (a_{21}, b_{21}) & (a_{22}, b_{22}) & \cdots & (a_{2n}, b_{2n}) \\ \vdots & \vdots & \ddots & \vdots \\ (a_{m1}, b_{m1}) & (a_{m2}, b_{m2}) & \cdots & (a_{mn}, b_{mn}) \end{array} \right) \end{matrix}$$

Let $I = \{1, \dots, m\}$ be the set of strategies for Player 1 (the row player) and $J = \{1, \dots, n\}$ the set of strategies for Player 2 (the column player).

Finite Game in Normal Form

Definition

A finite game Γ in *normal form* may be defined as a triple (N, S, u) where

- $N = \{1, \dots, n\}$ is the set of players
- $S = S_1 \times \dots \times S_n$ is the set of joint strategy profiles, S_i is the set of strategies available to player i
- $u : S \rightarrow \mathbb{R}^n$ is a utility function mapping each strategy profile to a payoff vector

A Solution for the Game

So is there a solution for the players?

John Nash (1951) proved the existence of equilibrium points for all finite games, provided players allow for randomized strategies.

In a Nash equilibrium, neither player has any incentive to deviate assuming the other player does not.

Definition

A pair of probability vectors (x^*, y^*) is a *Nash equilibrium* for bimatrix game (A, B) if

$$u_1(x, y^*) \leq u_1(x^*, y^*) \quad \forall x \in I \quad \text{and} \quad u_2(x^*, y) \leq u_2(x^*, y^*) \quad \forall y \in J$$

Nash Equilibria for “Chicken”

	<i>D</i>	<i>C</i>
<i>D</i>	(0, 0)	(7, 2)
<i>C</i>	(2, 7)	(6, 6)

There are three Nash equilibria for this game:

- Two *pure* equilibria, (D,C) and (C,D)
- One *mixed* equilibrium $((\frac{1}{3}, \frac{2}{3}), (\frac{1}{3}, \frac{2}{3}))$

Formalizing the Coordinated Solution

- A mediator makes a probability distribution over S known to all players.
- It randomly chooses an $s \in S$ according to the distribution and privately informs each player of his component strategy.
- If no player has an incentive to deviate knowing the other players' **conditional distributions**, the probability distribution is self-enforcing.

The Correlated Equilibrium

R. J. Aumann (1974) introduced *correlated equilibrium*.

Definition

A probability distribution μ over S is a *correlated equilibrium* if

$$\forall i \in N, \forall s_i \in S_i, \forall t_i \in S_i \setminus \{s_i\}$$
$$\sum_{s_{-i} \in S_{-i}} \mu(s) [u_i(s) - u_i(t_i, s_{-i})] \geq 0$$

The linear form

$$h_{s_i, t_i}(\mu) = \sum_{s_{-i} \in S_{-i}} \mu(s) [u_i(s) - u_i(t_i, s_{-i})]$$

is an *incentive constraint*.

First Proof of the Existence of CE

Aumann proved in 1974 that all games possess a correlated equilibrium.

Proof.

A Nash equilibrium $\sigma = (\sigma_1, \dots, \sigma_n)$ is a correlated equilibrium where

$$\mu(s) = \prod_{i \in N} \sigma_i(s_i)$$

and every game has a Nash equilibrium. □

The Correlated Equilibrium Set

The set of all correlated equilibria (\mathcal{C}) is defined by the following constraints:

$$\mu(s) \geq 0 \quad \forall s \in S \quad \text{Non-negativity}$$

$$\sum_{s \in S} \mu(s) = 1 \quad \text{Normalization}$$

$$h_{s_i, t_i}(\mu) \geq 0 \quad \forall i \in N, s_i \neq t_i \in S \quad \text{Incentives}$$

So \mathcal{C} is a polytope (the convex hull of a finite set of extreme points).

Extreme Correlated Equilibria for "Chicken"

	<i>D</i>	<i>C</i>
<i>D</i>	(0, 0)	(7, 2)
<i>C</i>	(2, 7)	(6, 6)


$$\blacksquare \mu_{CD} = \begin{bmatrix} 0 & 0 \\ 1 & 0 \end{bmatrix}$$

$$\blacksquare \mu_{DC} = \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix}$$

$$\blacksquare \mu_{N3} = \begin{bmatrix} 1/9 & 2/9 \\ 2/9 & 4/9 \end{bmatrix}$$

$$\blacksquare \mu_{C1} = \begin{bmatrix} 0 & 1/4 \\ 1/4 & 1/2 \end{bmatrix}$$

$$\blacksquare \mu_{C2} = \begin{bmatrix} 1/5 & 2/5 \\ 2/5 & 0 \end{bmatrix}$$



The Geometric Relationship Between NE and CE

- Hart and Schmeidler (1989) provided a new existence proof using linear methods which does not rely on the existence of Nash equilibria.
- We have already established that the set of Nash equilibria is a subset of \mathcal{C} .
- Raghavan and Evangelista (1996) have demonstrated that in bimatrix games the extreme points of so-called Nash sets (maximal convex sets of Nash equilibria) are extreme correlated equilibria.
- Nau et. al. (2004) have shown that in n -player games the Nash equilibria all lie on the (relative) boundary of \mathcal{C} .

Model for basic FOA (Brams-Merrill, 1983)

Player I (the minimizer) and Player II (the maximizer) each present a final offer. The arbitrator has an opinion of what he considers fair, and sides with whichever player's offer is closest to the fair settlement.

Assumptions:

- As far as players are concerned, the fair settlement is chosen randomly from a distribution F with density function f .
- F is common knowledge.
- WLOG, the median of the distribution is 0.
- The game is zero-sum.

Optimal Strategies for FOA

Say players choose x_1 and x_2 , while the arbitrator chooses ξ .

The payment made by Player I to Player II is

$$K(x_1, x_2 | \xi) = \begin{cases} x_1 & \text{if } |x_1 - \xi| < |x_2 - \xi| \\ x_2 & \text{if } |x_1 - \xi| > |x_2 - \xi| \end{cases}$$

If $|x_1 - \xi| = |x_2 - \xi|$, the payment is x_1 or x_2 with equal probability.

Brams-Merrill Theorem (1983)

Theorem

(1) If $f'(0)$ exists and $f(0) > 0$, then locally optimal strategies are


$$x_1^* = -\frac{1}{2f(0)} \quad \text{and} \quad x_2^* = \frac{1}{2f(0)}.$$

(2) If f is “sufficiently concentrated at the median”, then these represent the unique globally optimal strategy pair.


Brams and Merrill also provide a weaker condition for global optimality.

Both Normal and Uniform distributions satisfy the second condition.


Divergence of Global Optimal Pure Strategies




A. Double exponential




C. Logistic




E. Normal



B. Exponential



D. Triangular



F. Uniform



G. Cauchy

(Brams-Merrill, 1983)

Multiple-Issue FOA

When more than one issue is being arbitrated, two major variants of FOA have been used (Farber, 1980):

- **Issue by Issue:** Each party submits a vector of final offers and the arbitrator is free to compose a compromise by selecting some offers from each party
- **Whole Package:** Both parties submit a vector of final offers and the arbitrator must choose one or the other in its entirety

The Multi-Issue Game Setting

- Players I and II present final-offers $\mathbf{a}, \mathbf{b} \in \mathbb{R}^d$
- Judge selects $\xi \sim F$ as an ideal fair settlement.
- F is common knowledge.
- Judge uses reasonableness function $R(\mathbf{x}, \xi) : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}$ to compare final-offers.
- Game is zero-sum.
- Payoff is

$$K(\mathbf{a}, \mathbf{b} | \xi) = \begin{cases} \sum_i a_i & R(\mathbf{a}, \xi) > R(\mathbf{b}, \xi) \\ \sum_i b_i & R(\mathbf{a}, \xi) < R(\mathbf{b}, \xi) \end{cases}$$

Normal

$\xi \sim \mathcal{N}(\mu, \Sigma)$ (WLOG, assume $\mu = \mathbf{0}$)

Uniform

$\xi \sim \text{Unif} \left(\times_{j=1}^d [-\alpha_j, \alpha_j] \right)$, where $\alpha_j > 0$

Choice of R

Criterion	Reasonableness Function
Net Offer	$R_{NO}(\mathbf{x}, \boldsymbol{\xi}) = - \left \sum_{j=1}^d x_j - \xi_j \right $
L_1	$R_{L_1}(\mathbf{x}, \boldsymbol{\xi}) = - \sum_{j=1}^d x_j - \xi_j $
L_∞	$R_{L_\infty}(\mathbf{x}, \boldsymbol{\xi}) = - \max_j \{ x_j - \xi_j \}$
L_p	$R_{L_p}(\mathbf{x}, \boldsymbol{\xi}) = - \sum_{j=1}^d x_j - \xi_j ^p$
L_2	$R_{L_2}(\mathbf{x}, \boldsymbol{\xi}) = - \sum_{j=1}^d (x_j - \xi_j)^2$
Mahalanobis	$R_M(\mathbf{x}, \boldsymbol{\xi}) = -(\mathbf{x} - \boldsymbol{\xi})' \Sigma^{-1} (\mathbf{x} - \boldsymbol{\xi})$

Theorem

Any pure strategies from

$$S_i^* = \left\{ \left(x_1^*, (-1)^i \frac{\sqrt{2\pi(\sigma_x^2 + 2\rho\sigma_x\sigma_y + \sigma_y^2)}}{2} - x_i^* \right) : x_i^* \in \mathbb{R} \right\}$$

$i = 1, 2$ are independently optimal for Players I and II.

This is easily found as the game collapses to the one dimensional case.


Theorem

In $2NL_1, 2NL_\infty$ or $2NL_p$ if pure optimal strategies exist for Players $i = 1, 2$ then they are given by

$$(x_i^*, y_i^*) = ((-1)^i x^*, (-1)^i x^*),$$


where $x^* = \frac{\sqrt{2\pi(\sigma_x^2 + 2\rho\sigma_x\sigma_y + \sigma_y^2)}}{4}$.

$2NL_1$ Proof Sketch



As it is sub-optimal for either player to choose a pure strategy off the line $y = x$, the game reduces to the one-dimensional case.


$2NL_\infty$ Proof Sketch



As it is sub-optimal for either player to choose a pure strategy off the line $y = x$, the game reduces to the one-dimensional case.

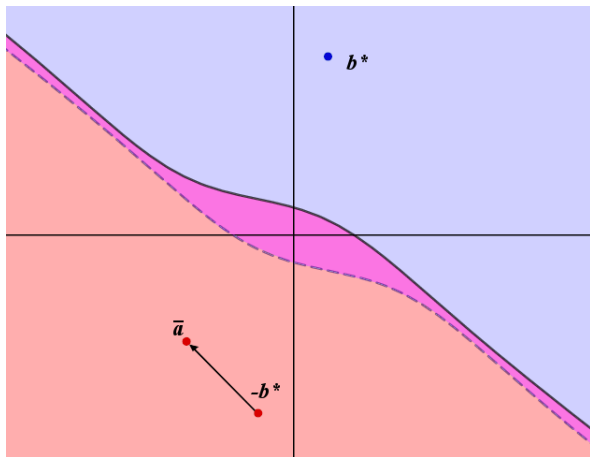
$2NL_p$ Circles

$$D_{L_p}(\mathbf{x}, \mathbf{y}) = \left(\sum_{i=1}^d |x_i - y_i|^p \right)^{1/p}$$




“Circles” and Midset curves under Minkowski metrics (L_p)

$2NL_p$ Proof Sketch



As it is sub-optimal for either player to choose a pure strategy off the line $y = x$, the game reduces to the one-dimensional case.

$2NL_p$ Proof Sketch (cont.)



This is because the midset curve between two such points $-\mathbf{x}^*, \mathbf{x}^*$ in L_p nowhere has a derivative equal to -1 .

$2NL_2$ Local Optimality

Theorem

In $2NL_2$, suppose $\rho > \max \left\{ -\frac{\sigma_x^2 + 3\sigma_y^2}{4\sigma_x\sigma_y}, -\frac{3\sigma_x^2 + \sigma_y^2}{4\sigma_x\sigma_y} \right\}$. The pure strategy pair in the previous theorem is locally optimal.

Because $Mid_{L_2}[\mathbf{a}, \mathbf{b}]$ is a straight line, we can essentially reduce the dimension of F .

$2NL_2$ Payoff Function

Let $(\xi, \eta) \sim \mathcal{N}(\mathbf{0}, \Sigma)$ be the opinion of the arbitrator, where

$$\Sigma = \begin{bmatrix} \sigma_x^2 & \rho\sigma_x\sigma_y \\ \rho\sigma_x\sigma_y & \sigma_y^2 \end{bmatrix}$$

$C_1(\mathbf{a}, \mathbf{b}) = \{(x, y) | (x_1 - x)^2 + (y_1 - y)^2 < (x_2 - x)^2 + (y_2 - y)^2\}$
and $C_2(\mathbf{a}, \mathbf{b})$ defined similarly.

$$P_i = P((\xi, \eta) \in C_i(\mathbf{a}, \mathbf{b}))$$

Assuming $\mathbf{a} \neq \mathbf{b}$, the expected payoff

$$K(\mathbf{a}, \mathbf{b}) = (x_1 + y_1)P_1 + (x_2 + y_2)P_2$$

may be written

$$K(\mathbf{a}, \mathbf{b}) = (x_2 + y_2) + (x_1 + y_1 - x_2 - y_2)P_1.$$

$2NL_2$ Local Optimality Proof Overview cont.

Express P_1 one-dimensionally:

$$P_1 = P((x_1 - \xi)^2 + (y_1 - \eta)^2 < (x_2 - \xi)^2 + (y_2 - \eta)^2) \quad (1)$$

$$= P\left((x_2 - x_1)\xi + (y_2 - y_1)\eta < \frac{x_2^2 + y_2^2 - x_1^2 - y_1^2}{2}\right) \quad (2)$$

$$= P(Z < z) \quad (3)$$

where

$$z = \frac{x_2^2 + y_2^2 - x_1^2 - y_1^2}{2\sqrt{(\mathbf{b} - \mathbf{a})^T \Sigma (\mathbf{b} - \mathbf{a})}} \quad (4)$$

we may write

$$K(\mathbf{a}, \mathbf{b}) = (x_2 + y_2) + (x_1 + y_1 - x_2 - y_2)\Phi(z) \quad (5)$$

where $\Phi(z)$ is the standard Gaussian cdf.

2NL₂ Local Optimality Proof Overview cont.

By solving the system of first-order equations

$$\frac{d}{dx_1} K|_{\mathbf{a}^*, \mathbf{b}^*} = \frac{d}{dy_1} K|_{\mathbf{a}^*, \mathbf{b}^*} = \frac{d}{dx_2} K|_{\mathbf{a}^*, \mathbf{b}^*} = \frac{d}{dy_2} K|_{\mathbf{a}^*, \mathbf{b}^*} = 0$$

we arrive at the unique \mathbf{a}^* , \mathbf{b}^* given in the theorem. It is *straightforward* to verify that the second order condition holds provided

$$\rho > \max \left\{ -\frac{\sigma_x^2 + 3\sigma_y^2}{4\sigma_x\sigma_y}, -\frac{3\sigma_x^2 + \sigma_y^2}{4\sigma_x\sigma_y} \right\}.$$

$2NL_2$ Global Optimality


Theorem

If $\rho > 0$, the solution points \mathbf{a}^ , \mathbf{b}^* given in the previous theorem are globally optimal.*

In other words, $K(\mathbf{a}, \mathbf{b}^*) \geq 0 \forall \mathbf{a} \in \mathbb{R}^2$, with equality only when $\mathbf{a} = \mathbf{a}^*$.


Thus players need not consider mixed strategies. The proof relies on a geometric interpretation of the players' strategies.

$2NL_2$ Global Optimality Proof Overview



$$z = \frac{2x^{*2} - x_1^2 - y_1^2}{2\sqrt{(\mathbf{b}^* - \mathbf{a})^T \Sigma (\mathbf{b}^* - \mathbf{a})}}$$


$2NL_2$ Global Optimality: Proof Overview cont.



If $K(\mathbf{a}, \mathbf{b}^*) \leq 0$ then $x_1 + y_1 < 0$ and either

$$x_1^2 + y_1^2 < 2x^{*2} \quad \text{or} \quad x_1 + y_1 \leq -2x^*.$$


$2NL_2$ Global Optimality: Proof Overview cont.



$$K(\mathbf{a}, \mathbf{b}^*) = 2x^* + \frac{1}{2}(x_1 + y_1 - 2x^*) = 2x^* + x_1 + y_1 \geq 0$$

when $z = 0$, with equality only when $\mathbf{a} = (-x^*, -x^*)$.

$2NL_2$ Global Optimality: Proof Overview cont.



Against Player II's strategy $\mathbf{b}^* = (x^*, x^*)$, any pure strategy

$\mathbf{a} = (x_1, y_1)$ may be represented as

$$\mathbf{a}(r, \theta) = (x^* + r \cos \theta, x^* + r \sin \theta).$$

$2NL_2$ Global Optimality: Polar Representation

Letting $t(\theta) = -\cos\theta - \sin\theta$,

$$K(\mathbf{a}, \mathbf{b}^*) = 2x^* - rt(\theta)\Phi(z)$$

So $K < 0$ is equivalent to

$$\phi(z) > \frac{2x^*}{rt(\theta)} = f(r, \theta)$$

$2NL_2$ Global Optimality: Two Tricks

To avoid the difficulties inherent in $\Phi(z)$, we use two tricks:


(1) For $z < 0$, the normal cdf is bounded by the sigmoidal

$$\Phi(z) < \frac{1}{1 + \exp\left(-\sqrt{\frac{8}{\pi}}z\right)}$$

(2) For $z > 0$, by its concavity, $\Phi(z) < y(z)$, the line tangent at $z = 0$

$2NL_2$ Global Optimality: Proof Overview cont.

For fixed $\theta \in \left[\frac{3\pi}{4}, \frac{7\pi}{4}\right]$:



$2UL_2$ - Globally Optimal Pure Strategies


Suppose ξ is drawn uniformly at random from $\Xi := [-\alpha, \alpha] \times [-\beta, \beta]$, where WLOG $0 < \alpha \leq \beta$, and the judge uses the L_2 metric.

Theorem

In $2UL_2$, the strategy pair $\mathbf{a}^ = (-\frac{\beta}{2}, -\frac{\beta}{2})$, $\mathbf{b}^* = (\frac{\beta}{2}, \frac{\beta}{2})$ is the unique globally optimal strategy pair.*


To prove this, we let Player II play \mathbf{b}^* and show that the expected payoff is minimized only when Player I plays \mathbf{a}^* .

$2UL_2$ Proof - Cases




$2UL_2$ Proof Case 1

In the first case, we can show directly that the payoff function is minimized only at $\mathbf{a} = \mathbf{a}^*$, which lies in this region.





$2UL_2$ Proof Case 2

Here we parameterize the strategy of Player I along line segments by slope m and $\lambda \in [0, 1]$, and show that the payoff function is a decreasing function of λ , and on the boundary the payoff is positive.




$2UL_2$ Proof Case 3

We parameterize the strategy by $p \in [.5, 1]$ (i.e. P_1) and $\bar{x} \in [0, 2\alpha]$ (the length of the upper boundary of C_1) to show there is no local minima.



$2UL_2$ Proof Case 4

The final case is handled directly; it is shown by the first order condition that no minimum to the payoff function exists in this region.



Generalize to N players

A unit must be split between the players. Each player i chooses a vector $\mathbf{x}^i \in \Delta^N$, where $\Delta^N = \{\mathbf{x} \mid \mathbf{x} \in \mathbb{R}^N, \sum x_i = 1, \mathbf{x} \geq 0\}$. Suppose each player i supplies evidence of strength $\lambda_i \geq 0$ in her favor to the judge. If $\lambda_i = 0$ then the player has supplied no evidence in her favor. Suppose that based on this evidence the judge decides on a fair split of the unit award. Let $\xi = (\xi_1, \xi_2, \dots, \xi_N)$ be the fair split, where $\xi \in \Delta^N$.

Dirichlet Distribution



Assume that it is common knowledge among the players that ξ will be drawn from a Dirichlet distribution with parameter $\alpha = (\alpha_1, \dots, \alpha_N)$, where $\alpha_i = \lambda_i + 1$. The density function is

$$f(\mathbf{x}) = \frac{\prod_{i=1}^N x_i^{\lambda_i}}{B(\alpha)}$$

Where $B(\alpha)$ is a normalizing constant.

Voronoi Cells in the Simplex

Given the N final-offers, we may partition Δ^N into N convex Voronoi cells. Call these V_i for $i = 1, \dots, N$.



If $\lambda_i = 0$ for all players, the probability distribution is uniform over the simplex. In this case, the payoff function is

$$K_i(\mathbf{x}^1, \dots, \mathbf{x}^N) = (N - 1)! \sum_{j=1}^N \mathbf{x}_i^j \iint_{V_j} \frac{N! \sqrt{2^N}}{\sqrt{N + 1}} dV_j$$

Conjecture: Let $N \geq 3$. Players $2, \dots, N$ demand β and offer $\frac{1-\beta}{N-1}$ to the opponents. Player 1 determines to demand α . Then P_1 is maximized when Player 1 offers $\frac{1-\alpha}{N-1}$ to each other player.

Theorem

For an N player FOA game where ξ is chosen uniformly at random, assuming the conjecture, a pure equilibrium strategy is for each player to demand $\frac{H_{N-1}}{N-1}$ for himself and offer the remaining portion equally to the other players.¹

References



Robert J. Aumann (1974)

Subjectivity and correlation in randomized strategies.

Journal of Mathematical Economics, 1(1):67-96, March 1974.



Fe.S. Evangelista and T.E.S. Raghavan (1996)

A Note on correlated equilibrium.

International Journal of Game Theory, 25(1):35-41.



Sergiu Hart and David Schmeidler (1989)

Existence of correlated equilibria.

Mathematics of Operations Research, 14(1):pp. 18-25.






John Nash (1951)

Non Cooperative Games.

The Annals of Mathematics, 54(2): 286-295, September 1951.

References

-  Robert Nau, Sabrina Gome Canovas, and Pierre Hansen (2004)
On the geometry of nash equilibria and correlated equilibria.
International Journal of Game Theory, 32(4):443-453.
-  Brams, S. J., & Merrill III, S. (1983)
Equilibrium strategies for final-offer arbitration: There is no median convergence.
Management Science, 29(8), 927-941.
-  Stevens, C. M. (1966)
Is compulsory arbitration compatible with bargaining?
Industrial Relations, 5(February): 38-52.

Thank You