Normal = Strategic Form Games

Prisoner's Dilemma Game

- two prisoner's in separate cells
- confess or not
- payoff matrix (matrix game)

	C	N
C	2,2	0,3
N	3,0	1,1

Or: contribute or not to a public good worth two each at a cost of three

Or: create an externality (fishing)

- no mind reading!!!

Finite Strategic = Normal Form Games

an N player game $i = 1 \ldots N$

finite strategy spaces S_i

$s \in S \equiv \times_{i=1}^{N} S_i$ are the strategy profiles

other useful notation $s_{-i} \in S_{-i} \equiv \times_{j \neq i} S_j$

$u_i(s)$ payoff or utility

Dominance and Rationalizability

σ_i weakly (strongly) dominates σ'_i if $u_i(\sigma_i, s_{-i}) \geq (>) u_i(\sigma'_i, s_{-i})$ with at least one strict

Prisoner's Dilemma Game

	C	N
C	2,2	0,3
N	3,0	1,1

a unique dominant strategy equilibrium (N,N)

this is Pareto dominated by (C,C) does it really occur??

Public Goods Experiment

Players randomly matched in pairs

May donate or keep a token

The token has a fixed commonly known public value of 15

It has a randomly drawn private value uniform on 10-20

V=private gain/public gain

So if the private value is 20 and you donate you lose 5, the other player gets 15; V= -1/3

If the private value is 10 and you donate you get 5 the other player gets 15; V=+1/3

Data from Levine/Palfrey, experiments conducted with caltech undergraduates

Based on Palfrey and Prisbey

V	donating a token
0.3	100%
0.2	92%
0.1	100%
0	83%
-0.1	55%
-0.2	13%
-0.3	20%

Weak Dominance and the Second Price Auction

a single item is to be auctioned.

value to the seller is zero.

$i = 1,\ldots,N$ buyers

value $v_i > 0$ to buyer *i*.

each buyer submits a bid b_i

the item is sold to the highest bidder at the second highest bid

- bidding your value is weakly dominant
- BDM mechanism with random “second highest bid”
- The endowment effect

This ticket is worth $2.00 to you.

You can sell it.

Name your offer price.

A price will be posted shortly

The posted price was drawn randomly between:

[$ 0 and $ 6]

If your offer price is **below** the posted price then you sell your ticket at the posted price.

If your offer price is **above** the posted price then you do not sell your ticket but you do collect the $2.00 value of the ticket.

You can view the posted price after you have named your price.

Indicate the appropriate amount .

My offer price is **below** the posted price.

Pay me the posted price of $__________.

My offer price is **above** the posted price.

Pay me $ 2.00.

Iterated Dominance: Example

	L	R-l	R-r
U-u	-1,-1	2,0	1,1
U-d	-1,-1	1,-1	0,0
D	1,1	1,1	1,1

Eliminate U-d then Eliminate L

Eliminate D (or) Eliminate R-l

Eliminate R-l

Notice that there can be more than one answer for iterated weak dominance

Not for iterated strong dominance

Cournot Duopoly Example

Two firms produce x_1, x_2 with aggregate output $x = x_1 + x_2$

Inverse demand is $17 - x$

marginal cost is constant and equal to 1

profits are therefore $\pi_i = [17 - (x_i + x_{-i})]x_i - x_i$

reaction function is derived by maximizing with respect to x_i

derivative $16 - 2x_i - x_{-i}$

observe that second derivative is negative

solve derivative equal to zero $x_i = 8 - x_{-i}/2$

Iterated Strict Dominance

$\pi_i = [17 - (x_i + x_{-i})]x_i - x_i$; reaction function $x_i = 8 - \frac{x_{-i}}{2}$

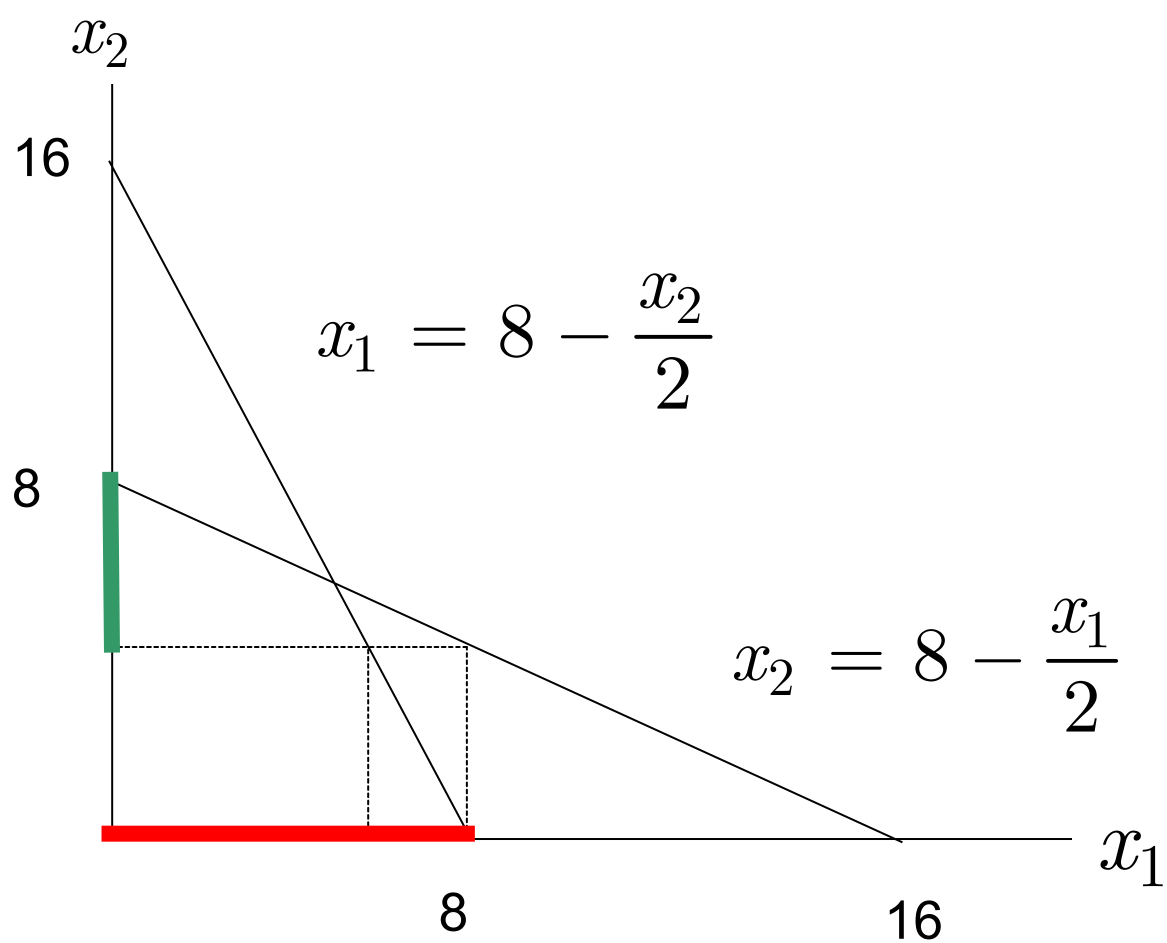

Mixed Strategies

$P(S)$ are probability measure on S

$\sigma_i \in \Sigma_i \equiv P(S_i)$ are mixed strategies, $\sigma \in \Sigma \equiv \times_{i=1}^{N} \Sigma_i$

$\sigma_{-i} \in \Sigma_{-i} \equiv \times_{j \neq i} \Sigma_j$

$u_i(\sigma) \equiv \sum_{s \in S} u_i(s) \prod_{j=1}^{N} \sigma_j(s_j)$ *is expected utility*

Nash Equilibrium: Definition

players can anticipate on another's strategies

σ is a *Nash equilibrium* profile if for each $i \in 1, \ldots N$
$u_i(\sigma) = \max_{\sigma'_i} u_i(\sigma'_i, \sigma_{-i})$

Voting

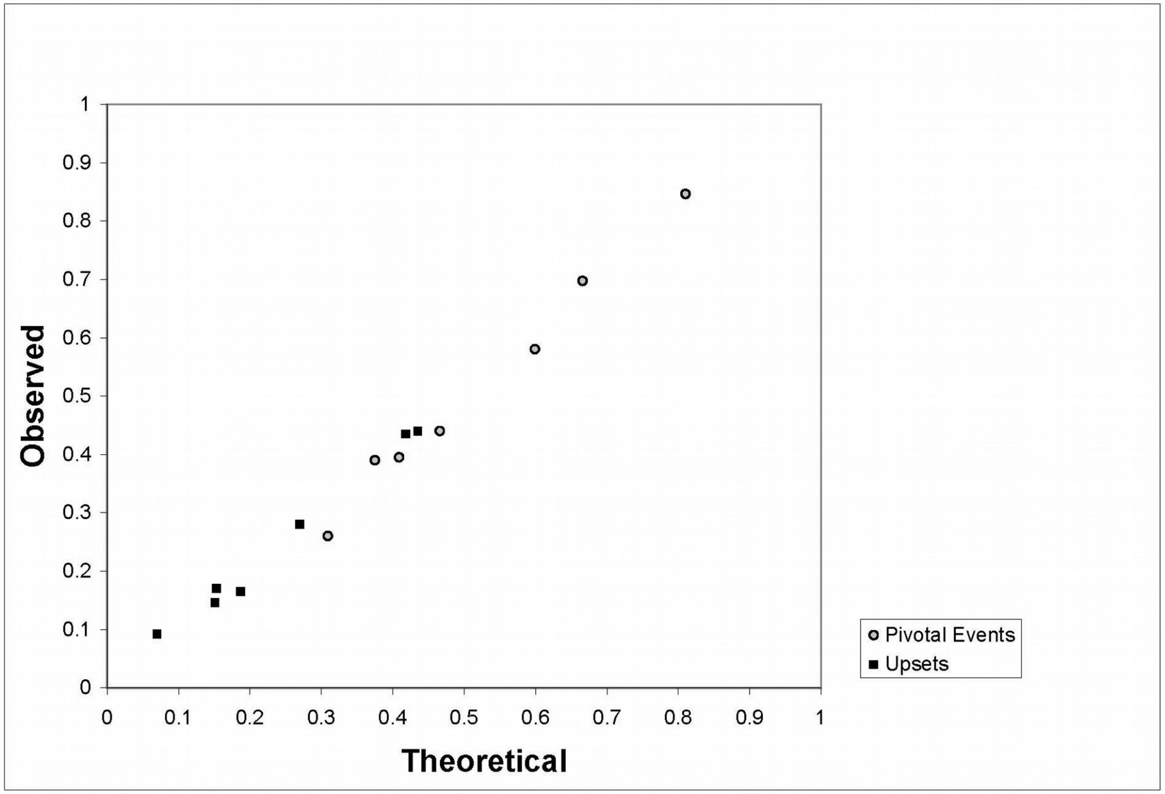

Observed
1
0.9
0.8
0.7
0.6
0.5
0.4
0.3
0.2
0.1
0
0 0.1 0.2 0.3 0.4 0.5 0.6 0.7 0.8 0.9 1
Theoretical
Pivotal Events
Upsets

Mixed Strategies: How Do Athletes Do It?

- Holmes, Moriarity, Canterbury and Dover
- once in Japan catchers equipped with mechanical randomization devices to call the pitch
- later ruled unsporting and banned from play
- good tennis players in important matches do it right
- professional soccer players do it right
- *submarine captains and the RAND corporation*

Existence

Theorem: a Nash equilibrium exists in a finite game

This theorem fails in pure strategies:

consider matching pennies; Holmes and Moriarity this is more or less why Kakutani's fixed point theorem was invented:

An upper hemi-continuous (UHC) convex valued correspondence B from a convex subset $B:\Sigma \to \Sigma$ to itself has a fixed point $\sigma \in B(\sigma)$

A correspondence $B:\Sigma \to \Sigma$ is UHC means if $\sigma^n \to \sigma$ such that $b^n \in B(\sigma^n), b^n \to b$ then $b \in B(\sigma)$

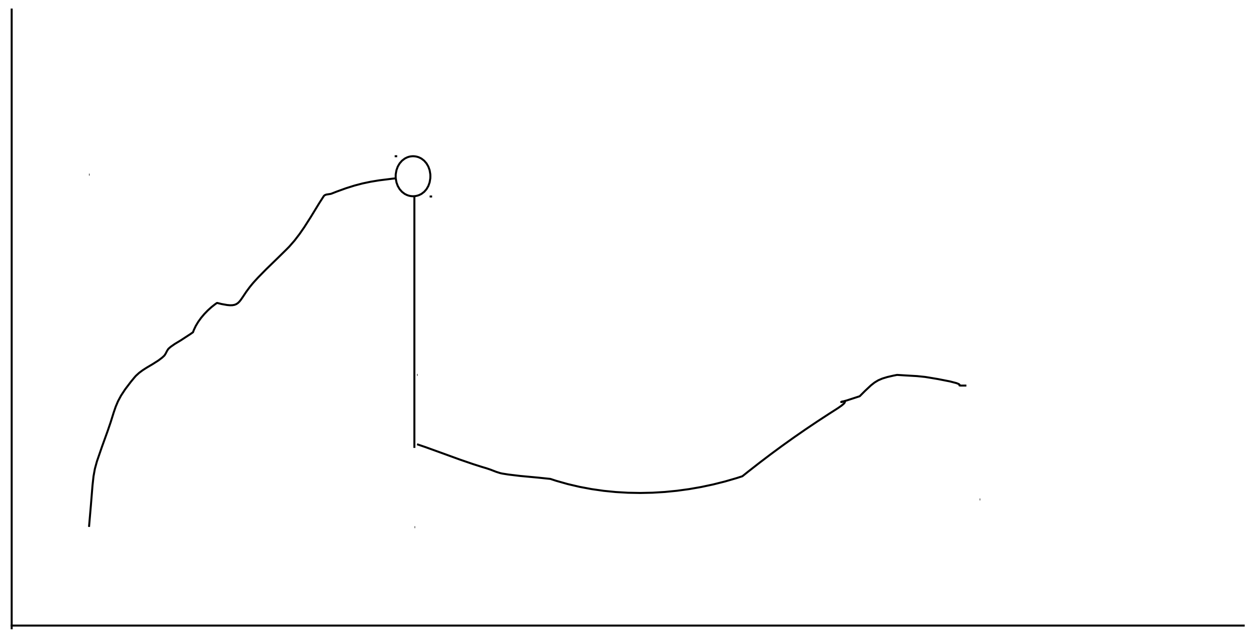

Proof: Let $B_i(\sigma)$ be the set of best responses of i to σ_{-i}

convex valued: convex combinations of a best response is a best response. Specifically, since you must be indifferent between all pure strategies played with positive probability, the best response set is the set of all convex combinations of the pure strategies that are best responses.

UHC: $b^n \in B(\sigma^n), b^n \to b$ means that $u^i(b_i^n, \sigma_{-i}^n) \geq u^i(\sigma_i, \sigma_{-i}^n)$. Suppose the conversely that $b \notin B(\sigma)$. This means for some $\widehat{\sigma}$ that $u^i(\widehat{\sigma}, \sigma_{-i}) > u^i(b_i, \sigma_{-i})$. Since $\sigma_{-i}^n \to \sigma_{-i}$ for n sufficiently large, since u^i is continuous (multi-linear in fact) in σ_{-i} we have $u^i(\widehat{\sigma}, \sigma_{-i}^n) > u^i(b_i, \sigma_{-i}^n)$. Since $b_i^n \to b_i$, since u^i is continuous (linear in fact) in σ_i, also for n sufficiently large $u^i(\widehat{\sigma}, \sigma_{-i}^n) > u^i(b_i^n, \sigma_{-i}^n)$. This contradicts .

$b_i^n \in B(\sigma^n)$

“a sequence of best-responses converges to a best-response”

Best Response Correspondence Example

	L ($\sigma_2(L) = q$)	R
U ($\sigma_1(U) = p$)	1,1	0,0
D	0,0	1,1

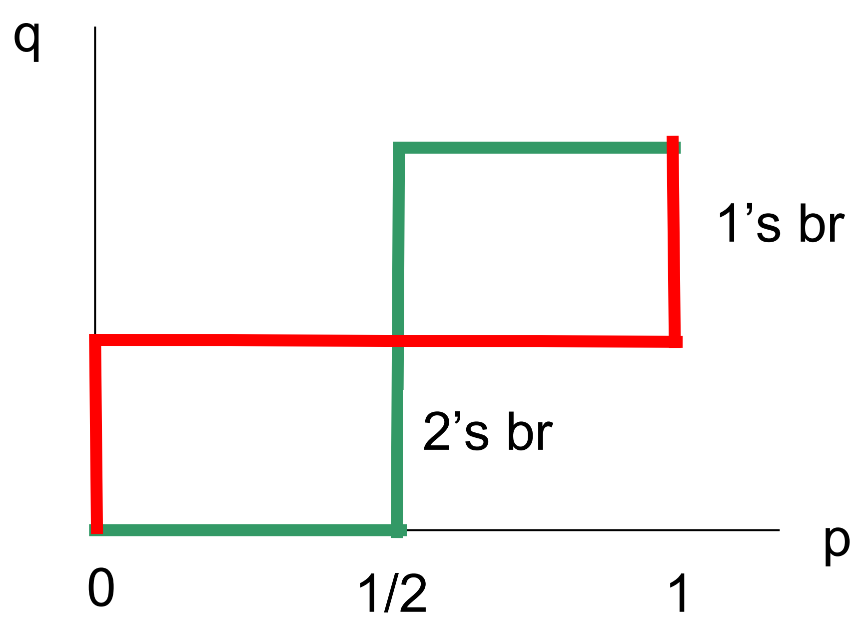

Mixed Strategies: The Kitty Genovese Problem

Description of the problem

Model:

n people all identical

benefit if someone calls the police is x

cost of calling the police is 1

Assumption: $x > 1$

Look for symmetric mixed strategy equilibrium where *p* is probability of each person calling the police

p is the symmetric equilibrium probability for each player to call the police

each player *i* must be indifferent between calling the police or not

if *i* calls the police, gets $x - 1$ for sure.

If *i* doesn't, gets 0 with probability $(1-p)^{n-1}$, gets x with probability $1-(1-p)^{n-1}$

so indifference when $x - 1 = x\left(1-(1-p)^{n-1}\right)$

solve for $p = 1-(1/x)^{1/(n-1)}$

probability police is called

$$1-(1-p)^n = 1-\left(\frac{1}{x}\right)^{\frac{n}{n-1}}$$

probability police are called

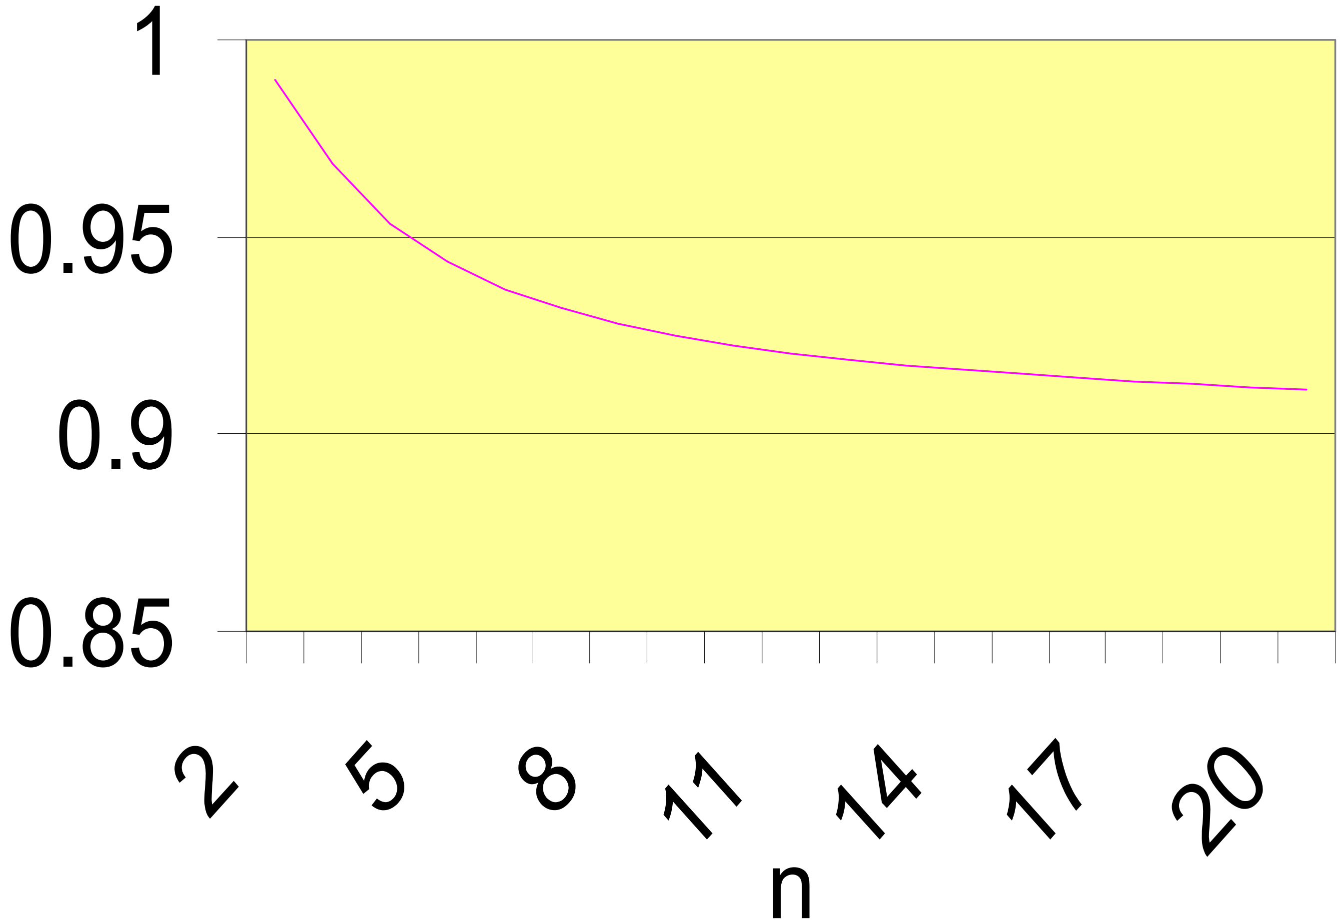

x=10

Coordination Games

	L	R
U	1,1	0,0
D	0,0	1,1

three equilibria (U,L) (D,R) (.5U,.5R)

too many equilibria?? introspection possible?

the rush hour traffic game – introspection clearly impossible, yet we seem to observe Nash equilibrium

equilibrium through learning?

Coordination Experiments

Van Huyck, Battalio and Beil [1990]

Actions $A = \{1, 2, \ldots \overline{e}\}$

Utility $u(a_i, a_{-i}) = b_0 \min(a_j) - b a_i$ where $b_0 > b > 0$

Everyone doing a' the same thing is always a Nash equilibrium

$a' = \overline{e}$ is efficient

the bigger is a' the more efficient, but the "riskier"

a model of "riskier" some probability of one player playing $a' = 1$

story of the stag-hunt game

$\bar{e} = 7$, 14-16 players

treatments: A $b_0 = 2b$; B $b = 0$

In final period treatment A:

77 subjects playing $a_i = 1$

30 subjects playing something else

minimum was always 1

In final period treatment B:

87 subjects playing $a_i = 7$

0 playing something else

with two players $a_i = 7$ was more common

1/2 Dominance

Coordination Game

	L (p_2)	R
U (p_1)	2,2	-10,0
D	0,-10	1,1

risk dominance:

indifference between U,D

$$2p_2 - 10(1-p_2) = (1-p_2)$$

$$13p_2 = 11, p_2 = 11/13$$

if U,R opponent must play equilibrium w/ 11/13

if D,L opponent must play equilibrium w/ 2/13

½ dominance: if each player puts weight of at least ½ on equilibrium strategy, then it is optimal for everyone to keep playing equilibrium

(same as risk dominance in 2x2 games)

Trembling Hand Perfection

σ is trembling hand perfect if there is a sequence $\sigma^n >> 0, \sigma^n \to \sigma$ such that

if $\sigma^i(s^i) > 0$ then s^i is a best response to σ^n

Note: thp is necessarily a Nash equilibrium

Examples:

strict Nash equilibrium is always thp

completely mixed Nash equilibrium is always thp

Example of Non-Trembling Hand Perfect Equilibrium

	L	R
U	-1,-1	2*,0*
D	0*,2*	0,2*

Correlated Equilibrium

Chicken

6,6	2,7
7,2	0,0

three Nash equilibria (2,7), (7,2) and mixed equilibrium w/ probabilities (2/3,1/3) and payoffs

(4 2/3, 4 2/3)

6,6	2,7
7,2	0,0

correlated strategy

1/3	1/3
1/3	0

is a correlated equilibrium giving utility (5,5)

What is public randomization?

Approximate Equilibria and Near Equilibria

- exact: $u_i(s_i|\sigma_i) \geq u_i({s_i}'|\sigma_i)$

  approximate: $u_i(s_i|\sigma_i) + \epsilon \geq u_i({s_i}'|\sigma_i)$

- Approximate equilibrium can be very different from exact equilibrium

Radner's work on finite repeated PD

gang of four on reputation

upper and lower hemi-continuity

A small portion of the population playing "non-optimally" may significantly change the incentives for other players causing a large shift in equilibrium behavior.

Quantal Response Equilibrium

(McKelvey and Palfrey)

propensity to play a strategy

$p_i(s_i) = \exp(\lambda_i u_i(s_i, \sigma_{-i}))$

$\sigma)i(s_i) = p_i(s_i)/\sum_{s_i{}'} p_i(s_i{}')$

as $\lambda_i \to \infty$ approaches best response

as $\lambda_i \to 0$ approaches uniform distribution

Smoothed Best Response Correspondence Example

	L ($\sigma_2(L) = q$)	R
U ($\sigma_1(U) = p$)	1,1	0,0
D	0,0	1,1

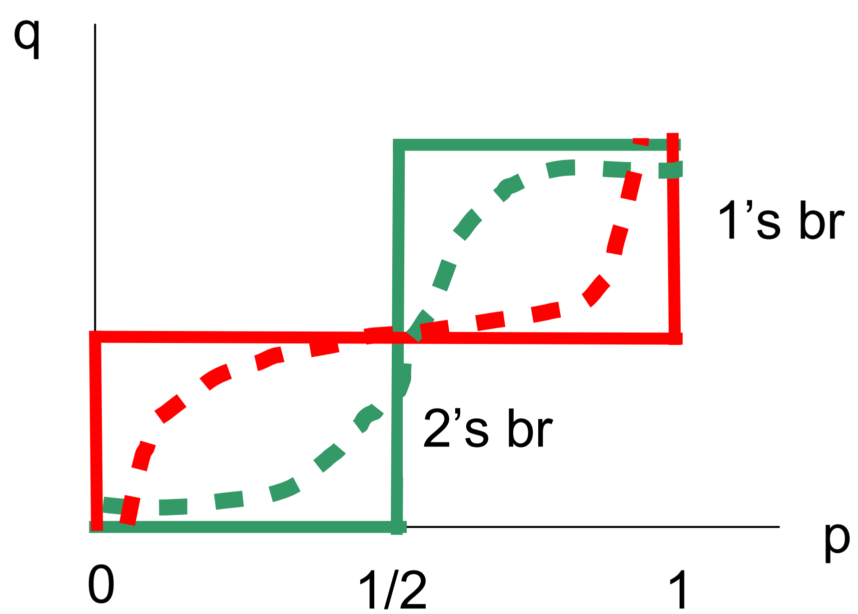

Voting: Individual Behavior

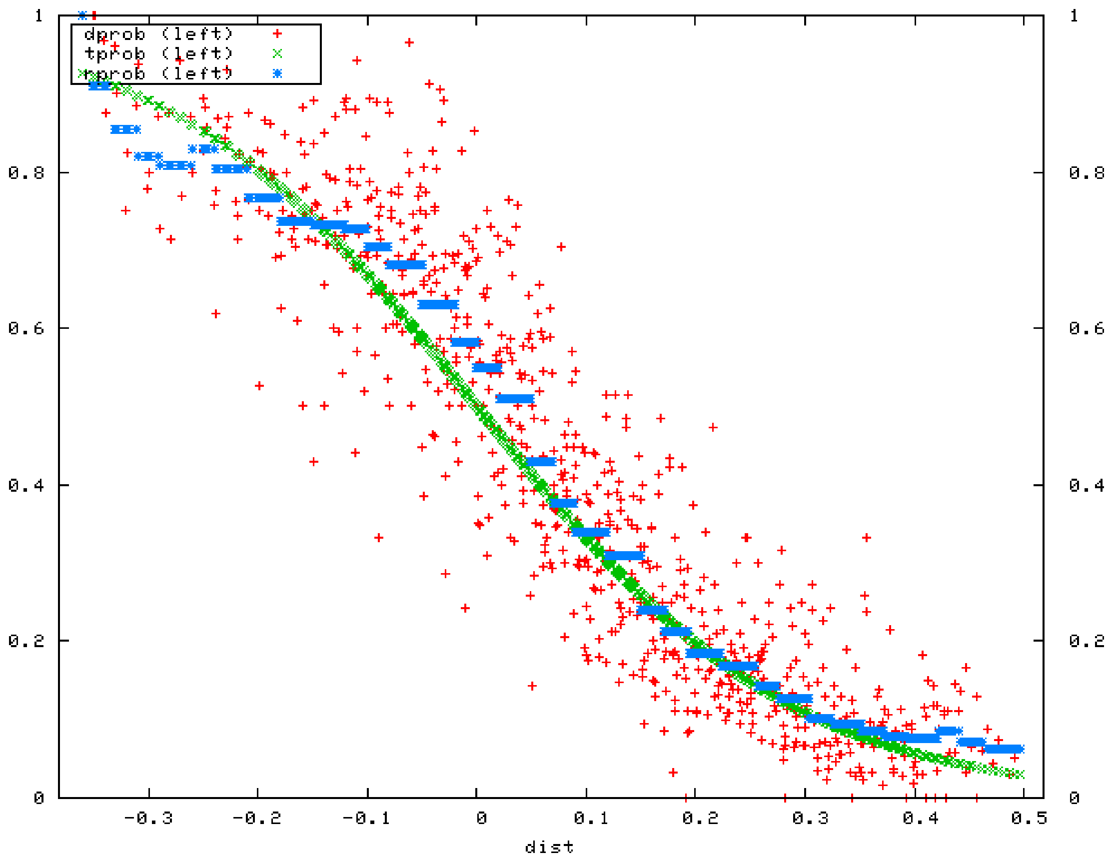

Observations

- contains an unknown preference parameter λ
- $\lambda = 0$ play is completely random
- as λ becomes large, the probability of playing the “best” response approaches one
- λ kind of index of rationality.
- in the voting experiment we can estimate a common value of λ for all players.
- corresponding equilibrium probabilities of play are given by the green curve
- does an excellent job of describing individual play
- it makes roughly the same predictions for aggregate play as Nash equilibrium

Limitations of QRE

- captures only the cost side of preferences
- recognizes – correctly – departures from standard “fully rational” selfish play are more likely if less costly in objective terms
- does not attempt to capture benefits of playing non-selfishly
- does not well capture, for example, the fact that under some circumstances players are altruistic, and in others spiteful.

Auctioning a Jar of Pennies

- surefire way to make some money
- put a bunch of pennies in a jar
- get together a group of friends
- auction off the jar of pennies
- with about thirty friends that you can sell a $3.00 jar of pennies for about $10.00

Winner's Curse

- friends all stare at the jar and try to guess how many pennies there are.
- Some under guess – they may guess that there are only 100 or 200 pennies. They bid low.
- Others over guess – they may guess that there are 1,000 pennies or more. They bid high.
- Of course those who overestimate the number of pennies by the most bid the highest – so you make out like a bandit.

Nash Equilibrium?

- According to Nash equilibrium this shouldn't happen
- Everyone should rationally realize that they will only win if they guess high
- they should bid less than their estimate of how many pennies there are in the jar
- they should bid a lot less – every player can guarantee they lose nothing by bidding nothing.
- in equilibrium, they can't on average lose anything, let alone $7.00.

QRE

- Recognize that there is small probability people aren't so rational
- Very different prediction
- some most possible profit anyone can make by getting the most number of pennies at zero cost: call this amount of utility U
- some least possible profit by getting a jar with no pennies at the highest possible bid: call this amount of utility u
- QRE says ratio of probability between two bids that give utility U, u is $\exp[\lambda(U-u)]$
- whatever is the difference in utility between two strategies it cannot be greater than that between U and u
- probability of highest possible bid is at least $p > 0$
- depends on how many bids are possible, not on how many bidders or their strategies

QRE with Many Bidders

- each bidder has at least a p probability of making the highest possible bid
- becomes a virtual certainty that one of the bidders will (unluckily for them) make this high bid

with enough bidders, QRE assures the seller a nice profit.