Economics 209B
Behavioral / Experimental Game Theory
(Spring 2008)

Lecture 4: Quantal Response Equilibrium (QRE)

Introduction

- Players do not choose best response with probability one (as in Nash equilibrium).

- Players choose responses with higher expected payoffs with higher probability – better response instead of best responses.

- Players have rational expectations and use the true mean error rate when interpreting others' actions.

- Modify Nash equilibrium to incorporate realistic limitations to rational choice modeling of games.

- Provide a statistical framework (structural econometric approach) to analyze game theoretic data (field and laboratory).

- If Nash had been a statistician, he might have discovered QRE rather then Nash equilibrium – Colin Camerer –

In practice, QRE often uses a logit or exponentiation payoff response function:

$$\mathsf{Pr}(a_i) = \frac{\mathsf{exp}[\lambda \sum_{a_{-i} \in A_{-i}} \mathsf{Pr}(a_{-i}) u_i(a_i, a_{-i})]}{\sum_{a'_i \in A_i} \mathsf{exp}[\lambda \sum_{a_{-i} \in A_{-i}} \mathsf{Pr}(a_{-i}) u_i(a'_i, a_{-i})]}.$$

The choice of action becomes purely random as $\lambda \rightarrow 0$, whereas the action with the higher expected payoff is chosen for sure as $\lambda \rightarrow \infty$.

- QRE does not abandon the notion of equilibrium, but instead replaces perfectly with imperfectly, or noisy, rational expectations.

- Players estimate expected payoffs in an unbiased way (expectations are correct, on average).

- As such, QRE provides a convenient statistical structure for estimation using either field or experimental data.

Normal-form games

Consider a finite n-player game in normal form:

- a set $N = \{1, ..., n\}$ of players,

- a strategy set $A_i = \{a_{i1}, ..., a_{iJ_i}\}$ consisting of J_i pure strategies for each player $i \in N$,

- a utility function $u_i : A \to \mathbb{R}$, where $A = \prod_{i \in N} A_i$ for every player $i \in N$.

Let Δ_i be the set of probability measures on A_i:

$$\Delta_i = \{(p_{i1}..., p_{iJ_i}) : \sum_{ij} p_{ij} = 1, p_{ij} \geq 0\}$$

where $p_{ij} = p_i(a_{ij})$.

The notation (a_{ij}, p_{-i}) represents the strategy profile where i adopts a_{ij} and all other players adopt their components of $p = (p_i, p_{-i})$.

A profile $p = (p_1, ..., p_n)$ is a Nash equilibrium if for all $i \in N$ and all $p_i' \in \Delta_i$

$$u_i(p) \geq u_i(p_i', p_{-i}).$$

Let $X_i = \mathbb{R}^{j_i}$ represent the space of possible payoffs for strategies that i can adopt and let $X = \prod_{i \in N} X_i$.

Then, define the function $\bar{u} : \mathbf{\Delta} \to X$ by

$$\bar{u}(p) = (\bar{u}_i(p), ..., \bar{u}_n(p)),$$

where

$$\bar{u}_{ij}(p) = u_i(a_{ij}, p_{-i}).$$

A quantal response equilibrium

A version of Nash equilibrium where each player's payoff for each action is subject to random error. Specifically:

[1] For each player i and each action $j \in \{1, ..., J_i\}$, and for any $p \in \Delta$, let

$$\hat{u}_{ij}(p) = \bar{u}_{ij}(p) + \epsilon_{ij}$$

where player i error vector $\epsilon_i = (\epsilon_{i1}, ..., \epsilon_{iJ_i})$ is distributed according to a joint PDF $f_i(\epsilon_i)$.

$f = (f_1, ..., f_n)$ is admissible if, for each i, the marginal distribution of f_i exists for each ϵ_{ij} and $\mathbb{E}(\epsilon_i) = 0$.

[2] For any $\bar{u} = (\bar{u}_1, ..., \bar{u}_n)$ with $\bar{u}_i \in \mathbb{R}^{j_i}$ for each i, define the ij-response set $\mathbf{R}_{ij} \subseteq \mathbb{R}^{j_i}$ by

$$\mathbf{R}_{ij}(\bar{u}_i) = \{\epsilon_i \in \mathbb{R}^{j_i} : \bar{u}_{ij}(p) + \epsilon_{ij} \geq \bar{u}_{ik}(p) + \epsilon_{ik} \forall k = 1, .., J_i\},$$

that is, given p, $\mathbf{R}_{ij}(\bar{u}_i(p))$ specifies the region of errors that will lead i to choose action j.

[3] Let the probability that player i will choose action j given $\bar{u}$ be equal

$$\sigma_{ij}(\bar{u}_i) = \int_{\mathbf{R}_{ij}(\bar{u}_i)} f(\epsilon) d\epsilon.$$

The function $\sigma_i : \mathbb{R}^{j_i} \longrightarrow \Delta^{J_i}$ is called the quantal response function (or statistical reaction function) of player i.

Let $G = \langle N, A, u \rangle$ be a normal form game, and let f be admissible. A QRE of G is any $\pi \in \Delta$ such that

$$\pi_{ij} = \sigma_{ij}(\bar{u}_i(\pi))$$

for all $i \in N$ and $1 \leq j \leq J_i$.

The quantal response functions

Properties of quantal response functions σ_{ij}:

[1] $\sigma \in \boldsymbol{\Delta}$ is non empty.

[2] σ_i is continuous in $\mathbb{R}^{j_i}$.

[1] and [2] imply that for any game G and for any admissible f, there exists a QRE.

[3] σ_{ij} is monotonically increasing in $\bar{u}_{ij}$.

[4] If, for each player i and every pair of actions $j, k = 1, .., J_i$, ϵ_{ij} and ϵ_{ik} are i.i.d., then

$$\bar{u}_{ij} \geq \bar{u}_{ik} \Longrightarrow \sigma_{ij}(\bar{u}) \geq \sigma_{ik}(\bar{u})$$

for all i and all $j, k = 1, .., J_i$.

[4] states that σ_i orders the probability of different actions by their expected payoffs.

A logit equilibrium

For any given $\lambda \geq 0$, the logistic quantal response function is defined, for $x_i \in \mathbb{R}^{j_i}$, by

$$\sigma_{ij}(x_i) = \frac{\exp(\lambda x_{ij})}{\sum_{k=1}^{J_i} \exp(\lambda x_{ik})},$$

and the QRE or logit equilibrium requires

$$\pi_{ij}(x_i) = \frac{\exp(\lambda \bar{u}_{ij}(\pi))}{\sum_{k=1}^{J_i} \exp(\lambda \bar{u}_{ik}(\pi))}$$

for each i and j.

Result I: Let σ be the logistic quantal response function; $\{\lambda_1, \lambda_2, ...\}$ be a sequence such that $\lim_{t\to\infty} \lambda_t = \infty$; $\{p_1, p_2, ...\}$ be a corresponding sequence with $p_t \in \pi^*(\lambda_t)$ for all t where

$$\pi^*(\lambda) = \left\{ \pi \in \Delta : \pi_{ij} = \frac{\exp(\lambda \bar{u}_{ij}(\pi))}{\sum_{k=1}^{J_i} \exp(\lambda \bar{u}_{ik}(\pi))} \forall i, j \right\}$$

is the logit correspondence.

Then, $p^* = \lim_{t\to\infty} p_t$ is a Nash equilibrium.

Proof: Assume p^* is not a Nash equilibrium. Then, there is some player $i \in N$ and some pair a_{ij} and a_{ik} with $p^*(a_{ik}) > 0$ and

$$u_i(a_{ij}, p^*_{-i}) > u_i(a_{ik}, p^*_{-i}) \text{ or } \bar{u}_{ij}(p^*) > \bar{u}_{ik}(p^*).$$

Since $\bar{u}$ is a continuous function, there exists some small ϵ and T, such that for all $t \geq T$,

$$\bar{u}_{ij}(p^t) > \bar{u}_{ik}(p^t) + \epsilon.$$

But as $t \to \infty$, $\sigma_k(\bar{u}_i(p^t))/\sigma_j(\bar{u}_i(p^t)) \to 0$ and thus $p^t(a_{ik}) \to 0$, which contradicts $p^*(a_{ik}) > 0$.

Result II: For almost any game G:

[1] $\pi^*(\lambda)$ is odd for almost all π.

[2] π^* is UHC.

[3] The graph of π^* contains a unique branch which starts at the centroid, for $\lambda = 0$, and converges the a unique NE, as $\lambda \to \infty$.

[3] implies that QRE defines a unique selection from the set of Nash equilibrium (the "tracing procedure" of Harsanyi and Selten, 1988).

Example I

Consider the game

	L	M	R
U	$1, 1$	$0, 0$	$1, 1$
M	$0, 0$	$0, 0$	$0, B$
D	$1, 1$	$A, 0$	$1, 1$

where $A > 0$ and $B > 0$.

The game has a unique THP (D, R), and the NE consists of all mixtures between U and D (resp. L and R) for player 1 (resp. 2).

The limit logit equilibrium selects $p = (\frac{1}{2}, 0, \frac{1}{2})$ and $q = (\frac{1}{2}, 0, \frac{1}{2})$ as the limit point.

QRE for example I with *A*=*B*=5

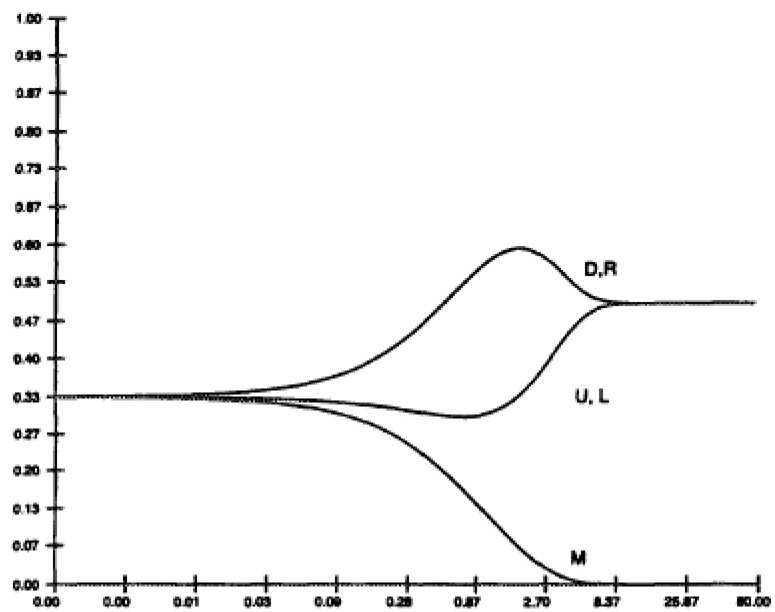

QRE for example I with *A*=*B*=100

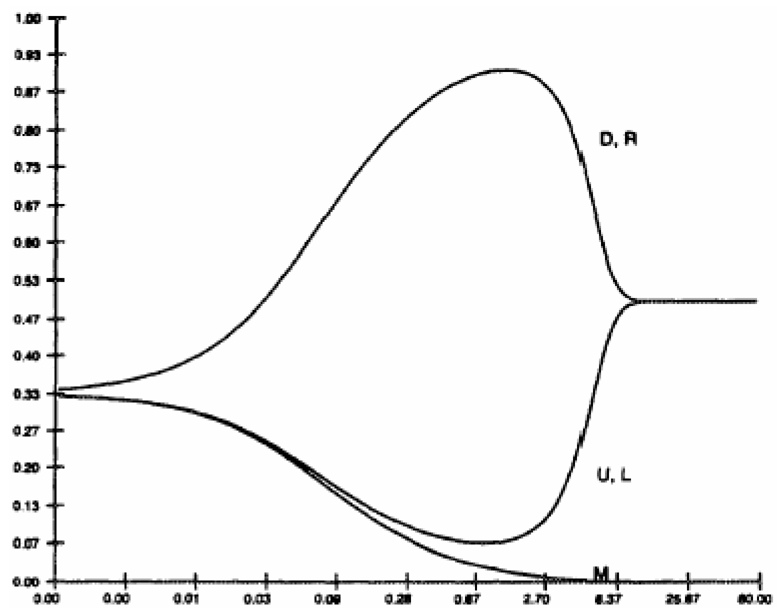

Example II

Consider the game

	R	L
T	$x, 1$	$1, 2$
B	$1, 2$	$2, 1$

All limit points are Nash equilibria but not all Nash equilibria are limit points (refinement). Computable in small finite games (Gambit).

QRE for example II

Properties of the QRE correspondence

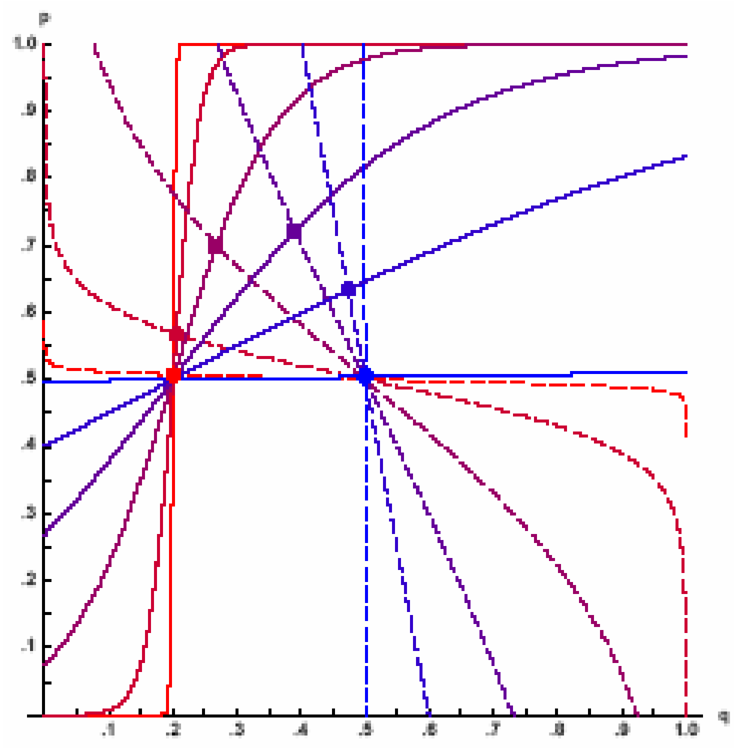

QRE for example II

Own-payoff Effects

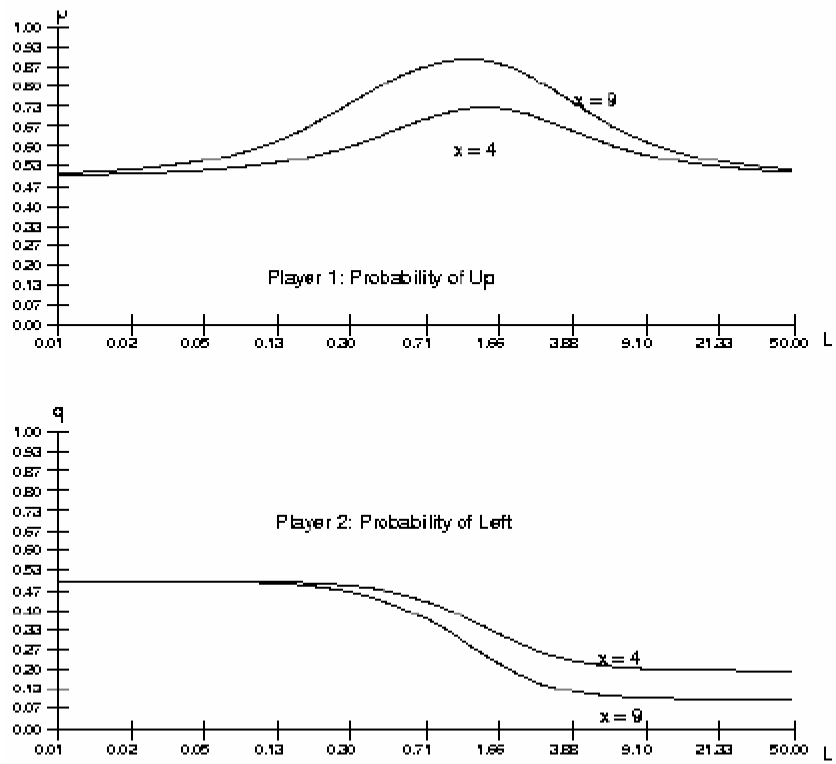

QRE for four-move centipede game

Extensive form

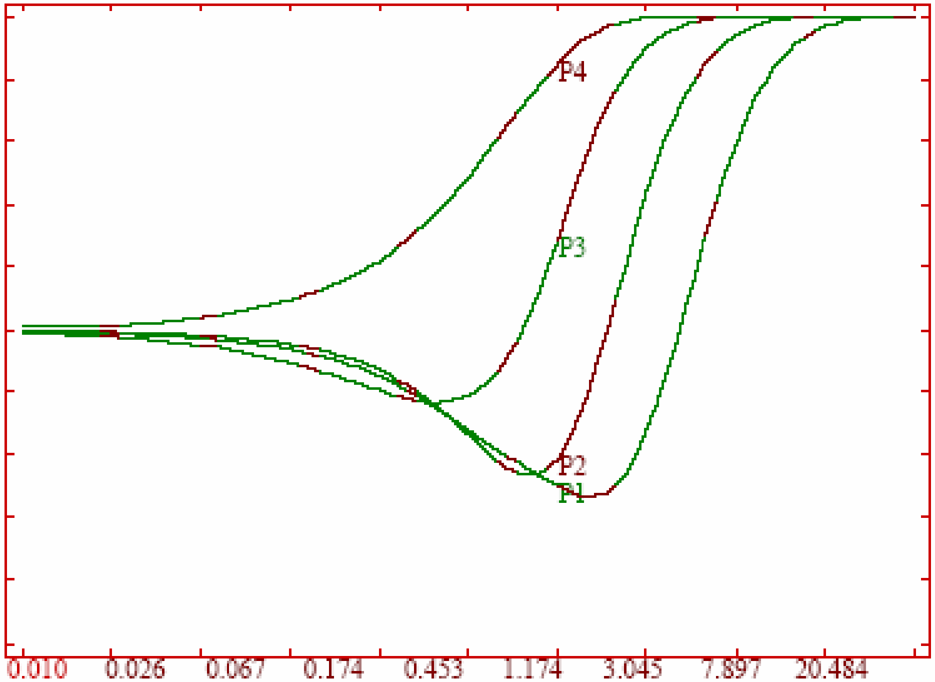

QRE for four-move centipede game

Normal form

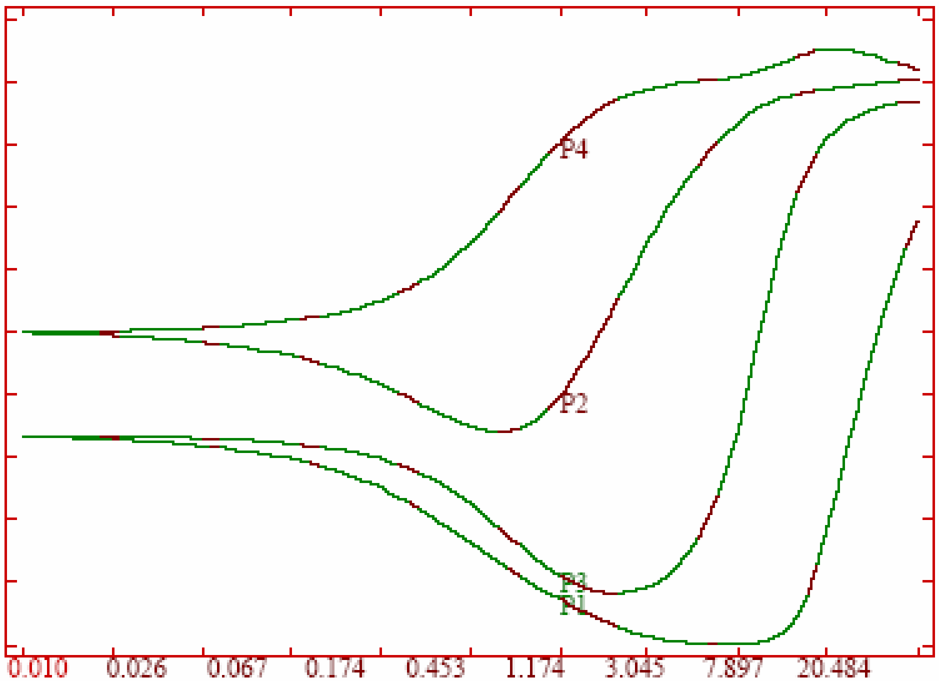

Relation to Bayesian equilibrium

In a Bayesian game (Harsanyi 1973), ϵ_i is viewed as a random disturbance to player i's payoff vector.

Suppose that for each $a \in A$, player i has a disturbance ϵ_{ij} added to $u_i(a_{ij}, a_{-i})$ and that each ϵ_{ij} is i.i.d. according to f.

Harsanyi (1973) assumes a separate disturbance $\epsilon_i(a)$ for i's payoff to each strategy profile $a \in A$, whereas here

$$\epsilon_i(a_i, a_{-i}) = \epsilon_i(a_i, a'_{-i})$$

for all i and all $a_{-i}, a'_{-i} \in A_{-i}$.

QRE inherits the properties of Bayesian equilibrium:

[1] An equilibrium exists.

[2] Best responses are “essentially unique” pure strategies.

[3] Every equilibrium is “essentially strong” and is essentially in pure strategies.

Data

Lieberman (1960)

	B_1	B_2	B_3
A_1	15	0	-2
A_2	0	15	-1
A_3	1	2	0

Ochs (1995)

	B_1	B_2
A_1	$1, 0$	$0, 1$
A_2	$0, 1$	$1, 0$

Game 1

	B_1	B_2
A_1	$9, 0$	$0, 1$
A_2	$0, 1$	$1, 0$

Game 2

	B_1	B_2
A_1	$4, 0$	$0, 1$
A_2	$0, 1$	$1, 0$

Game 3

QRE for Lieberman (1960)

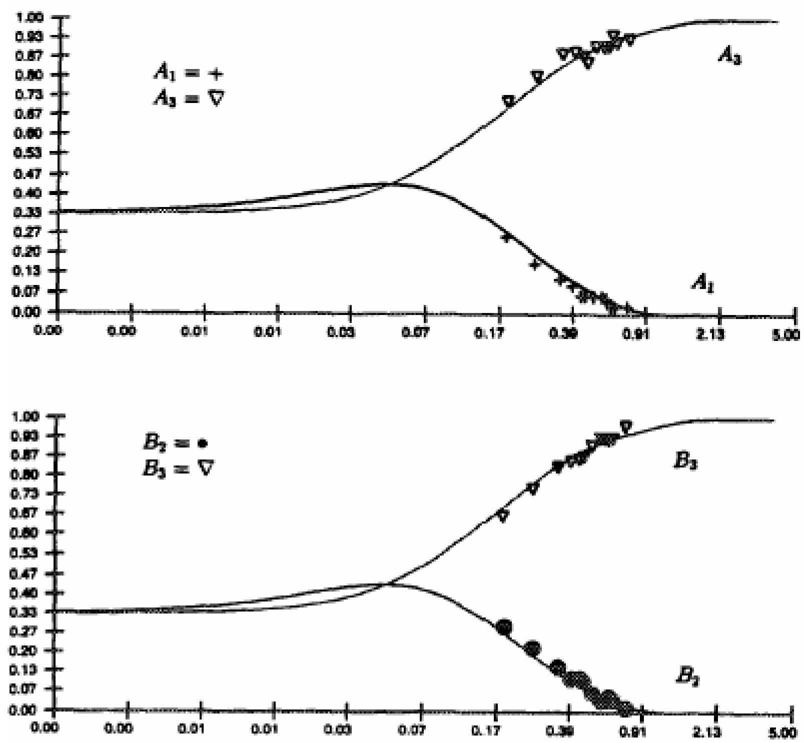

QRE for Ochs (1995)

Game 2

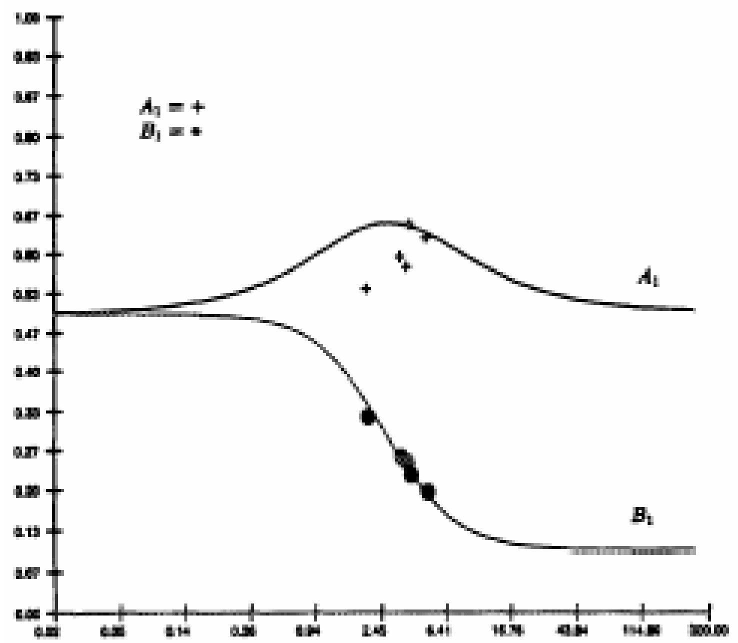

QRE for Ochs (1995)

Game 3