Title: Statistical Prediction of the Outcome of a Noncooperative Game

Date: Nov 13, 2008 04:00 PM

URL: http://pirsa.org/08110047

Abstract: Many statistics problems involve predicting the joint strategy that will be chosen by the players in a noncooperative game. Conventional game theory predicts that the joint strategy will satisfy an ``equilibrium concept\'\'. The relative probabilities of the joint strategies satisfying the equilibrium concept are not given, and all joint strategies that do not satisfy it are given probability zero. As an alternative, I view the prediction problem as one of statistical inference, where the ``data\'\' includes the details of the noncooperative game. This replaces conventional game theory\'s focus on how to specify a set of equilibrium joint strategies with a focus on how to specify a density function over joint strategies. I explore a Bayesian version of such a Predictive Game Theory (PGT) that provides a posterior density over joint strategies. It is based on the the entropic prior and on a likelihood that quantifies the rationalities of the players. The Quantal Response Equilibrium (QRE) is a popular game theory equilibrium concept parameterized by player rationalities. I show that for some games the local peaks of the posterior density over joint strategies approximate the associated QRE\'s, and derive the associated correction terms. I also discuss how to estimate parameters of the likelihood from observational data, and how to sample from the posterior. I end by showing how PGT can be used to specify a {it{unique}} equilibrium for any noncooperative game, thereby providing a solution to a long-standing problem of conventional game theory.

A PREDICTIVE THEORY OF GAMES

David H. Wolpert

NASA Ames Research Center

David.H.Wolpert@nasa.gov

http://ti.arc.nasa.gov/people/dhw/

NASA-ARC-05-097

ROADMAP

1) *Review statistics and game theory*

↓

2) *Apply statistics to games (as opposed to within games)*

↓

3) *Likelihood based on Quantal Response Eq.*

↓

4) *Ex.: Predicting airline behavior in bad weather*

↓

5) *New mathematical tools: rationality functions, cost of computation, varying numbers of players, etc.*

ONLY IDEA IN THIS TALK:

Human beings are physical objects*

**** - Physical object that are "goal-directed" though, whatever that means ...***

REVIEW OF GAME THEORY

- **N independent *players*, each with possible *moves*, $x_i \in X_i$**
- **Each i has a distribution $q_i(x_i)$; $q(x) = \prod_i q_i(x_i)$**
- **N *utility functions* $u^i(x)$; player i wants maximal $E_q(u^i)$**
- **$E_q(u^i)$ depends on q — but i only sets q_i**

***Equilibrium concept*: mapping from $\{u^i\} \rightarrow q$**

Strawman: Only equilibrium q can arise with humans.

"All we must do is find the right equilibrium concept."

REVIEW OF GAME THEORY - 2

Ex. 1: Nash Equilibrium (NE) q:

- **For all players i, $E_q(u^i)$ cannot rise by changing q_i**

Ex. 2: Quantal Response Equilibrium (QRE) q:

- **Simultaneously for all i, $q_i(x_i) \propto \exp[\beta_i E(u_i \mid x_i)]$**
- **Crude model of bounded rationality.**
- **In fair agreement with experiment.**
- **"Phase transitions" for finite systems.**

REVIEW OF GAME THEORY - 2

Ex. 1: Nash Equilibrium (NE) q:

- For all players i, $E_q(u^i)$ cannot rise by changing q_i

Ex. 2: Quantal Response Equilibrium (QRE) q:

- Simultaneously for all i, $q_i(x_i) \propto \exp[\beta_i E(u_i \mid x_i)]$
- Crude model of bounded rationality.
- In fair agreement with experiment.
- "Phase transitions" for finite systems.

REVIEW OF STATISTICS

1) Probability theory is the only consistent "calculus of uncertainty" for making predictions about physical world

2) In particular consistency forces *Bayes Theorem*:

$$P(\text{truth } z \mid \text{knowledge } \iota) \propto P(\iota \mid z)\, P(z)$$

3) Given a $P(z \mid \iota)$ and a *loss function* L(truth z, prediction y), the *Bayes-optimal* prediction is $\text{argmin}_y\, E_P[L(., y)]$ (Savage).

4) $\text{argmax}_z\, P(z \mid \iota)$ is an approximation; the *MAP* prediction

Probability theory to reason about physical objects.

Minimize expected loss to distill P(z) to a single z.

REVIEW OF STATISTICS

1) **Probability theory is the only consistent "calculus of uncertainty" for making predictions about physical world**

2) **In particular consistency forces *Bayes Theorem*:**

$$P(\text{truth } z \mid \text{knowledge } \iota) \propto P(\iota \mid z)\, P(z)$$

3) **Given a $P(z \mid \iota)$ and a *loss function* L(truth z, prediction y), the *Bayes-optimal* prediction is $\text{argmin}_y\ E_P[L(., y)]$ (Savage).**

4) **$\text{argmax}_z\ P(z \mid \iota)$ is an approximation; the *MAP* prediction**

Probability theory to reason about physical objects.

Minimize expected loss to distill P(z) to a single z.

EXAMPLE OF STATISTICS

1) Let the random variable we wish to predict itself be a probability distribution, $z = q(x)$.

2) Information theory tells us to use the *Entropic prior*

$$P(q) \propto exp[\alpha S(q)]$$

where $S(q)$ is the Shannon entropy of q, and $\alpha \in \Re^+$

3) Let the knowledge ι about q be $E_q(H) = h$ for some $H(x)$:

$$P(q \mid \iota) \propto exp[\alpha S(q)] \, \delta[E_q(H) - h]$$

REVIEW OF STATISTICS

1) Probability theory is the only consistent "calculus of uncertainty" for making predictions about physical world

2) In particular consistency forces *Bayes Theorem*:

$$P(\text{truth } z \mid \text{knowledge } \iota) \propto P(\iota \mid z)\, P(z)$$

3) Given a $P(z \mid \iota)$ and a *loss function* L(truth z, prediction y), the *Bayes-optimal* prediction is $\text{argmin}_y\, E_P[L(., y)]$ (Savage).

4) $\text{argmax}_z\, P(z \mid \iota)$ is an approximation; the *MAP* prediction

Probability theory to reason about physical objects.

Minimize expected loss to distill P(z) to a single z.

EXAMPLE OF STATISTICS

1) Let the random variable we wish to predict itself be a probability distribution, $z = q(x)$.

2) Information theory tells us to use the *Entropic prior*

$$P(q) \propto exp[\alpha\, S(q)]$$

where $S(q)$ is the Shannon entropy of q, and $\alpha \in \Re^+$

3) Let the knowledge ι about q be $E_q(H) = h$ for some $H(x)$:

$$P(q \mid \iota) \propto exp[\alpha\, S(q)]\, \delta[E_q(H) - h]$$

REVIEW OF STATISTICS

1) Probability theory is the only consistent "calculus of uncertainty" for making predictions about physical world

2) In particular consistency forces *Bayes Theorem*:

$$P(\text{truth } z \mid \text{knowledge } \iota) \propto P(\iota \mid z)\, P(z)$$

3) Given a $P(z \mid \iota)$ and a *loss function* L(truth z, prediction y), the *Bayes-optimal* prediction is $\text{argmin}_y\ E_P[L(., y)]$ (Savage).

4) $\text{argmax}_z\ P(z \mid \iota)$ is an approximation; the *MAP* prediction

Probability theory to reason about physical objects.

Minimize expected loss to distill P(z) to a single z.

EXAMPLE OF STATISTICS

1) Let the random variable we wish to predict itself be a probability distribution, $z = q(x)$.

2) Information theory tells us to use the *Entropic prior*

$$P(q) \propto \; exp[\alpha \, S(q)]$$

where $S(q)$ is the Shannon entropy of q, and $\alpha \in \Re^{+}$

3) Let the knowledge ι about q be $E_q(H) = h$ for some $H(x)$:

$$P(q \mid \iota) \;\propto\; exp[\alpha \, S(q)] \, \delta[E_q(H) - h]$$

REVIEW OF STATISTICS

1) **Probability theory is the only consistent "calculus of uncertainty" for making predictions about physical world**

2) **In particular consistency forces *Bayes Theorem*:**

   $$P(\text{truth } z \mid \text{knowledge } \iota) \propto P(\iota \mid z)\, P(z)$$

3) **Given a $P(z \mid \iota)$ and a *loss function* L(truth z, prediction y), the *Bayes-optimal* prediction is $\text{argmin}_y\, E_P[L(., y)]$ (Savage).**

4) **$\text{argmax}_z\, P(z \mid \iota)$ is an approximation; the *MAP* prediction**

Probability theory to reason about physical objects.

Minimize expected loss to distill P(z) to a single z.

EXAMPLE OF STATISTICS

1) Let the random variable we wish to predict itself be a probability distribution, $z = q(x)$.

2) Information theory tells us to use the *Entropic prior*

$$P(q) \propto exp[\alpha\, S(q)]$$

where $S(q)$ is the Shannon entropy of q, and $\alpha \in \Re^{+}$

3) Let the knowledge ι about q be $E_q(H) = h$ for some $H(x)$:

$$P(q \mid \iota) \propto exp[\alpha\, S(q)]\, \delta[E_q(H) - h]$$

REVIEW OF STATISTICS

1) Probability theory is the only consistent "calculus of uncertainty" for making predictions about physical world

2) In particular consistency forces *Bayes Theorem*:

$$P(\text{truth } z \mid \text{knowledge } \iota) \propto P(\iota \mid z)\, P(z)$$

3) Given a $P(z \mid \iota)$ and a *loss function* L(truth z, prediction y), the *Bayes-optimal* prediction is $\text{argmin}_y\, E_P[L(., y)]$ (Savage).

4) $\text{argmax}_z\, P(z \mid \iota)$ is an approximation; the *MAP* prediction

Probability theory to reason about physical objects.

Minimize expected loss to distill P(z) to a single z.

EXAMPLE OF STATISTICS

1) Let the random variable we wish to predict itself be a probability distribution, $z = q(x)$.

2) Information theory tells us to use the *Entropic prior*

$$P(q) \propto \; exp[\alpha\, S(q)]$$

where $S(q)$ is the Shannon entropy of q, and $\alpha \in \Re^{+}$

3) Let the knowledge ι about q be $E_q(H) = h$ for some $H(x)$:

$$P(q \mid \iota) \;\propto\; exp[\alpha\, S(q)]\; \delta[E_q(H) - h]$$

EXAMPLE OF STATISTICS - 2

4) So MAP q maximizes S(q') over the q' obeying $E_{q'}(H) = h$.

5) Let x be phase space position of a physical system with H(x) the Hamiltonian. The MAP q gives the Canonical Ensemble:

$$q(x) \propto exp[-\beta H(x)]$$

where β is a Lagrange parameter (it equals 1 / temperature)

6) If the numbers of particles of various types also varies stochastically, the MAP q is the Grand Canonical Ensemble.

EXAMPLE OF STATISTICS

1) Let the random variable we wish to predict itself be a probability distribution, $z = q(x)$.

2) Information theory tells us to use the *Entropic prior*

$$P(q) \propto \mathit{exp}[\alpha\, S(q)]$$

where $S(q)$ is the Shannon entropy of q, and $\alpha \in \Re^+$

3) Let the knowledge ι about q be $E_q(H) = h$ for some $H(x)$:

$$P(q \mid \iota) \propto \mathit{exp}[\alpha\, S(q)]\, \delta[E_q(H) - h]$$

EXAMPLE OF STATISTICS - 2

4) So MAP q maximizes S(q') over the q' obeying $E_{q'}(H) = h$.

5) Let x be phase space position of a physical system with H(x) the Hamiltonian. The MAP q gives the Canonical Ensemble:

$$q(x) \propto exp[-\beta H(x)]$$

where β is a Lagrange parameter (it equals 1 / temperature)

6) If the numbers of particles of various types also varies stochastically, the MAP q is the Grand Canonical Ensemble.

ROADMAP

1) *Review statistics and game theory*

↓

2) *Apply statistics* **to** *games (as opposed to within games)*

↓

3) *Likelihood based on Quantal Response Eq.*

↓

4) *Ex.: Predicting airline behavior in bad weather*

↓

5) *New mathematical tools: rationality functions, cost of computation, varying numbers of players, etc.*

ONLY IDEA IN THIS TALK:

Human beings are physical objects

I.e., you are a Scientist playing against Nature.

Whether Nature is

i) humans in a game, or

ii) interacting physical objects:

You should infer Nature's mixed strategy the same way.

GAME EQUILIBRIA

1) Humans are physical objects; to reason about the mixed strategy q of a game we *must* use a distribution P(q | ι):

Game theory equilibrium strawman is deficient

- N.b., bounded rationality automatic with distributions over q.

2) To distill P(q | ι) to a single q, can use a loss function L:

Predictive "Equilibrium" of a game meaningless without a loss function.

- L associated with the external scientist, *not* with the players.
- No need for refinements; equilibrium q is unique.

GAME EQUILIBRIA - 2

3) Alternative way to distill $P(q \mid \iota)$ to a single "mixed strategy":

$$P(x \mid \iota) = \int dq\, P(q \mid \iota) P(x \mid q, \iota) = \int dq\, P(q \mid \iota) q(x)$$

(I.e., average the q's)

4) Even if support of $P(q \mid \iota)$ restricted to NE q's, if $\exists$ multiple NE:

- **$P(x \mid \iota)$ is not a product distribution.**
  Players are independent under each NE q, but *to us* they appear coupled.
- **$P(x_i \mid \iota)$ is not best-response to $P(x_{-i} \mid \iota)$.**

GAME EQUILIBRIA

1) Humans are physical objects; to reason about the mixed strategy q of a game we *must* use a distribution $P(q \mid \iota)$:

Game theory equilibrium strawman is deficient

- N.b., bounded rationality automatic with distributions over q.

2) To distill $P(q \mid \iota)$ to a single q, can use a loss function L:

Predictive "Equilibrium" of a game meaningless without a loss function.

- L associated with the external scientist, *not* with the players.
- No need for refinements; equilibrium q is unique.

GAME EQUILIBRIA - 2

3) Alternative way to distill $P(q \mid \iota)$ to a single "mixed strategy":

$$P(x \mid \iota) = \int dq\, P(q \mid \iota) P(x \mid q, \iota) = \int dq\, P(q \mid \iota) q(x)$$

(I.e., average the q's)

4) Even if support of $P(q \mid \iota)$ restricted to NE q's, if ∃ multiple NE:

- $P(x \mid \iota)$ is not a product distribution.
  Players are independent under each NE q, but *to us* they appear coupled.
- $P(x_i \mid \iota)$ is not best-response to $P(x_{-i} \mid \iota)$.

GAME EQUILIBRIA

1) **Humans are physical objects; to reason about the mixed strategy q of a game we *must* use a distribution P(q | ι):**

Game theory equilibrium strawman is deficient

- **N.b., bounded rationality automatic with distributions over q.**

2) **To distill P(q | ι) to a single q, can use a loss function L:**

Predictive "Equilibrium" of a game meaningless without a loss function.

- **L associated with the external scientist, *not* with the players.**
- **No need for refinements; equilibrium q is unique.**

GAME EQUILIBRIA - 2

3) Alternative way to distill $P(q \mid \iota)$ to a single "mixed strategy":

$$P(x \mid \iota) = \int dq\, P(q \mid \iota) P(x \mid q, \iota) = \int dq\, P(q \mid \iota) q(x)$$

(I.e., average the q's)

4) Even if support of $P(q \mid \iota)$ restricted to NE q's, if $\exists$ multiple NE:

- **$P(x \mid \iota)$ is not a product distribution.**

  Players are independent under each NE q, but *to us* they appear coupled.
- **$P(x_i \mid \iota)$ is not best-response to $P(x_{-i} \mid \iota)$.**

THE POSTERIOR IN GAMES

1) So the "positive" problem is to infer the posterior

$$P(\text{joint mixed strategy } q \mid \{u^i : i = 1, N\})$$

2) Assume an entropic prior over q, $P(q) \propto exp[\alpha\, S(q)]$

3) How set the likelihood function $P(\iota \mid q)$ when "data" ι is the utility functions of the players?

EXAMPLE: LIKELIHOOD SAYS *q* IS A NE

	L	*R*
T	2	0
D	0	1

Row player move

Column player move | Player payoffs (identical)

3 NE, two shown in red

1) Entropic prior, P(q) ∝ *exp*[α S(q)]

2) P(ι | q) = 0 for non-NE q's, uniform over the three NE q's

3) So: P([L] × [T] | ι) = P([D] × [R] | ι) = 1 / [2 + w(α)]

P([(L, 2R)/3] × [(T, 2D)/3] | ι) = w(α) / [2 + w(α)]

where w(α) = *exp*[α {2ln[3] - (4/3)ln[2]}]

EXAMPLE: LIKELIHOOD SAYS q IS A NE

	L	*R*
T	2	0
D	0	1

Row player move

Column player move | Player payoffs (identical)

3 NE, two shown in red

1) Entropic prior, $P(q) \propto exp[\alpha S(q)]$

2) $P(\iota \mid q) = 0$ for non-NE q's, uniform over the three NE q's

3) So: $P([L] \times [T] \mid \iota) = P([D] \times [R] \mid \iota) = 1 / [2 + w(\alpha)]$

$$P([(L, 2R)/3] \times [(T, 2D)/3] \mid \iota) = w(\alpha) / [2 + w(\alpha)]$$

where $w(\alpha) = exp[\, \alpha \{2\ln[3] - (4/3)\ln[2]\}\,]$

LIKELIHOOD SAYS q IS A NE

	L	*R*
T	2	0
D	0	1

4) Recall $P(x \mid \iota) = \int dq\, P(q \mid \iota) q(x)$:

- **i) $P(T, L) = [9 + w(\alpha)] / [18 + 9w(\alpha)]$**
- **ii) $P(D, R) = [9 + 4w(\alpha)] / [18 + 9w(\alpha)]$**
- **iii) $P(D, L) = P(T, R) = 2w(\alpha) / [18 + 9w(\alpha)]$**

N.b., $P(D, R) > P(T, L)$, even though it's Pareto-inferior

5) Under this $P(x \mid \iota)$ player moves are coupled, and neither player's distribution is best-response to the other's.

6) Similarly, the likelihood

$P(\iota \mid q) = 0$ for non-QRE q's, uniform over the QRE q's

(for some particular values $\{\beta_1, \beta_2\}$)

gives a $P(q \mid \iota)$, which gives a (coupled) $P(x \mid \iota)$

ROADMAP

1) *Review statistics and game theory*

↓

2) *Apply statistics* to *games (as opposed to within games)*

↓

3) *Likelihood based on Quantal Response Eq.*

↓

4) *Ex.: Predicting airline behavior in bad weather*

↓

5) *New mathematical tools: rationality functions, cost of computation, varying numbers of players, etc.*

A MORE REALISTIC LIKELIHOOD

1) In real world, no q has *exactly* zero probability.

- How construct a likelihood allowing that?

2) Start with a QRE for a set of rationalities β, q^{β}

3) Utility theory says player i only cares about $E_q(u^i)$

4) So given q^{β}_{-i}, player i has no preference among q_i's obeying

$$E_{q_i, q^{\beta}_{-i}}(u^i) = E_{q^{\beta}_i, q^{\beta}_{-i}}(u^i) = E_{q^{\beta}}(u^i)$$

So allow all those q_i's in our new likelihood.

LIKELIHOOD SAYS q IS A NE

	L	*R*
T	2	0
D	0	1

4) Recall $P(x \mid \iota) = \int dq\, P(q \mid \iota) q(x)$:

   i) $P(T, L) = [9 + w(\alpha)] / [18 + 9w(\alpha)]$

   ii) $P(D, R) = [9 + 4w(\alpha)] / [18 + 9w(\alpha)]$

   iii) $P(D, L) = P(T, R) = 2w(\alpha) / [18 + 9w(\alpha)]$

   N.b., $P(D, R) > P(T, L)$, even though it's Pareto-inferior

5) Under this $P(x \mid \iota)$ player moves are coupled, and neither player's distribution is best-response to the other's.

6) Similarly, the likelihood

   $P(\iota \mid q) = 0$ for non-QRE q's, uniform over the QRE q's

   (for some particular values $\{\beta_1, \beta_2\}$)

   gives a $P(q \mid \iota)$, which gives a (coupled) $P(x \mid \iota)$

EXAMPLE: LIKELIHOOD SAYS *q* IS A NE

	L	*R*
T	2	0
D	0	1

Row player move

Column player move | Player payoffs (identical)

3 NE, two shown in red

1) Entropic prior, P(q) $\propto$ *exp*[α S(q)]

2) P(ι | q) = 0 for non-NE q's, uniform over the three NE q's

3) So: P([L] × [T] | ι) = P([D] × [R] | ι) = 1 / [2 + w(α)]

P([(L, 2R)/3] × [(T, 2D)/3] | ι) = w(α) / [2 + w(α)]

where w(α) = *exp*[α {2ln[3] - (4/3)ln[2]}]

LIKELIHOOD SAYS q IS A NE

	L	*R*
T	2	0
D	0	1

4) Recall $P(x \mid \iota) = \int dq\, P(q \mid \iota) q(x)$:

   i) $P(T, L) = [9 + w(\alpha)] / [18 + 9w(\alpha)]$

   ii) $P(D, R) = [9 + 4w(\alpha)] / [18 + 9w(\alpha)]$

   iii) $P(D, L) = P(T, R) = 2w(\alpha) / [18 + 9w(\alpha)]$

   N.b., $P(D, R) > P(T, L)$, even though it's Pareto-inferior

5) Under this $P(x \mid \iota)$ player moves are coupled, and neither player's distribution is best-response to the other's.

6) Similarly, the likelihood

   $P(\iota \mid q) = 0$ for non-QRE q's, uniform over the QRE q's

   (for some particular values $\{\beta_1, \beta_2\}$)

   gives a $P(q \mid \iota)$, which gives a (coupled) $P(x \mid \iota)$

ROADMAP

1) *Review statistics and game theory*

↓

2) *Apply statistics* to *games (as opposed to within games)*

↓

3) *Likelihood based on Quantal Response Eq.*

↓

4) *Ex.: Predicting airline behavior in bad weather*

↓

5) *New mathematical tools: rationality functions, cost of computation, varying numbers of players, etc.*

A MORE REALISTIC LIKELIHOOD

1) In real world, no q has *exactly* zero probability.

   - How construct a likelihood allowing that?

2) Start with a QRE for a set of rationalities β, q^{β}

3) Utility theory says player i only cares about $E_q(u^i)$

4) So given q^{β}_{-i}, player i has no preference among q_i's obeying

$$E_{q_i,q^{\beta}_{-i}}(u^i) = E_{q^{\beta}_i,q^{\beta}_{-i}}(u^i) = E_{q^{\beta}}(u^i)$$

So allow all those q_i's in our new likelihood.

A MORE REALISTIC LIKELIHOOD - 2

5) Do this even if $q_{-i} \neq q^{\beta}_{-i}$

6) So: Given *any* q_{-i}, the likelihood allows all q_i's such that

$$E_{q_i,q_{-i}}(u^i) = E_{q^{\beta}_i,q_{-i}}(u^i)$$

7) I.e., all q_i's are allowed which result in the same $E_{q_i,\, q_{-i}}(u^i)$ as

$$q^{\beta}_i(x_i) \propto \exp[\beta_i E_{q_{-i}}(u_i \mid x_i)]$$

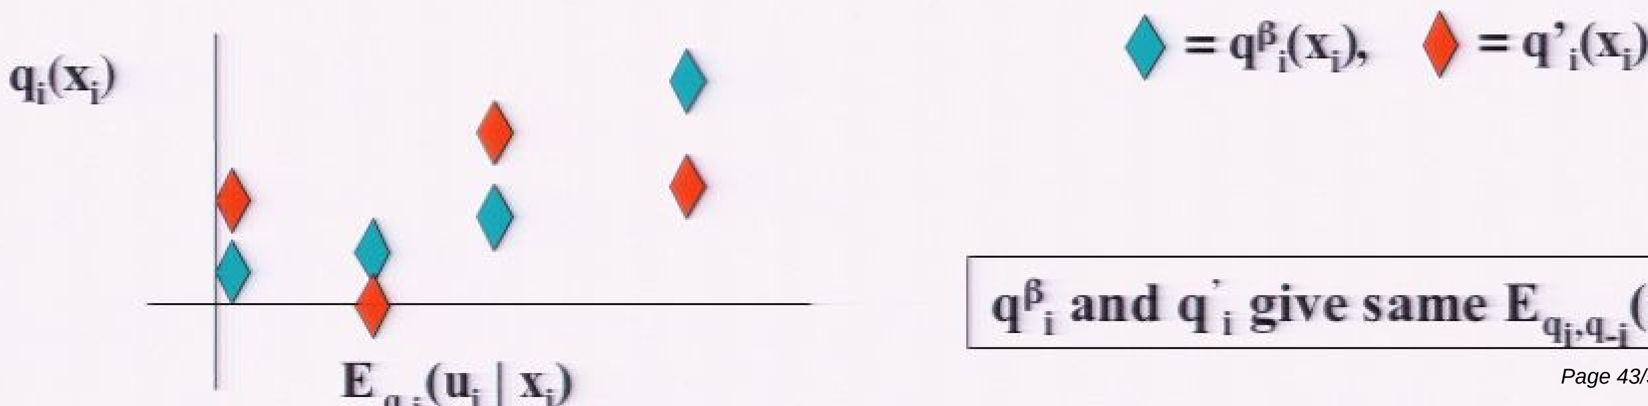

q^{β}_i and q'_i give same $E_{q_i,q_{-i}}(u^i)$

A MORE REALISTIC LIKELIHOOD - 3

8) β_i measures how high $E_{q_i,q_{-i}}(u^i)$ is for allowed q_i's, compared to how high it would be for unallowed q_i's

9) So β_i measures how "rational" i is.

10) Connection with Canonical Ensemble (CE) of physics:
    A smart player - high β, so low temperature - is *cold*.
    A dumb player is *hot*.

11) Of q allowed by $P(\iota \mid q)$, those locally maximizing $S(q)$ sometimes approximate the solution to the coupled equations

$$q_i(x_i) \propto exp[\beta_i U^i_{q_{-i}}(x_i)] \quad \forall i$$

- I.e., Quantal Response Eq. is sometimes a local approximation to maxent. (Can calculate correction terms to QRE.)

UNCERTAIN RATIONALITIES

1) This new likelihood allows infinitely more q's than the QRE - but still not all q's.

2) Solution: Note that in real world rationality uncertain.

3) Model this:

- Define $B_i(q) = B_i(q_i, q_{-i})$ as rationality value β_i of q_i given q_{-i}
- We know that if $B_i(q) = B_i(q')$, then

$$P(q \mid \iota) / P(q' \mid \iota) = exp[\alpha S(q)] / exp[\alpha S(q')]$$

- Require also that if $S(q) = S(q')$, then

$$P(q \mid \iota) / P(q' \mid \iota) = G(B_i(q)) / G(B_i(q'))$$

for some function G(.) (e.g., a hyperbolic tangent).

UNCERTAIN RATIONALITIES

4) So $$P(q \mid \iota) = \frac{e^{\alpha S(q)} G(B(q))}{\int dq \; e^{\alpha S(q)} G(B(q))}$$

- As desired, $P(q \mid \iota)$ is nowhere-zero (if G is positive)

- Therefore $$P(x \mid \iota) = \frac{\int dq \; q(x) e^{\alpha S(q)} G(B(q))}{\int dq \; e^{\alpha S(q)} G(B(q))}$$

- Use Monte Carlo to evaluate this.
  - In contrast, to find QRE must solve coupled set of equations
  - So PGT often scales far better than QRE.

ROADMAP

1) ***Review statistics and game theory***

↓

2) ***Apply statistics*** **to** ***games (as opposed to within games)***

↓

3) ***Likelihood based on Quantal Response Eq.***

↓

4) ***Ex.: Predicting airline behavior in bad weather***

↓

5) ***New mathematical tools: rationality functions, cost of computation, varying numbers of players, etc.***

Auction Scheme 1

- Pre-GDP: K flights over next T minutes.
- During GDP: reduced to K' flights over next T minutes.
  - Each airline affected by GDP submits as many bids (in $) as it had pre-GDP flights.
  - Following the order of bid sizes, airlines are allocated the earliest flight that fits their schedule.
- Airlines only pay their bids for allocated GDP flights.

Pre-GDP Schedule

Airline	Flight	ERTA	CTA
A	111	1:00	1:00
B	222	1:00	1:10
A	333	1:20	1:20
B	444	1:20	1:30
C	555	1:20	1:40

Ordered Bids

Airline	Bid
C	$15
B	$11
B	$10.99
A	$2
A	$1

Allocation Procedure

1. C has the highest bid and the earliest slot it can use is 1:20.
2. B is second and the earliest remaining slot it can use is 1:00.
3. B is third and the earliest remaining slot it can use is 1:40.

GDP Schedule

CTA*	Airline	Flight
1:00	B	222
1:20	C	555
1:40	B	444

LIKELIHOOD SAYS q IS A NE

	L	*R*
T	2	0
D	0	1

4) Recall $P(x \mid \iota) = \int dq\, P(q \mid \iota) q(x)$:

   i) $P(T, L) = [9 + w(\alpha)] / [18 + 9w(\alpha)]$

   ii) $P(D, R) = [9 + 4w(\alpha)] / [18 + 9w(\alpha)]$

   iii) $P(D, L) = P(T, R) = 2w(\alpha) / [18 + 9w(\alpha)]$

   N.b., $P(D, R) > P(T, L)$, even though it's Pareto-inferior

5) Under this $P(x \mid \iota)$ player moves are coupled, and neither player's distribution is best-response to the other's.

6) Similarly, the likelihood

   $P(\iota \mid q) = 0$ for non-QRE q's, uniform over the QRE q's (for some particular values $\{\beta_1, \beta_2\}$)

   gives a $P(q \mid \iota)$, which gives a (coupled) $P(x \mid \iota)$

EXAMPLE: LIKELIHOOD SAYS q IS A NE

	L	*R*
T	2	0
D	0	1

Row player move

Column player move — Player payoffs (identical)

3 NE, two shown in red

1) Entropic prior, $P(q) \propto exp[\alpha\, S(q)]$

2) $P(\iota \mid q) = 0$ for non-NE q's, uniform over the three NE q's

3) So: $P([L] \times [T] \mid \iota) \;=\; P([D] \times [R] \mid \iota) \;=\; 1 / [2 + w(\alpha)]$

$$P([(L, 2R)/3] \times [(T, 2D)/3] \mid \iota) \;=\; w(\alpha) / [2 + w(\alpha)]$$

where $w(\alpha) = exp[\, \alpha \,\{2\ln[3] - (4/3)\ln[2]\}\,]$