Calibration and Nash Equilibrium

Dean Foster
University of Pennsylvania

and

Sham Kakade
TTI

September 22, 2005

What is a Nash equilibrium?

- Cartoon definition of NE:
  - Leroy Lockhorn: "I'm drinking because she is driving."
  - Loretta Lockhorn: "I'm driving because he is drinking."
- Technical definition of NE:
  - If everyone else will play the Nash equilibrium, then I should play it also.
  - Holds for all players in a game.
- Equilibrium of what process?

Calibration: A form of unbiasedness

"Suppose in a long (conceptually infinite) sequence of weather forecasts, we look at all those days for which the forecast probability of precipitation was, say, close to some given value p and then determine the long run proportion f of such days on which the forecast event (rain) in fact occurred. If $f = p$ the forecaster may be termed well calibrated." Dawid [1982]

A minimal condition for performance

- On sequence: 0 1 0 1 0 1 0 ...
- A constant forecast of .5 is calibrated
- A constant forecast of .6 is not calibrated

Calibration: A form of unbiasedness

- see handout

- Left graph
  - Bridge players
  - Forecasts of winning a contract that was just bid.
  - Expert bridge players are more calibrated than beginners
  - Note: some experts play hands with 0 chance of winning!

- Right graph
  - College students
  - sports is more about utility than about probability. (I want my team to win.)

Bridge players

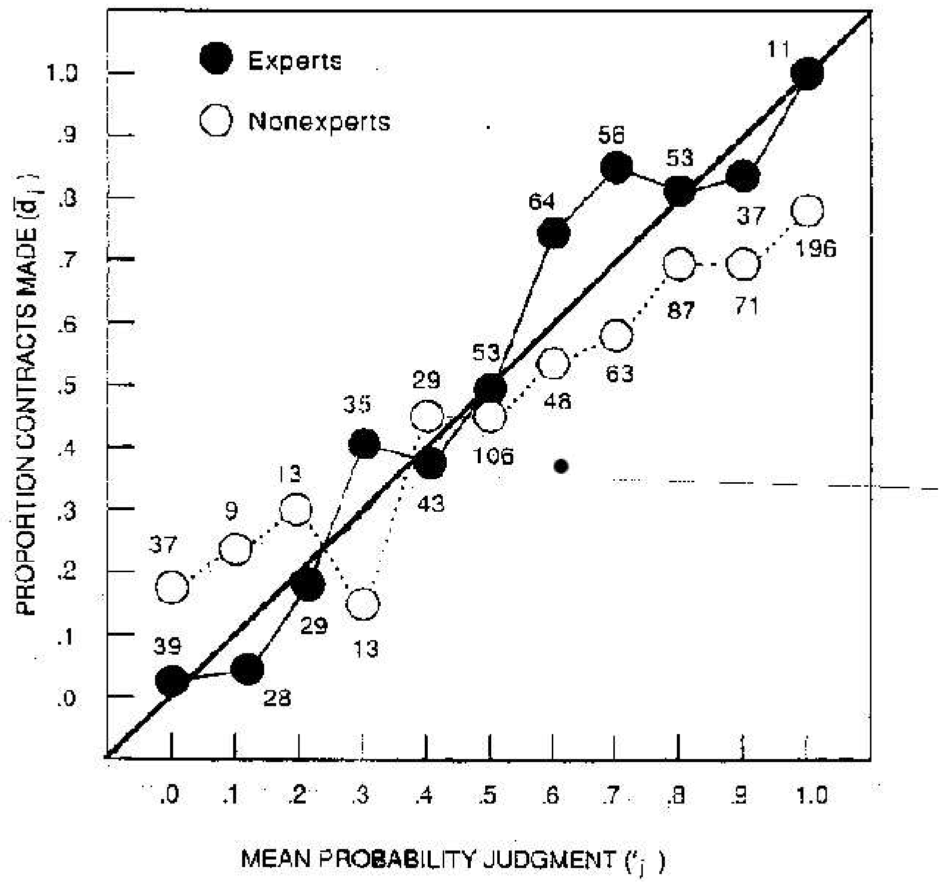

College students

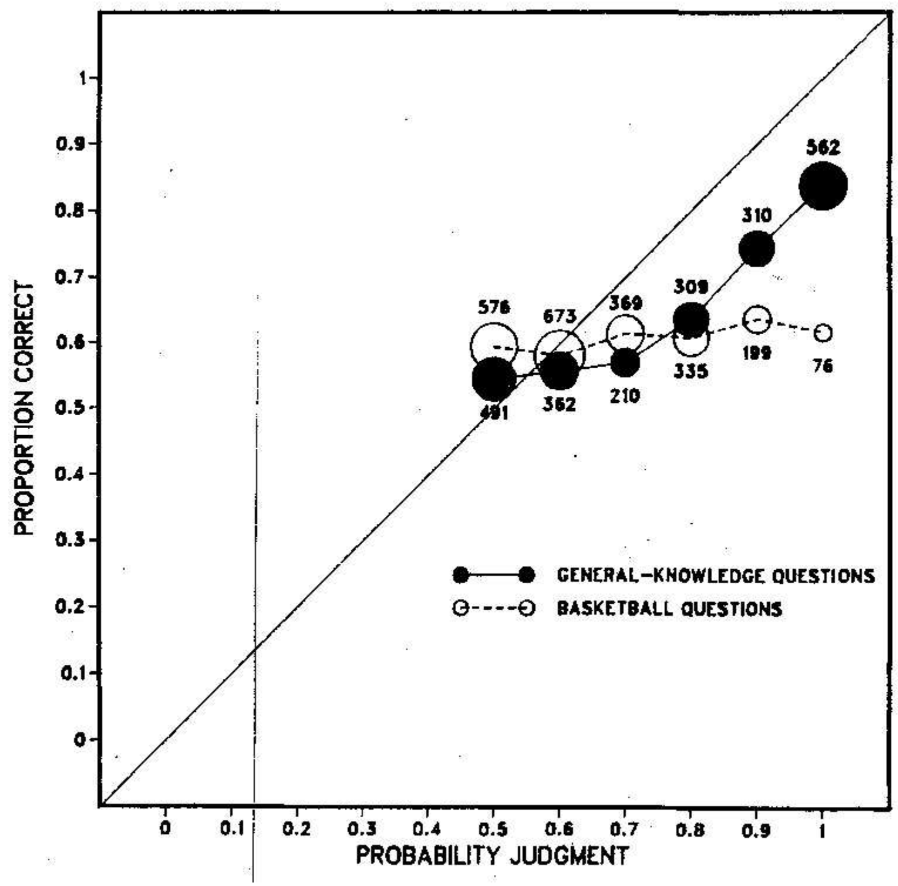

Learning in games

- Learning models for games:
  - Two players repeatedly play a game
  - Each views the sequence of the other person's plays as data
  - Each predicts what the other play will do
  - Each then plays a best response to the prediction
- We will discuss the equilbrium resulting from calibrated learning models

Traditional test functions for calibration

- Sequential prediction (t is time)

- X_t is forecast by p_t

- Traditional calibration, means

$$\frac{1}{T}\sum_{t=1}^{T}(X_t - p_t)\,w(p_t) \to 0$$

  holds for all possible $w()$.

  - Note: The class of $w()$ can be restricted to indicator functions.

  - Oakes proved without randomization, calibration is impossible.

  - With randomization calibration is possible.

New test functions

- X_t sequence to be forecast by p_t

- *Weak calibration*, means

$$\frac{1}{T}\sum_{t=1}^{T}(X_t - p_t)\, w(p_t) \to 0$$

  for all $w()$ which are continuous function.

Achieving weak calibration via polynomial regression

Algorithm:

- Fit the model

$$X_t = \sum_{i=0}^{d} \beta_i \, p_t^i + \text{noise}$$

  on $X_1, \ldots, X_{T-1}$ to estimate the β's.

- Solve fixed point equation:

$$p_T = \sum_{i=0}^{d} \beta_i p_T^i$$

  (If no solution exists, use arbitary rule, say $p_T = 0.5$.)

- Use p_T to forecast X_T.

Theorem: p_T is approximately weakly calibrated.

Algorithm: Solve the fixed point equation

$$p_T = \sum_{i=0}^{d} \widehat{\beta}_i p_T^i$$

where the $\widehat{\beta}$'s are determined by a polynomial regression of $X_1, \ldots, X_{T-1}$ on $p_1, \ldots, p_{T-1}$.

Theorem: p_T is approximately weakly calibrated.

Proof:

- Lemma (1991): regression does as well as any linear combination.
- Thus p_T will predict as well as any polynomial of p_T.
- Hence no polynomial change of p_T will do better.

Trivia: I talked about this lemma the last time I was here (1988).

Games as a good application for paranoid data analysis

- Learning in games has extensive literature
- Both emprical and theoretical
- Two players repeatedly play a game
- Do they converge to playing an equilibrium?
- Typical learning setup:
  - Player i uses $p_{i,t}$ to predict other's play at the round t
  - Player i computes best response distribution $s_i(p_{i,t})$
  - Player i then randomly action S_i from this distribution

Individual vs Public calibration

- Game setting for calibration
  - $X_{i,t}$ is the observable that player i cares about at time t
  - $p_{i,t}$ is a forecast of $X_{i,t}$
- Individual calibration:

$$(\forall i) \qquad \frac{1}{T}\sum_{t=1}^{T}(X_{i,t} - p_{i,t})\, w(p_{i,t}) \to 0$$

- Public calibration:

$$(\forall i) \qquad \frac{1}{T}\sum_{t=1}^{T}(X_{i,t} - p_{i,t})\, w(\vec{p}_t) \to 0$$

Sharp vs. smooth best response

- $s_i(p_{i,t})$ is the distrubtion player i will use for making a play at time t.

- Sharp best response means s_i maps to corners of simplex
  - Used in orginal research on learning
  - requires randomized forecasts to get convergence results
  - Obviously $p_{i,t}$ must be protected from being leaked

- Smooth best reply restricts $s_i(\cdot)$ to be Lipschitz
  - Only close to optimal
  - Randomization is now in the play

Observables

- Game setup:
  - Take $X_i = S_{-i}$ (i.e. all actions but player i)
  - $p_{i,t}$ is forecast of $X_{i,t}$

- Individual calibration:

$$(\forall i) \qquad \frac{1}{T}\sum_{t=1}^{T}(X_{i,t} - p_{i,t})\, w(p_{i,t}) \to 0$$

- Public calibration:

$$(\forall i) \qquad \frac{1}{T}\sum_{t=1}^{T}(X_{i,t} - p_{i,t})\, w(\vec{p}_t) \to 0$$

Convergence

- Suppose players play a smooth best reply to forecast $p_{i,t}$.
  - Traditional calibration $\rightarrow$ correlated equilibria
  - Public calibration $\rightarrow$ Nash equilibria

- Speed of convergence is related to dimension of the "Hilbert space" of the testing functions
  - For individual: dimension $(1/\epsilon)^{a^n}$
  - For public: dimension is $(1/\epsilon)^{na^n}$
  - Hence convergence is slow in both cases.

- Need lower dimensional space, but what can be changed?

Proof: Public calibration converges to NE

- Truth $\approx$ prediction
  - via calibration

- Truth is independent
  - Given $\vec{p}$ each player is in fact playing independently

- ϵ-rationality
  - ϵ-BR to prediction
  - p_i includes information about what all other players will do

- Independence + ϵ-rationality = ϵ-NE.

Utility estimation

- Take $X_{i,t}$ to be the vector of potential payoffs
  - $\vec{S}_{-i}$ is the vector of everyone else's play
  - $u_{i,t}(k) = u_i(k, \vec{S}_{-i,t})$
  - $X_{i,t} = (u_{i,t}(1), \ldots, u_{i,t}(a))$
- Calibration of utilities $\rightarrow$ correlated equilibria
- Public calibration of utilities $\rightarrow$ Nash equilibria

Conclusion: Today's talk in historical context

Method	**Forecast probability**	**Forecast utility**
Least squares (F. '91)	doesn't converge	doesn't converge
Blackwell Approach-ability	**CE** Calibration (F. and Vohra, '97)	**CE** No regret (F. and Vohra '97) (Hart and Mas-Colell '00)
Exhaustive search	**NE** Hypothesis testing (F. and Young '03)	**NE** Regret testing (F. and Young '05) (Germano & Lugosi '05)
Public methods	**NE** Weak calibration (Kakade and F. '04)	**NE** Weak utility estimation (Kakade and F. '05)