Regression, Calibration and Nash Equilibrium

Dean Foster
University of Pennsylvania

and

Sham Kakade
TTI

May 30, 2006

Three related questions

- How might people come to play a Nash equilibria?
- Can one be calibrated?
- Why should Least squares regression predict the future?

Nash equilibrium: Cartoon definition

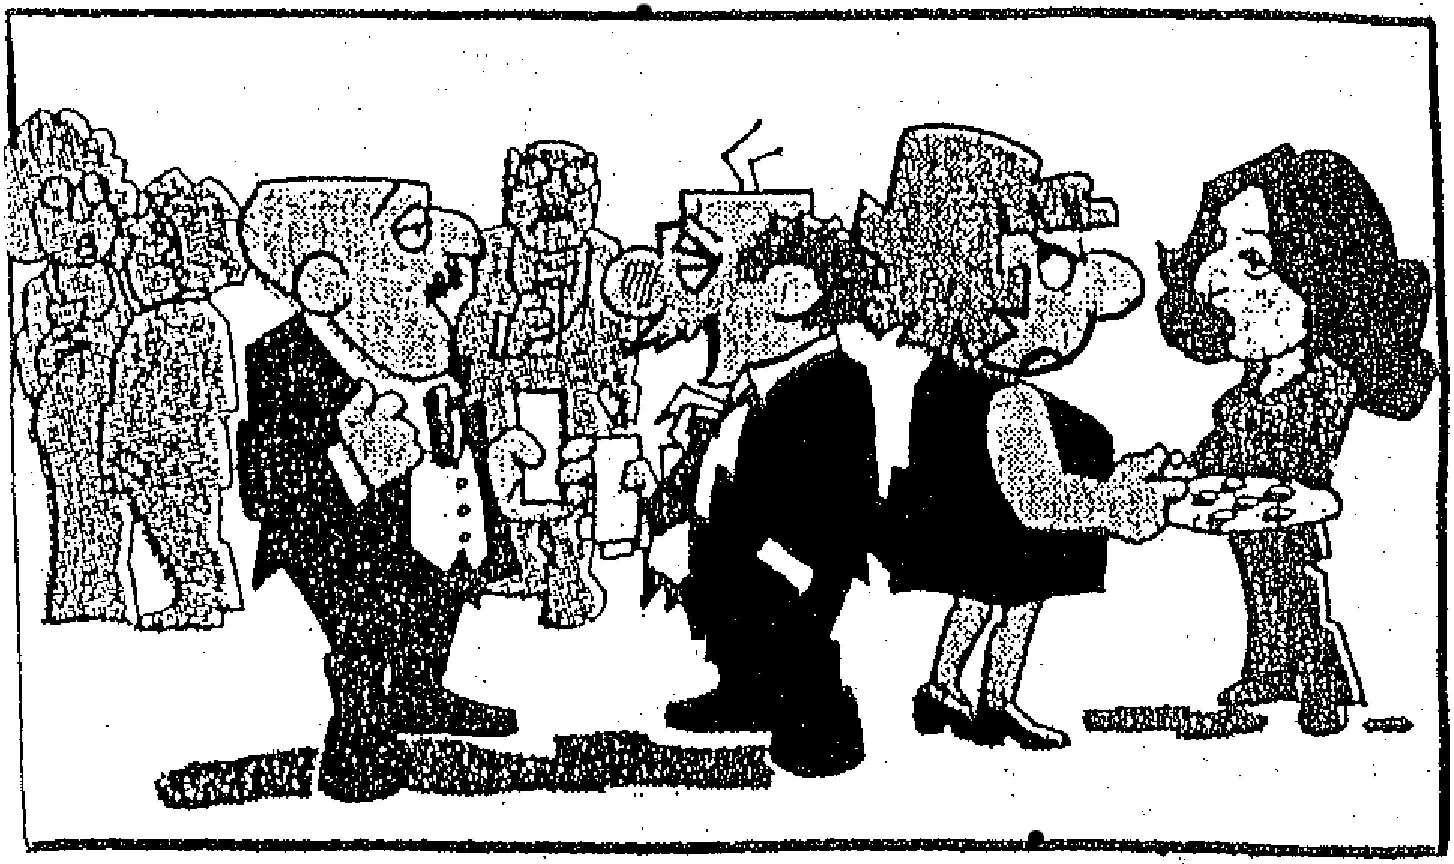

"LORETTA'S DRIVING BECAUSE I'M DRINKING, AND I'M DRINKING BECAUSE SHE'S DRIVING."

What is a Nash equilibrium? (1950)

- 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.

Question: Equilibrium of what process?

Least squares regression (1801)

- In 1801 Ceres (an asteroid) was found for 3 month and then lost
- On reading this, Gauss used the 24 observations to predict where it could be found.
- The search area is quadratic in accuracy
- Hence least squares regression was born

Question: Why should least squares on past data predict the future?

Calibration: A form of unbiasedness (1992)

"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

Question: Can calibration be guaranteed?

Bridge players

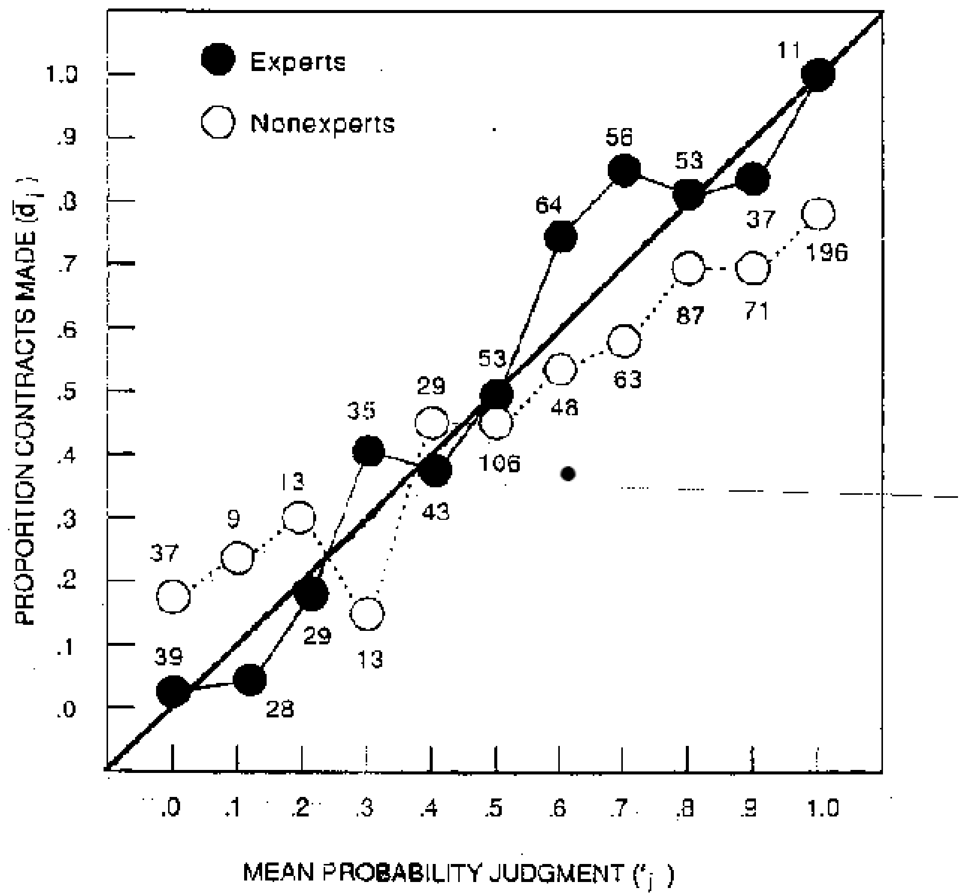

College students

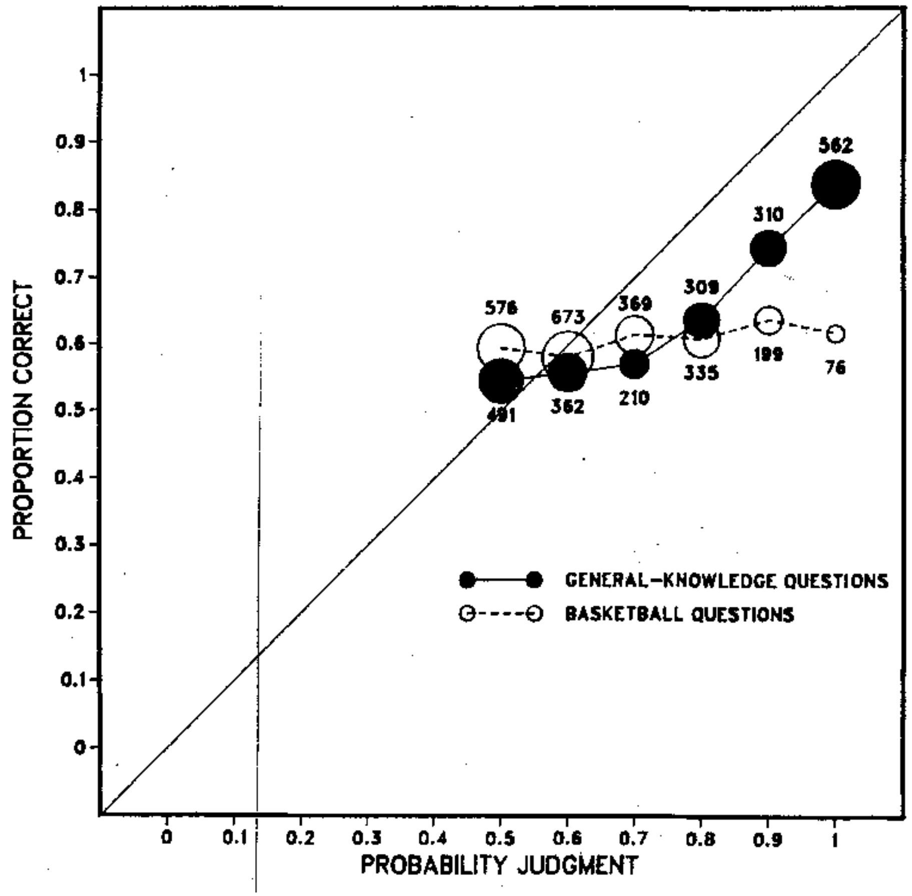

Answering the regression question

- Predict y_t by $\widehat{y}_t = \theta_t \cdot x_t$, using ridge regression.

$$\theta_t = \arg\min_{\theta} \sum_{i=1}^{t-1} \|\theta \cdot x_i - y_i\|^2 + \|\theta\|^2$$

(Clip predictions if outside range of possible values)

- Guarantee: for all θ and all sequences

$$\sum_{t=1}^{T} \|\widehat{y}_t - y_t\|^2 \leq \sum_{t=1}^{T} \|\theta \cdot x_t - y_t\|^2 + \|\theta\|^2 + \frac{d}{2} \log T$$

(Foster '91, Vovk '01, Azoury & Warmuth '01)

- The "Hannan regret" is sublinear

Goals for prediction

- Regret (Hannan)
  - compare our performance to the best predictor in some class
  - relative to given class
- Bias (Calibration)
  - check to see if our predictions are unbiased
  - relative to given class of checks
- This talk: Hannan regret to prove calibration

Answering the calibration question

- Sequential prediction: y_t is forecast by $\widehat{y}_t$

- Calibration, can be interpreted as saying that
$$\frac{1}{T}\sum_{t=1}^{T}(y_t - \widehat{y}_t)\, w(\widehat{y}_t) \to 0$$
for all $w()$.

  – **No!** Oakes (1985) without randomization. $(w(x) = I_{x \geq \frac{1}{2}} - \frac{1}{2})$

  – **Yes!** Foster and Vohra (~~1992~~ 1997) with randomization.

Relationship of normal equations to prediction

- For each regression variable, x_i

$$\sum_{t=1}^{T} x_{i,t}(y_t - \widehat{y}_t) \to 0$$

(Proof: otherwise $\widehat{y} + \epsilon x_i$ would be a better forecast than $\widehat{y}$.)

- Key idea: what if $x_{i,t} = w_i(\widehat{y}_t)$? Then

$$\frac{1}{T}\sum_{t=1}^{T} w_i(\widehat{y}_t)(y_t - \widehat{y}_t) \to 0$$

  - We are calibrated against this $w_i()$.
  - Need to be calibrated against many $w_i()$'s (call them $\vec{w}()$).
  - Fixed point: $\widehat{y}_t = \theta_{t-1} \cdot \vec{w}(\widehat{y}_t)$.

Weaker test functions

- All bounded functions is not a nice space
  - Too rich (i.e. Uncountable-dimensional Banach space.)
  - No fixed points

- Alternative: Use continuous functions with basis of continuous functions.

- New definition: *Weak calibration*, means

$$\frac{1}{T}\sum_{t=1}^{T}(y_t - \widehat{y}_t)\, w(\widehat{y}_t) \to 0$$

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

Algorithm for getting calibration from regression

- Pick family of continuous functions $w_i()$ which form a basis for the L2 bounded functions.
- Use $w_i(\widehat{y}_t)$ as regression variables
- use regression to fit historical data
- Solve fixed point equation for forecast next round

Theorem 1 (Foster and Kakade 2005) *The above procedure will be weakly calibrated.*

The Nash question: 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 equilibrium resulting from calibrated learning models

Games as a good application for paranoid data analysis

- 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 picks action S_i from this distribution
  - All players observe play vector and update their forecasts

First answer to Nash question

Theorem 2 (Foster and Vohra 1998) *If all players are using a calibrated forecast, then the empirical distribution of plays will be close to a correlated equilibrium.*

- CE is not as narrow as Nash equilibrium
- But basically on target
- Some even view the CE concept as more natural than a NE. (Roger Meyerson states it as "2 out of 3 intelligent species discover CE's first.")

Second answer to Nash question

Theorem 3 (Foster and Kakade 2005) *If all players play an ϵ-smooth response to the* same *calibrated forecast then each round play will be close to an ϵ-Nash equilibrium.*

- Stronger than first answer since it holds for each round.
- Requires coordination on the forecast.
- Can let ϵ slowly go to zero.

Proof: Public calibration converges to NE

- Truth $\approx$ prediction (by definition of calibration)
- ϵ-rationality
  - ϵ-BR to prediction
  - predictions close to truth
- Truth is independent (by construction of strategies)
- Independence $+$ ϵ-rationality $=$ ϵ-NE.

Speed of convergence

- Speed of convergence is related to dimension of the 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 using calibration
- Need lower dimensional space
- Answer: predict utility of each action
- Called "No-internal regret"
- In individual calibration, serious speed-up, not so here

History of equilibria in games

Method	Forecast probability	Forecast utility
Evolution	N.A.	Almost ($\rightarrow$ ESS)
Hannan regret	N.A.	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. '05)	**NE** Weak utility estimation (Kakade and F. '05)

Summary

- How might people come to play a Nash equilibria?
  - Coordinate on a calibrated forecast
  - Play best smooth reply to forecast
- Can one be calibrated?
  - Write calibration as normal equation
  - Use regression to get normal equations to converge to zero
- Why should Least squares regression predict the future?
  - It is Hannan consistent
  - individual sequences